Novelty Assessment Report

Paper: FALCON: Few-step Accurate Likelihoods for Continuous Flows

PDF URL: https://openreview.net/pdf?id=FbssShlI4N

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Scalable sampling of molecular states in thermodynamic equilibrium is a long-standing challenge in statistical physics. Boltzmann Generators tackle this problem by pairing a generative model, capable of exact likelihood computation, with importance sampling to obtain consistent samples under the target distribution. Current Boltzmann Generators primarily use continuous normalizing flows (CNFs) trained with flow matching for efficient training of powerful models. However, likelihood calculation for these models is extremely costly, requiring thousands of function evaluations per sample, severely limiting their adoption. In this work, we propose Few-step Accurate Likelihoods for Continuous Flows (FALCON), a method which allows for few-step sampling with a likelihood accurate enough for importance sampling applications by introducing a hybrid training objective that encourages invertibility. We show FALCON outperforms state-of-the-art normalizing flow models for molecular Boltzmann sampling and is \emph{two orders of magnitude faster} than the equivalently performing CNF model.

Core Task Landscape

This paper addresses: **Sampling Molecular Conformations from Boltzmann Distributions**

A total of **50 papers** were analyzed and organized into a taxonomy with **26 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Generative Model Architectures for Boltzmann Sampling**
- **Sampling Acceleration and Efficiency Techniques**
- **Surrogate Models and Transfer Operators**
- **Conformational Ensemble Prediction and Analysis**
- **Domain-Specific Applications and Benchmarks**
- **Theoretical Foundations and General Frameworks**
- **Energy-Weighted and Continuous Flow Training**

Complete Taxonomy Tree

- Sampling Molecular Conformations from Boltzmann Distributions Survey Taxonomy
- Generative Model Architectures for Boltzmann Sampling
  - Normalizing Flow-Based Generators
  - Boltzmann Generators with Continuous Normalizing Flows ★ (4 papers)
    - [0] FALCON: Few-step Accurate Likelihoods for Continuous Flows (Anon et al., 2026) View paper
    - [5] Transferable boltzmann generators (Klein, 2024) View paper
    - [6] Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning (F. NoÃ©, 2019) View paper
    - [8] Scalable equilibrium sampling with sequential boltzmann generators (Bose, 2025) View paper
  - Equivariant and Symmetry-Preserving Flows (1 papers)
    - [10] Equivariant flow matching (Klein, 2023) View paper
  - Rigid Body and Orientation-Aware Flows (1 papers)
    - [23] Rigid body flows for sampling molecular crystal structures (KÃ¶hler, 2023) View paper
  - Diffusion and Score-Based Models
  - Energy-Based Diffusion Training (3 papers)
    - [1] Iterated denoising energy matching for sampling from boltzmann densities (Rector-Brooks, 2024) View paper
    - [3] Consistent sampling and simulation: Molecular dynamics with energy-based diffusion models (Plainer, 2025) View paper
    - [12] Energy-based diffusion generator for efficient sampling of Boltzmann distributions. (Yan Wang, 2025) View paper
  - Torsional and Internal Coordinate Diffusion (2 papers)
    - [4] Torsional diffusion for molecular conformer generation (Jing, 2022) View paper
    - [28] Molecular Conformation Generation via Shifting Scores (Zhou, 2023) View paper
  - Consistency Models and Distillation for Diffusion (1 papers)
    - [29] Efficient and Unbiased Sampling of Boltzmann Distributions via Consistency Models (Zhang Feng-zhe, 2024) View paper
  - Potential Energy Guidance and Debiasing (1 papers)
    - [24] Potential Score Matching: Debiasing Molecular Structure Sampling with Potential Energy Guidance (Guo Liya, 2025) View paper
  - GFlowNets and Trajectory-Based Generators (3 papers)
  - [15] Torsional-GFN: a conditional conformation generator for small molecules (Volokhova, 2025) View paper
  - [21] Towards equilibrium molecular conformation generation with GFlowNets (Volokhova, 2024) View paper
  - [33] RL Boltzmann generators for conformer generation in data-sparse environments (Patel, 2022) View paper

