
GO-Diff: Data-free and amortized global structure optimization

Nikolaj Rønne[†], Tejs Vegge and Arghya Bhowmik

Technical University of Denmark

CAPeX Pioneer Center for Accelerating P2X Materials Discovery

[†]niron@dtu.dk

Abstract

We introduce GO-Diff, a diffusion-based method for global structure optimization that learns to directly sample low-energy atomic configurations without requiring prior data or explicit relaxation. GO-Diff is trained from scratch using a Boltzmann-weighted score-matching loss, leveraging only the known energy function to guide generation toward thermodynamically favorable regions. The method operates in a two-stage loop of self-sampling and model refinement, progressively improving its ability to target low-energy structures. Compared to traditional optimization pipelines, GO-Diff achieves competitive results with significantly fewer energy evaluations. Moreover, by reusing pretrained models across related systems, GO-Diff supports amortized optimization — enabling faster convergence on new tasks without retraining from scratch.

1 Introduction

The potential energy surface (PES) is a high-dimensional, non-convex function that encodes the stability of atomic configurations by mapping atomic positions to their associated potential energy. Exploring this surface to identify low-energy (and thus stable) structures is a fundamental challenge in computational materials science, chemistry, and catalysis.[1] This task, often referred to as global structure optimization, underpins applications ranging from the discovery of new catalytic surfaces to the design of functional materials.

Traditional global optimization methods — such as random structure search (RSS)[2], basin-hopping[3], genetic algorithms[4], and simulated annealing[5] — rely on local relaxation with gradient-based optimizers to identify nearby minima. While effective, these methods are computationally expensive due to the many energy and force evaluations required, and their reliance on local optimization can limit the exploration of complex energy landscapes. Machine learning interatomic potentials[6, 7, 8, 9, 10], whether pre-trained[11, 12] or learned on-the-fly[13, 14], can reduce the cost of local optimization. However, they require carefully selected training data to capture relevant minima; otherwise, the search risks getting trapped in self-reinforcing local minima (by only gathering new data near known regions of the PES) far from the true global optimum.

Score-based diffusion models[15] have shown promise for structure generation in molecular[16, 17, 18] and materials science[19, 20, 21, 22]. However, applying them to global optimization tasks is challenging: the goal is to sample rare, low-energy configurations corresponding to the global minimum of a PES, but the data distribution over such structures is typically unknown or inaccessible. This makes it difficult to train models that prioritize physically meaningful, low-energy samples.

In this work, we introduce GO-Diff, a diffusion-based framework for global structure optimization that operates without prior data by leveraging the known energy function during training. GO-Diff

learns to generate low-energy configurations without explicit relaxation using a Boltzmann-weighted score-matching loss combined with iterative self-sampling and annealed training.

Our key contributions are:

1. **A data-free generative optimization method** that directly samples minima of the potential energy surface.
2. **A Boltzmann-weighted loss with annealing** to guide sampling toward low-energy regions while maintaining exploration.
3. **Amortized optimization** through transfer of pretrained models across related systems.
4. **Empirical evidence** of superior sample efficiency compared to classical search methods.

GO-Diff departs from prior Boltzmann samplers (see SI) by using a direct Boltzmann-weighted score-matching loss that avoids force evaluations and Monte Carlo estimates, and by demonstrating amortized optimization through transfer of pretrained models across related systems.

2 Methods

Training loop. GO-Diff optimizes atomic structures by training a diffusion model to generate low-energy configurations without requiring data or local relaxation. It operates in a self-sampling loop: the model generates atomic structures via reverse-time diffusion, evaluates their energies, and uses the resulting samples to refine itself.

A replay buffer $\mathcal{B} = \{(\mathbf{x}_0^{(i)}, E^{(i)})\}$ stores generated configurations $\mathbf{x}_0^{(i)}$ along with their energies $E^{(i)}$. The buffer is initially seeded with N samples from an untrained model. At each iteration, the current model samples N new structures from the reverse stochastic differential equation (SDE), which are then evaluated using an energy function. These are merged with existing buffer entries.

This iterative process enables GO-Diff to learn directly from its own generations, progressively focusing on thermodynamically favorable regions of the energy landscape.

Boltzmann-weighted score matching. We wish to sample from the Boltzmann distribution[23] at temperature T

$$\pi_T(\mathbf{x}) = \frac{\exp(-E(\mathbf{x})/T)}{Z_T}, \quad Z_T = \int \exp(-E(\mathbf{x})/T) d\mathbf{x}, \quad (1)$$

where $E(\mathbf{x})$ is the potential energy of configuration $\mathbf{x}$.

If we had direct i.i.d. samples $\mathbf{x}_0 \sim \pi_T$, the denoising score matching objective for training a score network $s_\theta(\mathbf{x}_t, t)$ would follow standard denoising score matching[24]

$$\mathcal{L}_\theta = \mathbb{E}_{t \sim \mathcal{U}(0,1)} \left[\lambda(t) \mathbb{E}_{\mathbf{x}_0 \sim \pi_T, \mathbf{x}_t \sim p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} \left\| s_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) \right\|_2^2 \right], \quad (2)$$

where $s_\theta(\mathbf{x}_t, t)$ is the score model, $\mathbf{x}_t$ is a noisy version of $\mathbf{x}_0$ and $p_{t|0}$ denotes the marginal of the forward SDE conditioned on the initial data point usually referred to as the transition kernel.

