# Deep Generative Model in Machine Learning: Theory, Principle and Efficacy（2nd Workshop）

Website: https://delta-workshop.github.io/DeLTa2026/

## 1. Workshop Summary

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 autoregressive models [4] and diffusion models [5,6] 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.

### Motivation and Problem Statement

The continued advancement of DGMs will benefit from stronger theoretical foundations, which can provide deeper insight into their behavior and limitations. Current frameworks, however, fall short 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.

At the same time, several algorithmic challenges continue to limit the broader adoption and reliability of DGMs:

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.

### Goals and Themes

Building on these challenges, the 2nd **D**eep G**e**nerative Model in Machine **L**earning: **T**heories, Principles and Effic**a**cy (**DeLTa 2026**) workshop will convene researchers to advance discussion around two overarching questions:

***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?*

Following the success of the inaugural first DeLTa Workshop at ICLR 2025, which attracted over 400 participants and strong interdisciplinary engagement, this second edition aims to both deepen theoretical inquiries and broaden accessibility. In particular, we continue to introduce a dedicated tiny/short paper track, enabling contributions that present preliminary results, negative findings, or novel perspectives that may not fit into traditional long-paper formats. This format particularly supports early-career and under-resourced researchers, aligning with ICLR's ongoing commitment to diversity and inclusivity.

## Uniqueness and Urgency

Unlike prior workshops focused on specific DGM families or applications, DeLTa 2026 Workshop bridges theory and practice by emphasizing fundamental principles that guide algorithmic innovation. It serves as a forum for cross-pollination between theoretical, algorithmic, and applied research communities, ultimately aiming to establish robust, interpretable, and transferable generative frameworks.

The exponential growth of DGMs—in both scale and complexity—has outpaced our theoretical understanding, creating critical needs in stability, scalability, and interpretability. The field now stands at an inflection point: without a solid theoretical foundation, progress risks becoming empirical and ad hoc. DeLTa 2026 directly addresses this urgency by catalyzing rigorous discussion and collaboration toward principled theories that can sustain the next decade of generative AI innovation.

## Organizing Team and Expected Impact

Our organizing team comprises researchers with complementary expertise across generative modeling theory, scalable learning systems, and geometry-aware deep learning, with strong representation across continents and institutions. With confirmed commitments from five distinguished invited speakers, we are confident that the DeLTa 2026 will serve as both a catalyst for foundational insights and a model for inclusive, high-impact community building at ICLR.

## Topics

Building on the success of the inaugural DeLTa 2025 workshop, DeLTa 2026 expands its scope to address new theoretical and algorithmic frontiers emerging from the rapid evolution of modern deep generative models. Discussions will be organized along two major axes — Theoretical Foundations and Algorithms & Applications:

**Theory**:
(1) <u>Unified Theories of Generative Modeling:</u> Develop mathematical frameworks that unify diffusion, flow-matching, energy-based, and autoregressive paradigms, clarifying their equivalence, differences, and underlying stochastic principles.

(2) Optimization and Convergence in Flow-Matching and Diffusion Models: Analyze training and sampling dynamics under different solvers, discretization schemes, and noise schedules; derive convergence guarantees and variance bounds.
(3) Stochastic Control Perspectives: Explore connections between diffusion-based generative models and stochastic optimal control, enabling principled formulations for learning, sampling, and policy-guided generation.)
(4) Post-Training Theoretical Analysis: Study diffusion model post-training phases (e.g., reward-guided fine-tuning, preference alignment) through the lens of optimization, generalization, and stability.
(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) Information-Theoretic and Probabilistic Analysis: Apply tools from information theory and Bayesian inference to quantify uncertainty, mutual information, and representation disentanglement in DGMs.
(7) Geometry and Manifold Learning: Examine the geometric and topological structure of latent spaces; formalize manifold learning principles to improve controllability and representation quality.