- Sampling Acceleration and Efficiency Techniques
  - Importance Sampling and Reweighting (3 papers)
  - [9] Generation of conformational ensembles of small molecules via surrogate model-assisted molecular dynamics (Juan Viguera Diez, 2024) View paper
  - [11] â¦ Unnormalized Distributions With Deep Generative Models Toward the Acceleration of Molecular Design and Conformational Sampling With Deep Learning (Diez, 2024) View paper
  - [13] Sampling From Molecular Unnormalized Distributions With Deep Generative Models Toward the Acceleration of Molecular Design and Conformational â¦ (Diez, 2024) View paper
  - Walk-Jump and Hybrid Sampling Strategies (2 papers)
  - [2] Jamun: Transferable molecular conformational ensemble generation with walk-jump sampling (Ameya Daigavane, 2024) View paper
  - [49] Smart walking: A new method for Boltzmann sampling of protein conformations (Ruhong Zhou, 1997) View paper
  - Flow Perturbation and Jacobian Approximation (1 papers)
  - [18] Flow Perturbation to Accelerate Unbiased Sampling of Boltzmann distribution (Peng Xin, 2024) View paper
  - Enhanced Sampling with Collective Variables (2 papers)
  - [19] Accelerated Sampling of Boltzmann distributions (Orland, 2022) View paper
  - [38] Enhancing Diffusion-Based Sampling with Molecular Collective Variables (Nam, 2025) View paper
  - Parallel Tempering and Temperature Control (2 papers)
  - [20] Parallel tempering algorithm for conformational studies of biological molecules (Ulrich H. E. Hansmann, 1997) View paper
  - [40] Temperature and friction accelerated sampling of boltzmann-gibbs distribution (Tao, 2010) View paper
- Surrogate Models and Transfer Operators
  - Implicit Transfer Operator Learning (1 papers)
  - [14] Implicit transfer operator learning: Multiple time-resolution models for molecular dynamics (Winther, 2023) View paper
  - ML Force Field Surrogates for Accelerated Sampling (2 papers)
  - [36] BoostMD: Accelerating molecular sampling by leveraging ML force field features from previous time-steps (Batatia, 2024) View paper
  - [48] BoostMD: Accelerated Molecular Sampling Leveraging ML Force Field Features (LL Schaaf, 2024) View paper
- Conformational Ensemble Prediction and Analysis
  - AlphaFold2-Based Ensemble Approximation (2 papers)
  - [22] Approximating conformational Boltzmann distributions with AlphaFold2 predictions (Benjamin P. Brown, 2023) View paper
  - [26] Approximating Projections of Conformational Boltzmann Distributions with AlphaFold2 Predictions: Opportunities and Limitations (Benjamin P. Brown, 2024) View paper
  - Experimental Data-Driven Ensemble Refinement (4 papers)
  - [16] Modeling RNA secondary structure folding ensembles using SHAPE mapping data (Aleksandar Spasic, 2018) View paper
  - [32] Molecular dynamics simulations with replica-averaged structural restraints generate structural ensembles according to the maximum entropy principle (Andrea Cavalli, 2013) View paper
  - [44] Structure refinement with molecular dynamics and a Boltzmann-weighted ensemble (Jens Fennen, 1995) View paper
  - [50] A Structural Ensemble of Hen Egg-White Lysozyme in Aqueous Solution Based on 2043 NMR NOE, 3J-Coupling and S2 Order-Parameter Data. (Lorna J. Smith, 2025) View paper
  - Conformer Selection and Boltzmann Weighting (3 papers)
  - [17] Exploring the impacts of conformer selection methods on ion mobility collision cross section predictions (Nielson, 2021) View paper
  - [37] Determination of the Relative Stereochemistry of Flexible Organic Compounds by Ab Initio Methods: Conformational Analysis and BoltzmannâAveraged GIAO 13C â¦ (G Barone, 2002) View paper
  - [46] Optimization methods for conformational sampling using a Boltzmannâweighted mean field approach (Thomas Huber, 1996) View paper
- Domain-Specific Applications and Benchmarks
  - RNA Structure and Interaction Prediction (2 papers)
  - [7] The Boltzmann distributions of molecular structures predict likely changes through random mutations (Ahnert, 2023) View paper
  - [30] Prediction of RNA-RNA interaction structure by centroids in the Boltzmann ensemble (Chitsaz, 2022) View paper
  - Benchmark Sets and Validation Studies (1 papers)
  - [47] A diverse and chemically relevant solvation model benchmark set with flexible molecules and conformer ensembles. (Lukas Wittmann, 2025) View paper
  - Pre-training and Transfer Learning for Molecular GNNs (2 papers)
  - [34] Unifying Force Prediction and Molecular Conformation Generation Through Representation Alignment (L Pinede, 2025) View paper
  - [39] Pre-training of Molecular GNNs via Conditional Boltzmann Generator (Koge, 2023) View paper
- Theoretical Foundations and General Frameworks
  - Ab Initio and Phase Space Averaging (1 papers)
  - [27] Ab initio machine learning of phase space averages (Jan Weinreich, 2022) View paper
  - Maximum Entropy and Replica-Averaged Restraints (1 papers)
  - [43] Calculation of conformational ensembles from potentials of mena force: an approach to the knowledge-based prediction of local structures in globular proteins (Sippl, 1990) View paper
  - Methodological Reviews and Search Strategies (2 papers)
  - [25] On searching in, sampling of, and dynamically moving through conformational space of biomolecular systems: A review (Markus Christen, 2008) View paper
  - [41] Modeling Boltzmann weighted structural ensembles of proteins using AI based methods (Akashnathan Aranganathan, 2024) View paper
  - Macroscopic and Analog Systems (2 papers)
  - [31] A macroscopic device described by a Boltzmann-like distribution (Simon Tricard, 2013) View paper
  - [45] Monte Carlo simulation of nâbutane in water. Conformational evidence for the hydrophobic effect (W.L. Jorgensen, 1982) View paper

- Energy-Weighted and Continuous Flow Training (2 papers)
  - [35] Energy-Weighted Flow Matching: Unlocking Continuous Normalizing Flows for Efficient and Scalable Boltzmann Sampling (Dern, 2025) View paper
  - [42] Physically Interpretable Biomolecular Conformation Generation with A Deep Probabilistic Framework (Zhao, 2023) View paper

