# GEOMETRY-GROUNDED REPRESENTATION LEARNING AND GENERATIVE MODELING

**Alison Pouplin[1], Sharvaree Vadgama, Sékou-Oumar Kaba, Manuel Lecha, Jakub M. Tomczak, Robin Walters, Stefanie Jegelka, Erik Bekkers[2]**
Contacts for communication: [1]`alison.pouplin@gmail.com`, [2]`e.j.bekkers@uva.nl`

## ABSTRACT

Real-world data often originates from physical systems that are governed by geometric and physical laws. Yet, most machine learning methods treat this data as abstract vectors, ignoring the underlying structure that could improve both performance and interpretability. Geometry provides powerful guiding principles—from group equivariance to non-Euclidean metrics—that can preserve the symmetries or the structure inherent in data. We believe those geometric tools are well-suited, and perhaps essential, for representation learning and generative modeling. We propose GRaM, a workshop centered on the principle of *grounding in geometry*, which we define as:

> *An approach is geometrically grounded if it respects the geometric structure of the problem domain and supports geometric reasoning.*

This year, we aim to explore the relevance of geometric methods, particularly in the context of large models, focusing on the theme of *scale and simplicity*. We seek to understand when geometric grounding remains necessary, how to effectively scale geometric approaches, and when geometric constraints can be relaxed in favor of simpler alternatives.

## 1 MOTIVATION

Grounding machine learning models in geometry brings both theoretical clarity and practical benefits. This is especially true in fields such as physics (Batzner et al., 2022), biology (Jumper et al., 2021), and robotics (Wang et al., 2022), where the data is naturally endowed with a geometric structure, and can be expensive to generate. Riemannian geometry and Lie group theory can efficiently help modeling such data. For example, by construction, spherical manifolds capture periodic and directional features, with rotation groups encoding the symmetries (Hall, 2013). Hyperbolic spaces are well-suited for representing hierarchies (Gromov, 1987). Even the space of probability distributions has a geometric structure, as explored in information geometry (Amari, 1997).

Respecting those geometric constraints is essential in generative modeling—not only for generating accurate and valid samples, but also for studying the relationship between the data and the model training dynamics. Learning or modeling the data manifold is crucial for generating samples consistent with its intrinsic structure, avoiding artifacts that break symmetries present in data or distort latent properties Arvanitidis et al. (2017); De Bortoli et al. (2022); Bose et al. (2023); Davis et al. (2024). Beyond sample quality, geometry also provides principled tools to better understand the data (Stanczuk et al., 2022), and the generative models themselves Ross et al. (2024); Loaiza-Ganem et al. (2024).

Despite these advantages, geometric methods often require complex computations and non-trivial implementation. Their necessity can also be questioned in two key scenarios. First, large-scale models can empirically discover invariances on their own, raising questions about when explicit geometric biases remain necessary (Brehmer et al., 2024; Qu & Krishnapriyan, 2024; Abramson et al., 2024). Second, on manifolds with small curvature, Euclidean geometry may provide a sufficient approximation to more complex Riemannian approaches.

These observations raise central questions: How do geometric methods scale with model size and data complexity? Can we reduce architectural or theoretical overhead while maintaining the model's geometric inductive biases?

**The goals of our workshop** are threefold: (1) connect communities working in geometry, representation learning, generative modeling, and theoretical machine learning; (2) encourage dialogue between experts in academia and industry, in order to combine theoretical insights with practical impact; and (3) discuss and anticipate the limitations of geometric methods, with a focus on *scale and simplicity*.

## 2 Workshop Structure and Content

**GRaM 2024 and beyond** This proposal introduces the second edition of the Geometry-grounded Representation Learning and Generative Modeling (GRaM) workshop, building on the success of the inaugural edition held at ICML 2024 Vadgama et al. (2024). The first workshop attracted 300-400 in-person participants and received over 140 submissions across four thematic tracks. The original GRaM workshop included blog posts and tutorials, directly inspired by the ICLR blog series, aimed at making geometric mathematical concepts accessible to broader research communities. This year, we plan to expand these resources by creating a GitHub repository with user-friendly implementations of key models, curated paper collections, tutorials, open research questions, and relevant libraries. Based on last year's attendance, we expect $\sim$ 100-120 submissions and 300-400 attendees for this edition.