**Algorithms and Applications**:
(1) Large Language Diffusion Models (LLDMs): Explore architectures that integrate diffusion mechanisms into large language models; study their compositionality, scalability, and multimodal reasoning capabilities.
(2) One (few)-Step Generative Modeling: Develop efficient one-step and few-step generative algorithms via flow-matching or learned ODE/SDE solvers, reducing sampling cost while preserving quality.
(3) Diffusion Model Post-Training and Adaptation: Design post-training methods for alignment, preference optimization, or reinforcement learning within diffusion frameworks.
(4) Reliability, Interpretability, and Safety: Build tools for evaluating, explaining, and controlling DGM behavior to ensure responsible and robust deployment.
(5) Multimodal and Cross-Domain Generation: Integrate text, image, audio, and video modalities through coherent diffusion and flow-based architectures.
(6) 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.
(7) Generative models for scientific discovery (AI4Science): Develop DGMS for applications in biology, physics, chemistry, material science and environmental science.

## 2. Submission Instructions and Tentative Deadlines

We invite submissions in two formats:

- **short papers** (up to 4 pages, excluding references and appendices)
- **long papers** (up to 8 pages, excluding references and appendices)

All submissions must be prepared as a single PDF (main text and appendices combined) and uploaded via the official submission portal (Openreview). Submissions that have been previously published or are currently under review at other venues will not be considered.

**Tentative Deadlines**
(All dates are in 11.59pm AOE)

- Paper submission deadline: 30 January 2026
- Reviewer bidding dates: 1 - 5 February 2026
- Review Deadline: 26 February 2026
- Paper notification date: 1 March 2026
- Camera ready deadline: 6 March 2026

## 3. Workshop outlines

| Brasilia Time (GMT-3) | Event | Brasilia Time (GMT-3) | 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 - 17:30 | Invited Talk 7 |
| 12:00 - 12:30 | Oral Presentations | 17:30 - 18:00 | Oral Presentations |
| 12:30 - 14:00 | Lunch | 18:00 - 18:10 | Awards & Closing Remarks |

We anticipate receiving 100 long paper submissions plus 80 short paper submissions, with an acceptance rate of 60%~70%. Of the accepted papers, 3% will be invited for oral presentations, 1% will be selected for outstanding paper awards. We expect an audience of 400–500 participants, including both in-person attendees and remote viewers accessing recorded talks and posters

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

We invited 7 speakers, each specializing in key areas of deep generative models (DGMs), and 7 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 (e.g., Arnaud Doucet, University of Oxford & DeepMind; Sitan Chen, Harvard University; Nisha Chandramoorthy, University of Chicago)
- Algorithms for generative models (e.g., Atsushi Nitanda, ASTAR Singapore; Gabriele Steidl, TU Berlin)
- Multimodal and Structured Generative Modeling (e.g., Zahra Kadkhodaie, Flatiron Institute)
- Generative models for science (e.g., René Vidal, University of Pennsylvania; Nisha Chandramoorthy, University of Chicago)
- Safety in generative AI (e.g., René Vidal, University of Pennsylvania; Arnaud Doucet, DeepMind)

Our speaker lineup includes leading figures from both academia and industry, spanning junior and senior researchers across multiple continents. We are committed to diversity, equity, and inclusion: seven confirmed speakers collectively reflect diversity in gender, geography, and professional background. The workshop aims to foster an interdisciplinary and inclusive environment, promoting open dialogue between theorists and practitioners and amplifying the voices of early-career and underrepresented researchers in the DGM community.

- [Confirmed] **Zahra Kadkhodaie** (Postdoc, Flatiron institute, Simons foundation)
  Dr. Zahra Kadkhodaie received her Ph.D. in Data Science from NYU, advised by Profs. Eero Simoncelli and Carlos Fernandez-Granda. Her research focuses on understanding and improving deep neural networks through mathematical symmetries, as well as leveraging the priors of trained denoisers for solving computer vision problems. She holds a B.Sc. in Solid State Physics from K.N. Toosi University and an M.Sc. in Data Science and Psychology from NYU.

- [Confirmed] **Atsushi Nitanda** (Principal Scientist, A*STAR, Singapore)
  Dr. Atsushi Nitanda is a Principal Scientist at A*STAR (CFAR) and Associate Professor at NTU, Singapore. He received his Ph.D. from the University of Tokyo. His research spans stochastic and distributional optimization, mean-field methods, sampling, optimal transport, and diffusion models, with publications at top venues including NeurIPS, ICML, ICLR, and AISTATS. Notably, his work on averaged stochastic gradient descent under the neural tangent kernel regime received the ICLR 2021 Outstanding Paper Award.