Narrative

Core task: sampling molecular conformations from Boltzmann distributions. The field has organized itself around several complementary directions. Generative Model Architectures for Boltzmann Sampling explores diverse neural frameworks—normalizing flows (including continuous variants like those in Boltzmann Generators[6]), diffusion models (e.g., Torsional Diffusion[4]), and flow matching approaches (Equivariant Flow Matching[10])—each offering different trade-offs between expressivity and computational cost. Sampling Acceleration and Efficiency Techniques focuses on speeding up convergence through methods like parallel tempering (Parallel Tempering Algorithm[20]) and flow perturbation (Flow Perturbation Acceleration[18]). Surrogate Models and Transfer Operators leverage learned approximations of dynamics (Surrogate Model MD[9], Implicit Transfer Operator[14]) to bypass expensive simulations, while Conformational Ensemble Prediction and Analysis emphasizes extracting ensemble properties and validating generated distributions. Domain-Specific Applications and Benchmarks ground the methodology in real molecular systems, Theoretical Foundations provide rigorous underpinnings, and Energy-Weighted and Continuous Flow Training refine how models incorporate energetic information during learning (Energy-Weighted Flow Matching[35]).

Recent work has intensified around two contrasting themes: whether to use discrete-time diffusion or continuous-time flows, and how tightly to couple energy functions during training versus sampling. Iterated Denoising Energy[1] and Consistent Sampling Simulation[3] exemplify efforts to improve diffusion-based samplers, while Transferable Boltzmann Generators[5] and Sequential Boltzmann Generators[8] extend flow-based methods to new regimes. FALCON[0] sits within the normalizing flow branch alongside Boltzmann Generators[6], emphasizing continuous normalizing flows for direct Boltzmann sampling. Compared to Transferable Boltzmann Generators[5], which prioritizes generalization across molecular families, FALCON[0] focuses on refining the flow architecture itself for improved sampling fidelity. The interplay between architectural innovation and training strategies—whether to weight samples by energy or rely on post-hoc reweighting—remains an active question shaping how practitioners balance accuracy, efficiency, and transferability across diverse molecular systems.

Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

1. Transferable boltzmann generators

Authors: Klein, Leon, NoÃ©, Frank, Leon Klein, et al. (6 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

The generation of equilibrium samples of molecular systems has been a long-standing problem in statistical physics. Boltzmann Generators are a generative machine learning method that addresses this issue by learning a transformation via a normalizing flow from a simple prior distribution to the target Boltzmann distribution of interest. Recently, flow matching has been employed to train Boltzmann Generators for small molecular systems in Cartesian coordinates. We extend this work and propose a f...

Relationship Analysis

Both papers belong to the category of Boltzmann Generators with Continuous Normalizing Flows, employing CNFs trained via flow matching for molecular conformation sampling. They overlap in using flow-based architectures for Boltzmann sampling with importance sampling corrections, but differ fundamentally in scope: FALCON focuses on achieving few-step accurate likelihood computation through invertibility regularization for efficient single-molecule sampling, while the candidate paper addresses transferability across chemical space, enabling zero-shot generalization to unseen molecular systems without retraining.

2. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning

Authors: F. NoÃ©, Jonas KÃ¶hler, Hao Wu, Frank NoeÌ, Simon Olsson | **Year/Venue**: 2019 | **URL**: View paper

Abstract

Efficient sampling of equilibrium states Molecular dynamics or Monte Carlo methods can be used to sample equilibrium states, but these methods become computationally expensive for complex systems, where the transition from one equilibrium state to another may only occur through rare events. NoÃ© et al. used neural networks and deep learning to generate distributions of independent soft condensed-matter samples at equilibrium (see the Perspective by Tuckerman). Supervised training is used to cons...

Relationship Analysis

Both papers belong to the Boltzmann Generators with Continuous Normalizing Flows category, using flow-based generative models for molecular Boltzmann sampling. The original paper (FALCON) and the candidate paper (Boltzmann generators: Sampling equilibrium states) both address sampling molecular conformations from Boltzmann distributions using invertible neural networks with importance sampling for reweighting. The key difference is that FALCON focuses on achieving few-step accurate likelihoods through a hybrid training objective combining flow matching with cycle-consistency regularization, while the candidate paper introduces the foundational Boltzmann generator framework using invertible neural networks to transform between configuration space and a simple latent distribution.

3. Scalable equilibrium sampling with sequential boltzmann generators

Authors: Bose, Avishek Joey, Charlie B. Tan, Lin Chen, A. Bose, et al. (15 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Scalable sampling of molecular states in thermodynamic equilibrium is a long-standing challenge in statistical physics. Boltzmann generators tackle this problem by pairing normalizing flows with importance sampling to obtain uncorrelated samples under the target distribution. In this paper, we extend the Boltzmann generator framework with two key contributions, denoting our framework Sequential Boltzmann Generators (SBG). The first is a highly efficient Transformer-based normalizing flow operati...