**Workshop scope** We invite submissions spanning theory, methods, applications, critical analyses, open questions, and negative results in (but not limited to) the following areas:

1. **Preserving data geometry** · *Preservation of symmetries:* e.g., through equivariant operators. · *Geometric representation systems:* e.g., encoding data in intrinsically structured forms via Clifford algebras or steerable vectors with Clebsch-Gordan products. · *Isometric latent mappings:* e.g., learning latent representations of the data via pullback metrics.
2. **Inducing geometric structure** · *Geometric priors:* e.g., introducing curvature, symmetry, or topological constraints through explicit regularization. · *non-Euclidean generative models:* e.g., extending diffusion models or flow matching models to non-Euclidean domains with a predefined metric. · *Metric-preserving embeddings:* e.g., learning latent spaces where intrinsic geodesic distances are mapped to Euclidean ones.
3. **Geometry in theoretical analysis** · *Data and latent geometry* · Gaining insights on the data manifold, statistical manifold or the latent variables using geometric tools. · *Loss landscape geometry* · Viewing parameters and their optimization trajectory as lying on a manifold, enabling analysis of curvature, critical points, and generalization. · *Theoretical frameworks*· Using differential geometry, algebraic geometry, or group theory to provide a generalizing perspective on generation or representation learning. · *Open problems* · Identifying and addressing unresolved questions and challenges that lie at the intersection of geometry and learning.
4. **Scale and Simplicity** · *Geometry at scale*· Does equivariance retain value in large-scale models? · *Redundancy and minimality*· Evaluating when geometric structure is essential versus when simpler architectures suffice. · *Challenging assumptions* · Reporting negative results or limitations of geometric methods to guide future development.

**Workshop Tracks** We welcome submissions across three tracks:

- **Paper track**: Extended abstracts (tiny papers) or full papers (4–8 pages, excluding references and appendices). Accepted papers will be published in PMLR proceedings. Both abstracts and papers will be presented as posters or selected talks, with a best paper award.
- **Blogpost and Tutorial track**: Short blog posts (max 2000 words) clarifying key ideas or papers, and Colab-based tutorials. All accepted entries will be published on a dedicated site.

- **Competition track**: Code-based submissions via a shared GitHub repo, focused on replicating geometric/generative models with simplified architectures and preserved performance. Accepted entries will be published in PMLR and hosted online.

Each submission will be reviewed by three reviewers under ICLR-style guidelines. Our evaluation prioritizes scientific insight and clarity over the pursuit of state-of-the-art results.

**GRaM in relation to other workshops** We have observed that ICLR has supported a number of excellent workshops in recent years on generative modeling, including *Deep Generative Models in ML: Theory, Principle and Efficacy* and *Frontiers in Probabilistic Inference* (2025). It has also offered geometry-oriented venues such as *Physics for Machine Learning* (2023), though the most recent dedicated geometry workshop was *Geometrical and Topological Representation Learning* (2022). GRaM specifically targets the intersection of generative models and geometry. GRaM bridges generative modeling and geometry, motivated by scientific problems from physics to biology. It unites theory, methods, and applications grounded in geometric structure and offers a complementary space to prior workshops.

## 3 Planning and Logistics

**Timeline** The workshop will take place on April 26–27, with acceptance notifications required by March 1. The proposed submission deadline will be January 30, to allow a four-week review window. The camera-ready versions will be due by March 11.

**Logistics** Submissions will be double blind and managed through OpenReview (except for Blogpost and tutorial track submissions). Each PC member will review no more than three papers. Organizers will not review submissions from their own institutions, and no reviewer will evaluate work authored by individuals they have previously collaborated with. We aim to expand our existing program committee (we previously had 70 reviewers), request one author per submission to participate in reviewing, and thus support around **100-120 submissions** and finalize decisions by February 27, and notify all the authors before the 1st of March.

**Reviewing and LLM usage** GRaM will follow ICLR's Policies on LLM Usage: while we allow use of LLMs for rephrasing and polishing (as done partially in this proposal), we disallow any LLM use in reviewing, proofreading, or authoring scientific content. All authors will be required to include a clear disclosure statement of any LLM use in their submitted paper; violation may lead to desk rejection or revocation of acceptance.