- [Confirmed] **René Vidal** (Rachleff University Professor, University of Pennsylvania)
  Professor René Vidal is a Full Professor at University of Pennsylvania. He is internationally recognized for his contributions to machine learning and computer vision, particularly generalized principal component analysis (GPCA) and subspace clustering, which have influenced modern generative modeling. He has authored over 200 papers and a book on GPCA, served as Program Chair for ICCV and CVPR, and is a Fellow of IEEE and IAPR.

- [Confirmed] **Arnaud Doucet** (Senior Staff Research Scientist, Google DeepMind)
  Professor Arnaud Doucet is Professor of Statistics at the University of Oxford and a Senior Staff Research Scientist at Google DeepMind. He is internationally recognized for his contributions to Bayesian methods, Monte Carlo techniques, and computational statistics. His recent work focuses on generative modeling, particularly denoising diffusion models, and computational optimal transport. He is an IMS Fellow, an IMS Medallion Lecturer, and recipient of the Guy Silver Medal from the Royal Statistical Society.

- [Confirmed] **Sitan Chen** (Assistant Professor, Harvard University)
  Dr. Sitan Chen is an Assistant Professor of Computer Science at Harvard University, affiliated with the Theory of Computation group, the ML Foundations group, and the Harvard Quantum Initiative. His research focuses on algorithmic foundations of learning, with recent work on localization-based generative modeling—including

diffusion models, masked language models, and autoregressive models—as well as quantum protocols for learning physical systems. He received his Ph.D. in EECS from MIT, advised by Ankur Moitra, and was an NSF postdoctoral fellow at UC Berkeley. His work has been recognized with an NSF CAREER Award and multiple NSF grants.

- [Confirmed] **Nisha Chandramoorthy** (Assistant Professor, University of Chicago)
  Dr. Nisha Chandramoorthy is an Assistant Professor of Statistics and Computational and Applied Mathematics at the University of Chicago. Her research spans dynamical systems, Bayesian statistics, and machine learning theory, with a focus on understanding and engineering complex nonlinear dynamics arising in climate, biomedical, and AI systems. Before joining UChicago, she was a James C. Edenfield Early-Career Assistant Professor at Georgia Tech and a postdoctoral researcher at MIT's Institute for Data, Systems, and Society. She received her Ph.D. in Computational Dynamical Systems from MIT.

- [Confirmed] **Gabriele Steidl** (Professor at Technische Universität Berlin)
  Professor Gabriele Steidl is a Professor of Mathematics at Technische Universität Berlin. Her research spans computational harmonic analysis, convex optimization, and imaging science, with recent interest in mathematical foundations of machine learning and inverse problems. She is a Fellow of the Society for Industrial and Applied Mathematics (SIAM) for her contributions to computational harmonic analysis and imaging sciences. Before joining TU Berlin, she was a professor at the University of Mannheim and the Technical University of Kaiserslautern. She is also co-author of the book Numerical Fourier Analysis (Birkhäuser, 2018).

## 5. Organizers

Taiji Suzuki (Professor, University of Tokyo),
Mingyuan Bai (Postdoctoral Researcher, RIKEN AIP),
Andi Han (Lecturer, The University of Sydney),
Sara Fridovich-Keil (Assistant Professor, Georgia Tech),
Wei Huang (Research Scientist, RIKEN AIP),
Qing Qu (Assistant Professor, University of Michigan),
Kenji Fukumizu (Director, The Institute of Statistical Mathematics),
Valentin De Bortoli (Research Scientist, Google DeepMind).

**Diversity, background and experience**
Our organizing committee represents a diverse and interdisciplinary team spanning academia and industry, with members based across Asia, Europe, North America, and Oceania. The organizers collectively bring expertise that bridges theoretical foundations—including optimization, deep learning theory, and statistical learning—and application-driven research in diffusion models, sampling, and scientific machine learning. The team also reflects diversity in career stages, ranging from the postdoctoral researcher and research scientists to senior faculty members.

This year's team combines continuity and renewal. Four organizers—Taiji Suzuki, Andi Han, Wei Huang, and Mingyuan Bai—were members of the inaugural DeLTa 2025 workshop, providing continuity and experience from its successful first edition. The remaining four—Qing Qu, Kenji Fukumizu, Valentin De Bortoli, and Sara Fridovich-Keil—join as new organizers, bringing fresh perspectives and expanding the workshop's geographical and thematic reach.