Relationship Analysis

Both papers belong to the Boltzmann Generators with Continuous Normalizing Flows category, employing flow-based models for molecular Boltzmann sampling. While FALCON focuses on achieving few-step accurate likelihoods through a hybrid training objective combining flow matching with cycle-consistency to ensure invertibility, SBG (Sequential Boltzmann Generators) emphasizes inference-time scaling via sequential Monte Carlo and annealed Langevin dynamics to transport proposal samples toward the target distribution. The key distinction is that FALCON addresses the computational cost of likelihood evaluation in CNFs through invertible flow maps, whereas SBG uses non-equivariant normalizing flows with inference-time annealing strategies.

Contributions Analysis

Overall novelty summary. The paper proposes FALCON, a method for accelerating likelihood computation in continuous normalizing flows (CNFs) used for Boltzmann sampling. It sits within the 'Boltzmann Generators with Continuous Normalizing Flows' leaf, which contains four papers including the original work. This leaf is part of the broader 'Normalizing Flow-Based Generators' branch, which also includes equivariant flows and rigid body flows as sibling leaves. The taxonomy reveals this is a moderately populated research direction within the larger field of fifty papers, suggesting active but not overcrowded exploration of CNF-based approaches for molecular sampling.

The paper's position within the normalizing flow branch distinguishes it from neighboring diffusion-based methods (e.g., 'Energy-Based Diffusion Training', 'Torsional and Internal Coordinate Diffusion') and GFlowNets. The taxonomy shows clear boundaries: normalizing flows emphasize exact likelihood computation and invertibility, while diffusion methods trade off likelihood tractability for flexible denoising processes. FALCON's focus on few-step sampling with accurate likelihoods addresses a computational bottleneck specific to CNFs, contrasting with diffusion acceleration techniques like consistency models or distillation found in sibling branches. The 'Sampling Acceleration and Efficiency Techniques' category exists separately, indicating FALCON bridges architectural innovation with efficiency concerns.

Among twenty-seven candidates examined, the contribution-level analysis reveals mixed novelty signals. The core FALCON method (Contribution A) examined ten candidates with one appearing to provide overlapping prior work, suggesting some precedent for few-step flow acceleration exists within this limited search scope. The hybrid training objective for invertibility (Contribution B) examined seven candidates with one potential refutation, indicating prior exploration of invertibility constraints. The scalable equivariant architecture (Contribution C) examined ten candidates with none clearly refuting it, suggesting this component may be more distinctive. These statistics reflect a focused semantic search, not exhaustive coverage of the flow-based sampling literature.

Based on the limited search scope of twenty-seven top-ranked candidates, FALCON appears to combine known elements—CNF acceleration, invertibility training, equivariant architectures—in a novel configuration targeting a specific computational bottleneck. The taxonomy context shows this work incrementally advances a moderately active research direction rather than opening entirely new territory. The analysis cannot assess whether broader literature beyond the top-K semantic matches contains additional relevant prior work, particularly in adjacent fields like general normalizing flow acceleration or non-molecular applications of few-step flows.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: FALCON: Few-step accurate likelihoods for continuous flows

Description: The authors introduce FALCON, a continuous flow-based generative model that enables few-step sampling while providing fast and accurate likelihood computation for Boltzmann sampling. The method uses a hybrid training objective combining regression loss and a cycle-consistency term to encourage invertibility, making it suitable for importance sampling applications.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. PointFlow: 3D Point Cloud Generation With Continuous Normalizing Flows

URL: View paper

Brief Assessment

PointFlow[51] focuses on 3D point cloud generation using continuous normalizing flows in a two-level hierarchical framework (shape distribution and point distribution), not on few-step sampling with accurate likelihoods for Boltzmann generation or importance sampling applications.

---

2. Building Normalizing Flows with Stochastic Interpolants

URL: View paper

Brief Assessment

Stochastic Interpolants[53] focuses on building normalizing flows using stochastic interpolants with a variational objective for the velocity field. FALCON addresses a different problem: enabling few-step sampling with accurate likelihoods for Boltzmann generation through a hybrid training objective combining regression loss and cycle-consistency. The candidate does not demonstrate prior work on FALCON's specific contribution of few-step sampling with importance-sampling-ready likelihoods for molecular applications.

---

3. Amortized Sampling with Transferable Normalizing Flows

URL: View paper

Brief Assessment

Amortized Sampling Flows[57] focuses on transferable normalizing flows for peptide systems with zero-shot generalization across sequences, while FALCON addresses few-step sampling with accurate likelihoods for Boltzmann generation. The architectural approaches and primary objectives differ substantially.

---

4. Joint Distillation for Fast Likelihood Evaluation and Sampling in Flow-based Models

URL: View paper

Prior Art Analysis

Joint Distillation Flows[52] demonstrates that prior work exists addressing the same core problem: enabling few-step sampling with accurate likelihood computation in continuous normalizing flows. Both papers tackle the computational bottleneck of likelihood evaluation in flow-based models, which traditionally requires hundreds to thousands of function evaluations. Joint Distillation Flows[52] presents F2D2, a method that 'simultaneously reduces the number of nfes required for both sampling and likelihood evaluation by two orders of magnitude,' directly paralleling FALCON's claim of being 'two orders of magnitude faster than the equivalently performing cnf model' while providing 'few-step sampling with a likelihood accurate enough for importance sampling applications.'