**Advertising and outreach strategy** We are planning on advertising GRaM through mailing lists, social media, and academic networks, and personally reach out to researchers in generative modeling, and geometry. GRaM has already its own Twitter account.

**Virtual access to workshop materials** In addition to the PMLR proceedings, the blogpost track and competition track, we aim to promote on social media the posters being presented, in agreement with the respective authors.

## 4 Diversity and Inclusion

**Gender balance** We have prioritized gender equity; 3/8 of our organization team are women, and **half** of our invited speakers and panelists are women. **Career-stage diversity.** Speakers, panelists, and organizers range from early-career researchers to senior professors and industry leads. **Academic–industry balance.** We have a mix of speakers from academia and from industry labs, highlighting how both sectors contribute to advances in geometry and generative modeling. **Accessibility and inclusion.** We are fully committed to following ADA guidelines and we will proactively implement them throughout the workshop. This includes ensuring visual and audio accessibility for all posters, presentations, and panel discussions, as well as securing wheelchair-accessible stages, seating, and entrances.

| Time | Speaker | | Affiliations |
|---|---|---|---|
| 09:00 - 10:30 | **Invited Talks A** (30 min each) | | |
| | Yaron Lipman | (tentative) | FAIR Meta & Weizmann Institute |
| | Kathlén Kohn | (confirmed) | KTH Royal Institute of Technology |
| | Arash Vahdat | (tentative) | NVIDIA Research |
| 10:30 - 11:00 | Coffee break | | |
| 11:00 - 12:00 | **Contributed talks (x4)** (15 min each) | | |
| 12:00 - 13:30 | Poster session 1 | | |
| 13:30 - 14:00 | Lunch break | | |
| 14:00 - 15:30 | **Invited Talks B** (30 min each) | | |
| | Maya Bechler-Speicher | (confirmed) | Tel-Aviv University & Meta |
| | Gabriel Loaiza-Ganem | (confirmed) | Layer 6 AI |
| | Tess Smidt | (tentative) | Massachusetts Institute of Technology |
| 15:30 - 16:00 | Coffee break | | |
| 16:00 - 17:00 | **Panel discussion** | | |
| | Aditi Krishnapriyan | (tentative) | University of California Berkeley |
| | Gabriel Loaiza-Ganem | (confirmed) | Layer 6 AI |
| | Alexander Tong | (confirmed) | AITHYRA |
| | Tess Smidt | (tentative) | Massachusetts Institute of Technology |
| 17:00 - 18:30 | Poster session 2 | | |

## 5 Speakers and panelists

**Maya Bechler-Speicher** (Scholar) Maya Bechler-Speicher is an AI researcher at Meta, and a PhD student at Tel-Aviv University. She investigates the expressivity of graph neural networks, with a focus on how spectral features, positional encodings, and symmetry constraints, all grounded in geometric structure, influence model performance. Her contributions range from theoretical analysis (Bechler-Speicher et al., 2024) to broader position pieces (Bechler-Speicher et al., 2025) on the state of the field.

**Kathlén Kohn** (Scholar) Kathlén Kohn is an Associate Professor at KTH and an algebraic geometer whose research connects nonlinear and computational geometry with applications in vision, deep learning, and invariant theory (Henry et al., 2024; Marchetti et al.). Widely recognized as an emerging leader at the interface of algebraic geometry and machine learning, she has contributed significant mathematical insights and received several prizes for her work. Her expertise in algebraic constraints and geometric methods makes her perspective particularly relevant to generative modeling.

**Aditi Krishnapriyan** (Scholar) Aditi Krishnapriyan is an Assistant Professor at UC Berkeley (Chemical Engineering and EECS). She develops physics-inspired machine learning methods with applications in molecular dynamics, fluid mechanics, and PDEs, offering both deep theoretical insight and ground-breaking applications in chemistry (Krishnapriyan et al., 2021; Raja et al., 2025). Her recent work links generative models with statistical mechanics and develops scalable neural potentials with physical constraints.