In addition to our team's diverse backgrounds, we have extensive experience in organizing high-impact workshops. For example, (i) 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. (ii) Professor Qing Qu has organized multiple symposiums on generative AI at the University of Michigan, and he has served as the program chair of the inaugural conference on parsimony and learning (CPAL 2024). He co-organized the IMSI Workshop on Computational Imaging in 2024, the Midwest ML Symposium in 2024, and DeepMath Conference in 2025. (iii) Professor Fridovich-Keil has co-organized a special session at Sampling Theory and Applications (SampTA 2025), was on the organizing team for CPAL 2025, and served as an Area Chair at ICCP 2025 and ICASSP 2026. (iv) Valentin De Bortoli has organized the NeurIPS Workshop on Diffusion Models in 2022, 2023, the Neurips workshop on Scientific Understanding of Deep Learning at Neurips 2024, the ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems in 2023 and co-organized the AISSAI Workshop Machine Learning-Assisted Sampling for Scientific Computing in October 2022 in Paris. (v) Professor Kenji Fukumizu served as Co-Program Chair for AISTATS 2021, and has been a Board Member of the Society for Artificial Intelligence and Statistics, the steering organization of AISTATS, since 2022.

This combination of diverse perspectives, geographical representation, and a balanced mix of experienced and new organizers ensures that our team is well positioned to deliver a scientifically rigorous, inclusive, and high-impact workshop at ICLR 2026.

**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 RIKEN AIP. 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.

**Mingyuan Bai** (E-mail: mingyuan.bai@riken.jp)
Google Scholar: https://scholar.google.com/citations?user=Io0_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.

**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 Lecturer from the School of Mathematics and Statistics at University of Sydney. He was a Postdoctoral Researcher from Continuous Optimization Team in RIKEN AIP, 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. He has served as Area chair for ICLR and AISTATS.

**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. He obtained his Ph.D. in Computer Science from the University of Technology Sydney, 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. His contributions to the field are documented in publications such as NeurIPS, ICLR, and ICML et al. He has served as area chair for NeurIPS, ICLR, and AISTATS, and received the CPAL 2025 and KAUST 2025 Rising Star Awards. Dr. Huang also keeps a comprehensive and popular blog on the latest Deep Learning Theory works on social media, with more than 10k followers.

**Qing Qu** (Email: qingqu@umich.edu)
Website: https://qingqu.engin.umich.edu/
Google Scholar: https://scholar.google.com/citations?user=JfblW3MAAAAJ&hl=zh-CN

Dr. Qing Qu is an assistant professor in the EECS department at the University of Michigan. Before that, he was a Moore-Sloan data science fellow at the Center for Data Science, New York University, from 2018 to 2020. He received his Ph.D from Columbia University in Electrical Engineering in Oct. 2018. He received his B.Eng. from Tsinghua University in Jul. 2011, and a M.Sc. from Johns Hopkins University in Dec. 2012, both in Electrical and Computer Engineering. His research interest lies at the intersection of the foundation of data science, machine learning, numerical optimization, and signal/image processing. His current research interests focus on deep representation learning and diffusion models. He is the recipient of the Best Student Paper Award at SPARS'15, the recipient of the Microsoft PhD Fellowship in machine learning in 2016, and the Best Paper Award in the NeurIPS Diffusion Model Workshop in 2023. He received the NSF Career Award in 2022, and Amazon Research Award (AWS AI) in 2023, a UM CHS Junior Faculty Award in 2025, and a Google Research Scholar Award in 2025. He was the program chair of the new Conference on Parsimony & Learning (CPAL'24), area chair of NeurIPS, ICML, and ICLR, and action editor of TMLR.

**Sara Fridovich-Keil** (Email: sfk@gatech.edu)
Website: https://sarafridov.github.io
Google Scholar: https://scholar.google.com/citations?user=9xF7M6wAAAAJ&hl=en

Dr. Sara Fridovich-Keil is an assistant professor in the School of Electrical and Computer Engineering at the Georgia Institute of Technology. Previously, she was an NSF Mathematical Sciences Postdoctoral Research Fellow at Stanford University. Dr. Fridovich-Keil received her PhD from UC Berkeley in Electrical Engineering and Computer Sciences in 2023, and her BSE from Princeton University in Electrical Engineering in 2018. Her research involves foundations of machine learning, signal processing, and optimization applied to inverse problems in computer vision as well as medical and scientific imaging. She is a member of the IEEE Signal Processing Society Computational Imaging Technical Committee and an Associate Editor of IEEE Transactions on Computational Imaging.