Evidence

Evidence 1 - **Rationale**: Both papers present methods that enable few-step sampling with accurate likelihood computation in continuous flows, addressing the same fundamental problem of computational efficiency in flow-based models. - **Original**: in this work, we propose few-stepaccuratelikelihoods forcontinuousflows (falcon), a method which allows for few-step sampling with a likelihood accurate enough for importance sampling applications by introducing a hybrid training objective that encourages invertibility. - **Candidate**: we present fast flow joint distillation (f2d2), a framework that simultaneously reduces the number of nfes required for both sampling and likelihood evaluation by two orders of magnitude.

Evidence 2 - **Rationale**: Both papers claim to solve the computational bottleneck of likelihood evaluation in flow-based models while maintaining quality, with similar magnitude of speedup claims. - **Original**: we show falcon outperforms state-of-the-art normalizing flow models for molecular boltzmann sampling and is two orders of magnitude fasterthan the equivalently performing cnf model. - **Candidate**: experiments demonstrate f2d2's capability of achieving accurate log-likelihood with few-step evaluations while maintaining high sample quality, solving a long-standing computational bottleneck in flow-based generative models.

Evidence 3 - **Rationale**: Both papers identify the same specific problem: the computational cost of likelihood evaluation requiring hundreds to thousands of function evaluations in flow-based models. - **Original**: likelihood calculation for these models is extremely costly, requiring thousands of function evaluations per sample, severely limiting their adoption. - **Candidate**: yet paradoxically, some of today's best generative models -- diffusion and flow-based models -- still require hundreds to thousands of neural function evaluations (nfes) to compute a single likelihood.

Evidence 4 - **Rationale**: Joint Distillation Flows[52] explicitly addresses the gap that FALCON claims to fill: existing few-step methods either abandon likelihood computation or require expensive integration, demonstrating awareness of the same technical challenge. - **Original**: we introduce few-stepaccuratelikelihoods forcontinuousflows(falcon): a new continuous flow-based generative model for boltzmann sampling that is invertible, trainable with a regression loss, and supports free-form architectures, while enabling both few-step generation and efficient likelihood evalua... - **Candidate**: while recent distillation methods have successfully accelerated sampling to just a few steps, they achieve this at the cost of likelihood tractability: existing approaches either abandon likelihood computation entirely or still require expensive integration over full trajectories.

---

5. Entropy-Informed Weighting Channel Normalizing Flow for Deep Generative Models

URL: View paper

Brief Assessment

Entropy-Informed Weighting[58] focuses on multi-scale architecture design for normalizing flows in image generation, not on few-step sampling with accurate likelihoods for Boltzmann generation in molecular physics applications.

---

6. Deep Generative Models for Fast, Efficient and Personalized Speech Synthesis

URL: View paper

Brief Assessment

Personalized Speech Synthesis[60] focuses on speech synthesis using flow-based generative models for audio waveforms and mel-spectrograms, not on molecular Boltzmann sampling or few-step likelihood computation for continuous normalizing flows in statistical physics applications.

---

7. Verlet Flows: Exact-Likelihood Integrators for Flow-Based Generative Models

URL: View paper

Brief Assessment

Verlet Flows[54] addresses exact likelihood computation in CNFs through symplectic integrators on augmented state-space, whereas FALCON focuses on few-step sampling with a hybrid training objective combining regression loss and cycle-consistency for invertibility in Boltzmann generation applications.

---

8. Towards Climate Variable Prediction with Conditioned Spatio-Temporal Normalizing Flows

URL: View paper

Brief Assessment

Climate Variable Flows[59] focuses on spatio-temporal prediction for climate data using conditional normalizing flows, not on few-step sampling with accurate likelihoods for molecular Boltzmann generation or importance sampling applications.

---

9. Flow-based generative models as iterative algorithms in probability space

URL: View paper

Brief Assessment

Iterative Algorithms Flows[56] focuses on theoretical frameworks for flow-based models as iterative algorithms in probability space using Wasserstein metrics and JKO schemes, not on few-step sampling with accurate likelihoods for Boltzmann generation applications.

---

10. Stochastic normalizing flows

URL: View paper

Brief Assessment

Stochastic Normalizing Flows[55] focuses on combining stochastic sampling blocks (MCMC/Langevin dynamics) with deterministic normalizing flows for asymptotically unbiased sampling, not on few-step deterministic flow generation with accurate likelihoods for Boltzmann sampling.

Contribution 2: Hybrid training objective for invertibility

Description: The authors propose a hybrid training objective that combines a regression loss for stable few-step generation with a cycle-consistency term to encourage invertibility prior to convergence. This design allows the model to be invertible, trainable with regression loss, and compatible with free-form architectures while supporting efficient likelihood evaluation.

This contribution was assessed against **7 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Semi-supervised biomedical translation with cycle wasserstein regression GANs

URL: View paper

Brief Assessment

Cycle Wasserstein GANs[73] focuses on semi-supervised biomedical translation tasks using cycle-consistency with adversarial losses, not on training invertible flow models with regression loss for likelihood evaluation in continuous normalizing flows.