**Yaron Lipman** (Scholar) Yaron Lipman Yaron is a Research Scientist at Meta Fundamental AI Research and a professor at the Weizmann Institute of Science. He leads work on geometric deep learning and generative flows. His Flow Matching work (Lipman et al., 2022) and his subsequent generalization on data manifolds (Chen & Lipman, 2023) are widely regarded as foundational advances in geometric generative modeling.

**Gabriel Loaiza-Ganem** (Scholar) Gabriel Loaiza-Ganem is a Research Scientist at Layer 6 AI whose research sharply investigates the geometry underpinning generative models. His theoretical insights on generative models (Loaiza-Ganem et al., 2024) provide principled explanations for issues such as out-of-distribution likelihoods (Kamkari et al., 2024) and memorization (Ross et al., 2024). Gabriel's contributions are also unusually pedagogical.

**Tess Smidt** (Scholar) Tess Smidt is an Associate Professor in MIT's EECS department and leads the Atomic Architects group at MIT's RLE, where she blends physics, geometry, and ML. Her work on Euclidean neural networks (Geiger & Smidt, 2022) has been extremely influential for E(3)-equivariant modeling of 3D data, with many applications in physics and chemistry. She has also advanced this work with clever methods for controlled symmetry breaking in equivariant networks and generative models (Unke et al., 2021; Costa et al., 2024).

**Alexander Tong** (Scholar) Alexander Tong is a PI at Aithyra working at the interface of generative modeling, optimal transport, and applications in the life sciences. His recent work combines generative models with geometry (Huguet et al., 2022; Atanackovic et al., 2024), leading to impactful applications in biology (Bose et al., 2023).

**Arash Vahdat** (Scholar) is a Research Director at NVIDIA, where he leads the GenAIR team on fundamental generative AI research. His work spans diffusion models, latent variable models, and flows, with applications ranging from vision and text to proteins, molecules, and weather. Notable contributions include Score-Based Generative Modeling in Latent Space (Vahdat et al., 2021) and recent advances in protein generation and accelerated diffusion sampling (Geffner et al., 2025; Cachay et al., 2025). Beyond his academic impact, Arash brings an important industry perspective on scaling and applying generative models, making him a valuable voice for the workshop.

# 6 Organizer teams



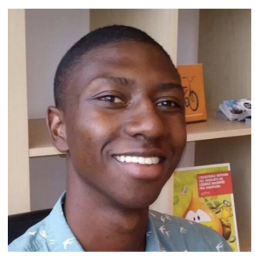




**Sharvaree Vadgama** (email, Scholar) is a PhD student at Amsterdam Machine Learning Lab, University of Amsterdam. Her work is focused on learning representations primarily at the intersection of geometric deep learning and generative modeling. She was a co-organizer of Generative Modeling Summer School (GeMSS) 2023/2024 as well as a co-organizer of Women in AI meetups (Amsterdam chapter). She is also organizer for TAG-DS event co-located with NeurIPS 2025.

**Alison Pouplin** (email, Scholar) is a research scientist at Bayer, Berlin. Prior that, she was a postdoctoral researcher at Aalto University and the Finnish Center of AI. Her current research interests involve Riemannian geometry, information geometry, and generative modeling. She defended her PhD on differential geometric approaches to machine learning.

**Sékou-Oumar Kaba** (email, Scholar) is a PhD student at McGill University and Mila, the Quebec AI institute. His interests are at the intersection of physics and deep learning, with a focus on symmetry and geometric inductive biases. In parallel, he also works on applications of machine learning to material science. He is a co-organizer of the Mila Quantum and AI meetups. He is a recipient of the DeepMind PhD scholarship.

**Manuel Lecha** (email, Scholar) is a PhD student in the ELLIS program, jointly affiliated with the Italian Institute of Technology and Oxford University, under the supervision of Alessio Del Bue and Michael Bronstein. His research explores the intersection between topological and geometric methods. Manuel won multiple awards, including first place at the ICML Topological Deep Learning Challenge in both 2023 and 2024.

**Robin Walters** (email, Scholar) is an assistant professor at Northeastern University, where he leads the Geometric Learning Lab. Robin's research aims to understand symmetry in deep learning and exploit this to improve generalization and data efficiency of deep learning methods. This includes designing equivariant neural networks, symmetry discovery methods, and a symmetry theory for model parameters, applying these approaches to complex dynamic domains such as climate science, transportation, and robotics.