**Kenji Fukumizu** (Email: fukumizu@ism.ac.jp)
Website: https://www.ism.ac.jp/~fukumizu/
Google Scholar:https://scholar.google.com/citations?user=Dav2k7cAAAA

Kenji Fukumizu is a Professor at The Institute of Statistical Mathematics, where he serves as Director of the Research Center for Statistical Machine Learning, the Department of Advanced Data Sciences. He received the B.S. and Ph.D. degrees from Kyoto University in 1989 and 1996, respectively. He joined the Research and Development Center, Ricoh Co., Ltd in 1989, worked as a Researcher at the Institute of Physical and Chemical Research (RIKEN) from 1998, and became an Associate Professor at The Institute of Statistical Mathematics in 2000. He was a Visiting Scholar at the Department of Statistics, UC Berkeley in 2002--2003, and a Humboldt Fellow at Max Planck Institute for Biological Cybernetics in 2006--2007. He has been serving as Technical Advisor at Preferred Networks, Inc. since 2018. He has served as an organizing committee member for many conferences, including Program Co-Chair of AISTATS 2021, and Senior Area Chair on the Program Committees of NeurIPS and ICML. He has also served as Chief Editor of the Annals of the Institute of Statistical Mathematics from 2011 to 2017 and currently serves as an Action Editor for the Journal of Machine Learning Research.

**Valentin De Bortoli** (Email: vdebortoli@google.com)
Website: https://vdeborto.github.io/
Google Scholar: https://scholar.google.com/citations?user=dn_F9I4AAAAJ&hl=en&oi=ao
Valentin De Bortoli is a research scientist at Google DeepMind and a chargé de recherche (equiv. to assistant professor) in the Center for Data Science in Ecole Normale Supérieure in Paris (on leave). He was previously a postdoctoral researcher at Oxford University and received his PhD from ENS Paris-Saclay. He has published papers in Nature, ICASSP, COLT, UAI, ICML, NeurIPS, TMLR and JMLR. His research lies at the intersection between applied probability, statistics and machine learning with a recent focus on the interplay between stochastic control, optimal transport and generative modeling.

## 6. Anticipated audience size

Based on the strong participation and community engagement at the first DeLTa Workshop at ICLR 2025, which attracted more than 400 attendees (including both in-person and remote participants), we anticipate an audience of approximately 400–500 participants for the 2nd DeLTa Workshop at ICLR 2026.

Given the rapidly growing interest in deep generative modeling—particularly in areas such as diffusion, flow matching, and theoretical foundations—we expect broad attendance from both academic researchers and industry practitioners. The anticipated audience includes participants from machine learning theory, optimization, statistics, and applied domains such as computer vision, natural language processing, and AI for science.

To accommodate this audience, we plan to provide accessible online materials, including recorded talks and poster PDFs, to engage those who are unable to attend in person.

## 7. Plan to obtain audience for a workshop

We expect strong interest in the 2nd DeLTa Workshop given the rapidly growing attention to deep generative modeling and the success of the first edition at ICLR 2025. To further broaden participation and engagement, we outline the following targeted strategy:

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/DeLTa2026/.

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).

## 8. 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 4:3 for speakers and 6:2 for organizers.
- **Geographic and institutional diversity**: The invited speakers and organizers represent a diverse range of geographic locations. This includes 1 speaker and 5 organizers from the Asia-Pacific region, 2 speakers and 1 organizer from Europe, and 4 speakers and 2 organizers from North America. Additionally, both academia and industry are well-represented, with a speaker ratio of 5:2 and an organizer ratio of 6:2.
- **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 1 postdoctoral researcher, 2 research scientists, 3 assistant professors and 2 professors.
- **Encouragement of Early-Career and Underrepresented Researchers**: The workshop explicitly aims to amplify the voices of early-career researchers and individuals from underrepresented groups. We will particularly encourage submissions to our **tiny/short paper track**, designed to provide visibility to preliminary, resource-limited, or emerging work that may not fit the traditional full-paper format. This aligns with ICLR's mission of promoting accessibility and inclusivity in the ML community.
- **Inclusive Call for Papers and Participation**: Our Call for Papers and announcements will explicitly welcome participation from all individuals, regardless of gender, race, ethnicity, nationality, disability, or socioeconomic background. We will publicize this commitment clearly on our website, ensuring that our workshop remains open and welcoming to everyone.
- **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.