---

2. Unsupervised Domain Transfer with Conditional Invertible Neural Networks

URL: View paper

Brief Assessment

Unsupervised Domain Transfer[74] focuses on domain transfer between simulated and real spectral imaging data using conditional invertible neural networks with maximum likelihood and adversarial training. The candidate does not address the specific problem of training flow models for few-step generation with accurate likelihoods for Boltzmann sampling, which is the core focus of the original paper's hybrid objective combining regression loss and cycle-consistency.

---

3. Enhancing Bayesian Inference-Based Damage Diagnostics Through Domain Translation With Application to Miter Gates

URL: View paper

Brief Assessment

Miter Gates Diagnostics[70] uses a conditional invertible neural network (CINN) for Bayesian damage diagnostics in structural engineering, not for training flow models with regression loss and cycle-consistency for molecular sampling or general likelihood evaluation.

4. Conditional invertible neural networks for guided image generation

URL: View paper

Prior Art Analysis

Conditional Invertible Networks[72] demonstrates prior work on training invertible flow models using a hybrid objective that combines regression loss with cycle-consistency terms to encourage invertibility. The paper explicitly describes using maximum likelihood training combined with a cycle-consistency regularization term to ensure invertibility, which directly parallels the ORIGINAL paper's claimed novel hybrid training approach. Both papers address the same fundamental challenge: achieving invertibility in flow models while maintaining stable training through regression-based objectives.

Evidence

Evidence 1 - **Rationale**: The candidate paper describes invertible neural networks that are trainable with regression-based objectives and support flexible architectures, establishing prior work on combining invertibility with regression training. - **Original**: we introduce few-stepaccuratelikelihoods forcontinuousflows(falcon): a new continuous flow-based generative model for boltzmann sampling that is invertible, trainable with a regression loss, and supports free-form architectures, while enabling both few-step generation and efficient likelihood evalua... - **Candidate**: we propose a third approach, by extending invertible neural networks (inns, dinh et al. (2016); kingma & dhariwal (2018); ardizzone et al. (2019)) for the task of conditional image generation, by adding conditioning inputs to their core building blocks. inns are neural networks which are by construc...

Evidence 2 - **Rationale**: The candidate demonstrates maintaining invertibility while using regression-based training through architectural design, showing prior work on combining these objectives. - **Original**: falcon leverages a hybrid training objective that combines a regression loss for stable and efficient few-step generation with a cycle-consistency term to encourage invertibility prior to convergence. - **Candidate**: we adapt the design of eqs. (1) and (2) to produce a conditional version of the coupling block. because the subnetworks sj and tj are never inverted, we can concatenate conditioning datac to their inputs without losing the invertibility, replacing s1(u2) with s1(u2,c) etc.

Evidence 3 - **Rationale**: Both papers establish theoretical guarantees for invertibility under their respective training objectives, demonstrating prior theoretical work on this approach. - **Original**: minimizing this loss has a less strict requirement for the correctness of the boltzmann generator specifically: proposition 2.let u⋆ θ be a minimizer oflinv (eq.(8)) with respect to some v. then, for sufficiently smooth u⋆ θ and v and for any (s, t)∈[0,1] 2, xu(·, s, t)is an invertible map everywher... - **Candidate**: training a network with this loss yields an estimate of the maximum likelihood network parameters ˆθml. from there, we can perform conditional generation for a fixed c by sampling z and using the inverted network g: xgen = g(z; c,ˆθml), with z ~pz(z).

5. Latent Space Regression in GANs for Invertible Image Generation

URL: View paper

Brief Assessment

Latent Space Regression[76] focuses on inverting GAN generators for image generation tasks using a separate encoder network, not on training invertible flow models with regression and cycle-consistency for likelihood evaluation in continuous flows.

6. Inverse design of compressor/fan blade profiles based on conditional invertible neural networks

URL: View paper

Brief Assessment

Compressor Blade Design[71] uses conditional invertible neural networks (cINN) for blade profile design, which is a different application domain (aerodynamic engineering) with different objectives (mapping design parameters to performance metrics). The invertibility in cINN serves a different purpose than FALCON's hybrid objective combining regression loss with cycle-consistency for flow matching models.

7. Inverse Molecule Design with Invertible Neural Networks as Generative Models

URL: View paper

Brief Assessment

Inverse Molecule Design[75] uses flow-based invertible neural networks for molecule generation but does not describe a hybrid training objective combining regression loss with cycle-consistency. The candidate focuses on bidirectional learning for chemical property prediction and molecule generation, not on training objectives for achieving invertibility in flow models.

Contribution 3: Scalable softly equivariant continuous flow architecture

Description: The authors introduce a simple and scalable softly equivariant continuous flow architecture that significantly improves over existing state-of-the-art equivariant flow model architectures. This architectural contribution enables the use of larger and more expressive models for molecular sampling tasks.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. EC-Conf: A ultra-fast diffusion model for molecular conformation generation with equivariant consistency

URL: View paper

Brief Assessment

EC-Conf[67] focuses on consistency models for molecular conformation generation, not on continuous flow architectures for Boltzmann sampling. The architectural approaches and application domains differ fundamentally.

2. E (n) equivariant normalizing flows

URL: View paper

Brief Assessment

Equivariant Normalizing Flows[65] focuses on E(n) equivariant normalizing flows for molecular generation using continuous-time flows with EGNN dynamics. The original paper's contribution is about a 'softly equivariant' architecture using diffusion transformers with data augmentation for few-step sampling, which differs from the strictly equivariant EGNN-based continuous flows in the candidate.