**Jakub M. Tomczak** (email, Scholar) is a Senior Staff Research Scientist (Generative AI x Science) at Chan Zuckerberg Initiative. He is a Generative AI Leader with 15+ years of experience in machine learning, deep learning, and Generative AI. Previously, he was a group leader (Generative AI) at the Eindhoven University of Technology. He has also served as a (fractional) AI leader (head/director/CTO) for multiple companies (e.g., eBay, Qualcomm, startups) and authored the first fully comprehensive book on GenAI (”Deep Generative Modeling”). He served as a Program Chair for NeurIPS 2024.

**Stefanie Jegelka** (email, Scholar) is an Associate Professor (on leave) at MIT EECS, and a Humboldt Professor at TU Munich. At MIT, she is a member of CSAIL, IDSS, and the Center for Statistics and Machine Learning. She previously held a postdoc position at UC Berkeley's AMPlab and Computer Vision Group. Her research in algorithmic machine learning—spanning modeling, optimization, theory, and applications—focuses on exploiting mathematical structure for discrete and combinatorial problems to improve robustness and scale ML algorithms.

**Erik Bekkers** (email, Scholar) is an associate professor in Geometric Deep Learning (GDL) in the Machine Learning Lab of the University of Amsterdam (AMLab, UvA). His main interests include generalizations of group convolutional neural networks and their improvements through sparse, adaptive, and geometric learning mechanisms. He has organized several international workshops and conference sessions (e.g., GeoMedIA'22, GSI '21 and '23, ELLIS workshops). He is a recipient of a VENI and VIDI research grant (Dutch Research Council) and recognized as an ELLIS Scholar in Geometric Deep Learning.

## REFERENCES


Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, pp. 1–3, 2024.

SI Amari. Information geometry. *Contemporary Mathematics*, 203:81–96, 1997.

Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the curvature of deep generative models. *arXiv preprint arXiv:1710.11379*, 2017.

Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the wasserstein manifold. *arXiv preprint arXiv:2408.14608*, 2024.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1): 2453, 2022.

Maya Bechler-Speicher, Ido Amos, Ran Gilad-Bachrach, and Amir Globerson. Graph neural networks use graphs when they shouldn't. In *Forty-first International Conference on Machine Learning*, 2024.

Maya Bechler-Speicher, Ben Finkelshtein, Fabrizio Frasca, Luis Müller, Jan Tönshoff, Antoine Siraudin, Viktor Zaverkin, Michael M Bronstein, Mathias Niepert, Bryan Perozzi, et al. Position: Graph learning will lose relevance due to poor benchmarks. *arXiv preprint arXiv:2502.14546*, 2025.

Avishek Joey Bose, Tara Akhound-Sadegh, Guillaume Huguet, Kilian Fatras, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael Bronstein, and Alexander Tong. Se (3)-stochastic flow matching for protein backbone generation. *arXiv preprint arXiv:2310.02391*, 2023.

Johann Brehmer, Sönke Behrends, Pim de Haan, and Taco Cohen. Does equivariance matter at scale? *arXiv preprint arXiv:2410.23179*, 2024.

Salva Rühling Cachay, Miika Aittala, Karsten Kreis, Noah Brenowitz, Arash Vahdat, Morteza Mardani, and Rose Yu. Elucidated rolling diffusion models for probabilistic weather forecasting. *arXiv preprint arXiv:2506.20024*, 2025.

Ricky TQ Chen and Yaron Lipman. Flow matching on general geometries. *arXiv preprint arXiv:2302.03660*, 2023.

Allan Dos Santos Costa, Ilan Mitnikov, Franco Pellegrini, Ameya Daigavane, Mario Geiger, Zhonglin Cao, Karsten Kreis, Tess Smidt, Emine Kucukbenli, and Joseph Jacobson. Equijump: Protein dynamics simulation via so (3)-equivariant stochastic interpolants. *arXiv preprint arXiv:2410.09667*, 2024.

Oscar Davis, Samuel Kessler, Mircea Petrache, Ismail Ceylan, Michael Bronstein, and Joey Bose. Fisher flow matching for generative modeling over discrete data. *Advances in Neural Information Processing Systems*, 37:139054–139084, 2024.

Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling. *Advances in neural information processing systems*, 35:2406–2422, 2022.

Tomas Geffner, Kieran Didi, Zhonglin Cao, Danny Reidenbach, Zuobai Zhang, Christian Dallago, Emine Kucukbenli, Karsten Kreis, and Arash Vahdat. La-proteina: Atomistic protein generation via partially latent flow matching. *arXiv preprint arXiv:2507.09466*, 2025.

Mario Geiger and Tess Smidt. e3nn: Euclidean neural networks. *arXiv preprint arXiv:2207.09453*, 2022.

Mikhael Gromov. Hyperbolic groups. In *Essays in group theory*, pp. 75–263. Springer, 1987.

Brian C Hall. Lie groups, lie algebras, and representations. In *Quantum Theory for Mathematicians*, pp. 333–366. Springer, 2013.

Nathan W Henry, Giovanni Luca Marchetti, and Kathlén Kohn. Geometry of lightning self-attention: Identifiability and dimension. *arXiv preprint arXiv:2408.17221*, 2024.

Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

Hamidreza Kamkari, Brendan Leigh Ross, Jesse C Cresswell, Anthony L Caterini, Rahul G Krishnan, and Gabriel Loaiza-Ganem. A geometric explanation of the likelihood ood detection paradox. *arXiv preprint arXiv:2403.18910*, 2024.

Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in neural information processing systems*, 34:26548–26560, 2021.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Gabriel Loaiza-Ganem, Brendan Leigh Ross, Rasa Hosseinzadeh, Anthony L Caterini, and Jesse C Cresswell. Deep generative models through the lens of the manifold hypothesis: A survey and new connections. *arXiv preprint arXiv:2404.02954*, 2024.

Giovanni Luca Marchetti, Vahid Shahverdi, Stefano Mereta, Matthew Trager, and Kathlén Kohn. Position: Algebra unveils deep learning-an invitation to neuroalgebraic geometry. In *Forty-second International Conference on Machine Learning Position Paper Track*.

Eric Qu and Aditi S. Krishnapriyan. The importance of being scalable: Improving the speed and accuracy of neural network interatomic potentials across chemical domains, 2024. URL `https://arxiv.org/abs/2410.24169`.

Sanjeev Raja, Martin Šípka, Michael Psenka, Tobias Kreiman, Michal Pavelka, and Aditi S Krishnapriyan. Action-minimization meets generative modeling: Efficient transition path sampling with the onsager-machlup functional. *arXiv preprint arXiv:2504.18506*, 2025.

Brendan Leigh Ross, Hamidreza Kamkari, Tongzi Wu, Rasa Hosseinzadeh, Zhaoyan Liu, George Stein, Jesse C Cresswell, and Gabriel Loaiza-Ganem. A geometric framework for understanding memorization in generative models. *arXiv preprint arXiv:2411.00113*, 2024.

Jan Stanczuk, Georgios Batzolis, Teo Deveney, and Carola-Bibiane Schönlieb. Your diffusion model secretly knows the dimension of the data manifold. *arXiv preprint arXiv:2212.12611*, 2022.

Oliver Unke, Mihail Bogojeski, Michael Gastegger, Mario Geiger, Tess Smidt, and Klaus-Robert Müller. Se (3)-equivariant prediction of molecular wavefunctions and electronic densities. *Advances in Neural Information Processing Systems*, 34:14434–14447, 2021.

Sharvaree Vadgama, Erik Bekkers, Alison Pouplin, Sekou-Oumar Kaba, Robin Walters, Hannah Lawrence, Tegan Emerson, Henry Kvinge, Jakub Tomczak, and Stephanie Jegelka. Preface to geometry-grounded representation learning and generative modeling (gram) workshop. In *Proceedings of the Geometry-grounded Representation Learning and Generative Modeling Workshop (GRaM)*, volume 251, pp. 1–6. PMLR, 29 Jul 2024.

Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in neural information processing systems*, 34:11287–11302, 2021.

Dian Wang, Robin Walters, Xupeng Zhu, and Robert Platt. Equivariant $q$ learning in spatial action spaces. In *Conference on Robot Learning*, pp. 1713–1723. PMLR, 2022.