## 9. 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). **Pre-Workshop Social Media Engagements**: Before the workshop, we will actively advertise for our workshop on social media and showcase our speakers and panelists, in order to collect questions for panelists or speakers from potential audiences, especially those who cannot attend our workshop in person. We will also release the decisions for papers on social media, within a week after the accepted paper notification date. Each one or two days, we plan to promote a category of accepted papers, such as deep generative models for security, theories on deep generative models, etc., on social media platforms. Hence we can enhance discussions before our workshop and boost virtual participation for our audience who cannot physically attend.

2). **Live Streaming and Virtual Participation Options**: We will notify all registered participants on 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.

3). **Workshop materials**: We will make all workshop materials (speaker slides, workshop papers, and all posters) available on the workshop website. Our workshop web page will be constantly maintained with all the aforementioned resources. Furthermore, we will share video recordings of talks where possible such as YouTube after the workshop. We will also design all workshop materials with accessibility and usability in mind.

## 10. LLM Usage Policy and Compliance

In accordance with the ICLR 2026 Policies on Large Language Model Usage, this workshop strictly prohibits AI-generated submissions in all tracks, including the tiny and short paper formats. AI-assisted tools (e.g., for grammar or style refinement) may be used under full human authorship and oversight.

For this proposal, large language models were used only for minor language editing and formatting, with all substantive content, analysis, and organization prepared by human authors.

Additionally, workshop organizers may employ LLMs to assist with logistical or administrative tasks—such as reviewer matching or email communication—under human supervision.

## 11. Relevant Workshops

The 2nd DeLTa Workshop on Deep Generative Models in Machine Learning: Theories, Principles, and Efficacy builds directly on the success of its inaugural edition at ICLR 2025, which attracted over 400 participants and fostered rich discussions across theory, algorithms, and applications of deep generative models (DGMs).

The second edition expands the scope and depth of these discussions, incorporating new topics such as flow matching, discrete diffusion models, post-training, and one-step generative modeling, while maintaining its central mission of bridging theoretical understanding and practical efficacy.

Unlike other workshops that focus on specific model families or applied settings, DeLTa uniquely integrates theoretical, algorithmic, and empirical perspectives, aiming to establish a principled and unified framework for DGMs across paradigms including VAEs, GANs, flow-based models, and diffusion models.

Below we outline several relevant workshops and how DeLTa complements and extends them:

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 2026*, 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. 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 2026* 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.

3. 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 2026* 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.

4. Frontiers in Probabilistic Inference: Learning Meets Sampling @ ICLR 2025 [https://sites.google.com/view/fpiworkshop/about?authuser=0] centered on probabilistic inference and sampling-based learning methods, aiming to bridge communities working on scalable sampling and data-driven inference across fields such as physics, biology, and chemistry. While sharing a theoretical foundation, *DeLTa 2026* extends beyond sampling-based approaches by unifying diverse deep generative paradigms—including adversarial, variational, flow-based, and diffusion models—under a principled framework. In contrast to FIP's focus on inference and sampling as methodological tools, DeLTa emphasizes the theoretical, algorithmic, and empirical principles that connect and advance these approaches as part of a cohesive view of deep generative modeling.

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 Integrating Generative and Experimental Platforms for Biomolecular Design (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) Integrating Generative and Experimental Platforms for Biomolecular Design (ICLR 2025 https://www.gembio.ai/) focus on specific generative applications.

By contrast, DeLTa 2026 focuses on the theoretical, algorithmic, and principled foundations that unify and advance these disparate efforts, continuing to serve as a central hub for the DGM theory and practice community.

## 12. 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] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng and Liwei Wang, Visual autoregressive modeling: scalable image generation via next-scale prediction. In Proceedings of the 38th International Conference on Neural Information Processing Systems (2024).
[5] Song, Yang, and Stefano Ermon. "Generative modeling by estimating gradients of the data distribution." Advances in neural information processing systems 32 (2019).
[6] Ho, Jonathan, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models." Advances in neural information processing systems 33 (2020): 6840-6851.