3. Equivariant neural diffusion for molecule generation

URL: View paper

Brief Assessment

Equivariant Neural Diffusion[66] focuses on learnable forward processes for diffusion models in molecule generation, not on continuous flow architectures for molecular sampling tasks as in the original paper's Boltzmann generator context.

4. Geodiff: A geometric diffusion model for molecular conformation generation

URL: View paper

Brief Assessment

GeoDiff[61] focuses on diffusion models for molecular conformation generation using equivariant graph neural networks, not continuous flow architectures for Boltzmann sampling. The architectural approaches and problem domains differ fundamentally.

5. Equivariant flow matching with hybrid probability transport for 3d molecule generation

URL: View paper

Brief Assessment

Hybrid Probability Transport[69] focuses on equivariant flow matching for 3D molecule generation with hybrid probability paths for coordinates and categorical features. The original paper's contribution is about a softly equivariant continuous flow architecture for molecular sampling in Boltzmann generators, which is a different application domain and architectural approach.

6. Mdm: Molecular diffusion model for 3d molecule generation

URL: View paper

Brief Assessment

MDM[68] focuses on diffusion models for 3D molecule generation using dual equivariant encoders to model different interatomic forces, not on continuous flow architectures for molecular sampling tasks as described in the original paper.

7. Equivariant flow matching for molecular conformer generation

URL: View paper

Brief Assessment

Equivariant Flow Matching[64] focuses on molecular conformer generation using equivariant transformers with flow matching, not on Boltzmann generators or the specific architectural innovations for importance sampling applications described in the original paper.

8. Et-flow: Equivariant flow-matching for molecular conformer generation

URL: View paper

Brief Assessment

Et-flow[62] focuses on equivariant flow matching for molecular conformer generation using a different architectural approach (equivariant transformers) and application domain (conformer generation from molecular graphs), rather than Boltzmann sampling from energy distributions. The architectural innovations are distinct from FALCON's hybrid training objective for few-step flows.

9. Equivariant flow matching

URL: View paper

Brief Assessment

Equivariant Flow Matching[10] focuses on training objectives (equivariant optimal transport flow matching) rather than architectural innovations. The candidate uses standard equivariant graph neural networks similar to prior work, not a novel 'softly equivariant' architecture as claimed in the original paper.

10. Equivariant diffusion for molecule generation in 3d

URL: View paper

Brief Assessment

Equivariant Diffusion Molecules[63] focuses on diffusion models for 3D molecule generation with E(3) equivariance, not on continuous flow architectures for molecular sampling. The candidate uses EGNN layers within a diffusion framework, while the original contribution describes a flow-based architecture for Boltzmann generation.

Appendix: Text Similarity Detection

Textual similarity detection checked 30 papers and found 6 similarity segment(s) across 3 paper(s).

The following **3 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Scalable equilibrium sampling with sequential boltzmann generators

Detected in: Core Task (sibling)

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

2. Amortized Sampling with Transferable Normalizing Flows

Detected in: Contribution: contribution_1

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

3. Transferable boltzmann generators

Detected in: Core Task (sibling)

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] FALCON: Few-step Accurate Likelihoods for Continuous Flows View paper
- [1] Iterated denoising energy matching for sampling from boltzmann densities View paper
- [2] Jamun: Transferable molecular conformational ensemble generation with walk-jump sampling View paper
- [3] Consistent sampling and simulation: Molecular dynamics with energy-based diffusion models View paper
- [4] Torsional diffusion for molecular conformer generation View paper

- [5] Transferable boltzmann generators View paper
- [6] Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning View paper
- [7] The Boltzmann distributions of molecular structures predict likely changes through random mutations View paper
- [8] Scalable equilibrium sampling with sequential boltzmann generators View paper
- [9] Generation of conformational ensembles of small molecules via surrogate model-assisted molecular dynamics View paper
- [10] Equivariant flow matching View paper
- [11] â¦ Unnormalized Distributions With Deep Generative Models Toward the Acceleration of Molecular Design and Conformational Sampling With Deep Learning View paper
- [12] Energy-based diffusion generator for efficient sampling of Boltzmann distributions. View paper
- [13] Sampling From Molecular Unnormalized Distributions With Deep Generative Models Toward the Acceleration of Molecular Design and Conformational â¦ View paper
- [14] Implicit transfer operator learning: Multiple time-resolution models for molecular dynamics View paper
- [15] Torsional-GFN: a conditional conformation generator for small molecules View paper
- [16] Modeling RNA secondary structure folding ensembles using SHAPE mapping data View paper
- [17] Exploring the impacts of conformer selection methods on ion mobility collision cross section predictions View paper
- [18] Flow Perturbation to Accelerate Unbiased Sampling of Boltzmann distribution View paper
- [19] Accelerated Sampling of Boltzmann distributions View paper
- [20] Parallel tempering algorithm for conformational studies of biological molecules View paper
- [21] Towards equilibrium molecular conformation generation with GFlowNets View paper
- [22] Approximating conformational Boltzmann distributions with AlphaFold2 predictions View paper
- [23] Rigid body flows for sampling molecular crystal structures View paper
- [24] Potential Score Matching: Debiasing Molecular Structure Sampling with Potential Energy Guidance View paper
- [25] On searching in, sampling of, and dynamically moving through conformational space of biomolecular systems: A review View paper
- [26] Approximating Projections of Conformational Boltzmann Distributions with AlphaFold2 Predictions: Opportunities and Limitations View paper
- [27] Ab initio machine learning of phase space averages View paper
- [28] Molecular Conformation Generation via Shifting Scores View paper
- [29] Efficient and Unbiased Sampling of Boltzmann Distributions via Consistency Models View paper
- [30] Prediction of RNA-RNA interaction structure by centroids in the Boltzmann ensemble View paper
- [31] A macroscopic device described by a Boltzmann-like distribution View paper
- [32] Molecular dynamics simulations with replica-averaged structural restraints generate structural ensembles according to the maximum entropy principle View paper
- [33] RL Boltzmann generators for conformer generation in data-sparse environments View paper
- [34] Unifying Force Prediction and Molecular Conformation Generation Through Representation Alignment View paper
- [35] Energy-Weighted Flow Matching: Unlocking Continuous Normalizing Flows for Efficient and Scalable Boltzmann Sampling View paper
- [36] BoostMD: Accelerating molecular sampling by leveraging ML force field features from previous time-steps View paper
- [37] Determination of the Relative Stereochemistry of Flexible Organic Compounds by Ab Initio Methods: Conformational Analysis and Boltzmannâ Averaged GIAO 13C â¦ View paper
- [38] Enhancing Diffusion-Based Sampling with Molecular Collective Variables View paper
- [39] Pre-training of Molecular GNNs via Conditional Boltzmann Generator View paper
- [40] Temperature and friction accelerated sampling of boltzmann-gibbs distribution View paper
- [41] Modeling Boltzmann weighted structural ensembles of proteins using AI based methods View paper
- [42] Physically Interpretable Biomolecular Conformation Generation with A Deep Probabilistic Framework View paper
- [43] Calculation of conformational ensembles from potentials of mena force: an approach to the knowledge-based prediction of local structures in globular proteins View paper
- [44] Structure refinement with molecular dynamics and a Boltzmann-weighted ensemble View paper
- [45] Monte Carlo simulation of nâbutane in water. Conformational evidence for the hydrophobic effect View paper
- [46] Optimization methods for conformational sampling using a Boltzmannâweighted mean field approach View paper
- [47] A diverse and chemically relevant solvation model benchmark set with flexible molecules and conformer ensembles. View paper
- [48] BoostMD: Accelerated Molecular Sampling Leveraging ML Force Field Features View paper
- [49] Smart walking: A new method for Boltzmann sampling of protein conformations View paper
- [50] A Structural Ensemble of Hen Egg-White Lysozyme in Aqueous Solution Based on 2043 NMR NOE, 3J-Coupling and S2 Order-Parameter Data. View paper
- [51] PointFlow: 3D Point Cloud Generation With Continuous Normalizing Flows View paper
- [52] Joint Distillation for Fast Likelihood Evaluation and Sampling in Flow-based Models View paper
- [53] Building Normalizing Flows with Stochastic Interpolants View paper
- [54] Verlet Flows: Exact-Likelihood Integrators for Flow-Based Generative Models View paper
- [55] Stochastic normalizing flows View paper
- [56] Flow-based generative models as iterative algorithms in probability space View paper
- [57] Amortized Sampling with Transferable Normalizing Flows View paper
- [58] Entropy-Informed Weighting Channel Normalizing Flow for Deep Generative Models View paper
- [59] Towards Climate Variable Prediction with Conditioned Spatio-Temporal Normalizing Flows View paper
- [60] Deep Generative Models for Fast, Efficient and Personalized Speech Synthesis View paper
- [61] Geodiff: A geometric diffusion model for molecular conformation generation View paper
- [62] Et-flow: Equivariant flow-matching for molecular conformer generation View paper
- [63] Equivariant diffusion for molecule generation in 3d View paper
- [64] Equivariant flow matching for molecular conformer generation View paper
- [65] E (n) equivariant normalizing flows View paper
- [66] Equivariant neural diffusion for molecule generation View paper
- [67] EC-Conf: A ultra-fast diffusion model for molecular conformation generation with equivariant consistency View paper
- [68] Mdm: Molecular diffusion model for 3d molecule generation View paper

- [69] Equivariant flow matching with hybrid probability transport for 3d molecule generation View paper
- [70] Enhancing Bayesian Inference-Based Damage Diagnostics Through Domain Translation With Application to Miter Gates View paper
- [71] Inverse design of compressor/fan blade profiles based on conditional invertible neural networks View paper
- [72] Conditional invertible neural networks for guided image generation View paper
- [73] Semi-supervised biomedical translation with cycle wasserstein regression GANs View paper
- [74] Unsupervised Domain Transfer with Conditional Invertible Neural Networks View paper
- [75] Inverse Molecule Design with Invertible Neural Networks as Generative Models View paper
- [76] Latent Space Regression in GANs for Invertible Image Generation View paper