Because direct samples from π_T are unavailable, we approximate expectations with samples $\mathbf{x}_0 \sim q$, where q denotes the empirical buffer distribution. From importance sampling (IS)[25, 26], we have that for any integrable g ,

$$\mathbb{E}_{\pi_T}[g(\mathbf{x})] = \mathbb{E}_q \left[g(\mathbf{x}) \frac{\pi_T(\mathbf{x})}{q(\mathbf{x})} \right]. \quad (3)$$

In practice, q is not known analytically and we only store a buffer $\mathcal{B}$ of evaluated configurations. From self-normalized importance sampling (SNIS)[26], we have that the expectation can be approximated by

$$\mathbb{E}_{\pi_T}[g(\mathbf{x})] \approx \sum_i w(E^{(i)}) g(\mathbf{x}^{(i)}), \quad (4)$$

where the weights are the Boltzmann weights normalized over the buffer samples

$$w(E) = \frac{\exp(-E/T)}{\sum_{E^{(i)} \in \mathcal{B}} \exp(-E^{(i)}/T)}. \quad (5)$$

Substituting $g = \|s_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)\|_2^2$ from Eq. 2 yields the *Boltzmann-weighted score matching* loss

$$\mathcal{L}_\theta^{\text{Boltzmann}} = \mathbb{E}_{t \sim \mathcal{U}(0,1)} \left[\lambda(t) \mathbb{E}_{\mathbf{x}_0 \sim q, \mathbf{x}_t \sim p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} w(E) \|s_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)\|_2^2 \right], \quad (6)$$

where $\mathbf{x}_0$ now represents buffer samples. The loss emphasizes low-energy structures without requiring force labels or sampling from the true Boltzmann distribution.

Annealing strategy. To balance exploration and exploitation, the temperature T is annealed from a high initial value to a low final value over training. This encourages broad exploration early on and convergence toward deep minima in later stages, analogous to simulated annealing.

Force-field guidance. Atomic forces are typically available alongside energies at negligible additional cost. Although they are not used directly in the Boltzmann-weighted loss, we exploit them through a force-prediction head attached to the shared representation backbone of the score network. The predicted forces $F_\theta(\mathbf{x})$ are incorporated into a predictor–corrector sampling scheme: each reverse diffusion (predictor) step is followed by a force-based correction step,

$$\Delta \mathbf{x} = \alpha(1-t)^\zeta F_\theta(\mathbf{x}), \quad (7)$$

where t denotes the diffusion time, ζ is a scalar hyperparameter, and α is an adaptive step size. As $t \rightarrow 0$, the corrective term plays a progressively larger guiding role, steering samples toward low-energy configurations, while early in sampling the stochastic diffusion dynamics promote exploration. The force head is trained jointly on all evaluated configurations during GO-Diff training. This approach, similar in spirit to Ref. [21], provides physically grounded force-field guidance (FFG) that accelerates convergence toward equilibrium configuration.

Diffusion process. We follow the methodology of Ref. [21, 27] for the atomistic diffusion process.

3 Results

We evaluate GO-Diff on two atomistic optimization tasks using the MACE-MP0 universal potential[12], following the established benchmark systems from Ref. [28]. The first task involves optimizing the placement of a Pt addatom on a stepped Pt surface. Although the system is three-dimensional, the energy landscape can be effectively visualized by projecting optimized addatom positions along the surface normal, yielding a two-dimensional representation. The second task targets the discovery of a stable Pt-heptamer cluster on a large 6×6 Pt(111) surface. This more complex system serves both as a benchmark against classical RSS and as a testbed for amortized optimization via transfer of the pretrained score model from the first task.

Pt-addatom on stepped Pt surface. Figure 1(a) shows the projected 2D potential energy surface for different Pt addatom placements. The colored points represent GO-Diff buffer samples at various temperatures during annealing. Initially, samples from the untrained model are broadly scattered. As training progresses and temperature decreases, samples increasingly concentrate around low-energy basins near the global minimum.

Figure 1(b) illustrates the corresponding energy evolution of buffer structures. Lower temperatures bias sampling toward deeper minima, validating the effectiveness of the Boltzmann-weighted loss and annealing schedule.

Pt-heptamer on Pt(111). We next evaluate GO-Diff (with and without FFG) on the more challenging task of discovering the Pt-heptamer cluster on a 6×6 Pt(111) surface. We compare against RSS by measuring both success and the average number of evaluations required to find the target structure. In addition, we test amortized optimization: transferring a pretrained GO-Diff model from the addatom task to initialize the score model in this new system.

Figure 2 reports the mean best energy achieved over time for each method. RSS shows limited progress, often failing to improve beyond its initial state. In contrast, GO-Diff consistently discovers the Pt-heptamer within 2,560 energy evaluations. We observe an improvement using force-field guidance both in terms of evaluated energies and success. With model transfer, amortized GO-Diff




Figure 1: **Addatom optimization on a stepped Pt surface.** (a) Projected 2D energy surface for different addatom placements. Colored dots indicate buffer samples, with color denoting annealing temperature (red = high, blue = low). (b) Evolution of buffer structure energies during annealing. Sampling gradually concentrates in low-energy basins.




Figure 2: **Benchmarking GO-Diff on Pt-heptamer discovery.** Mean best energy over eight independent experiments with faint lines showing the moving average evaluate energy of each run (Only plotted for GO-Diff and GO-Diff + Transfer). Crosses indicate discovery of the Pt-heptamer. GO-Diff rapidly improves, while RSS stagnates. Transfer learning further accelerates convergence.

(without FFG) achieves even faster convergence and better performance, finding the target in seven out of eight runs using ~ 600 energy evaluation on average. See Table 1 for details.

This acceleration is expected: the transferred model already captures bonding preferences, such as the stability of hollow sites below step edges. As a result, the optimization task reduces to adjusting interatomic geometry, rather than learning bonding from scratch. These results highlight GO-Diff’s ability to reuse generative knowledge across systems, enabling amortized optimization in more complex scenarios.

Table 1 summarizes these findings. GO-Diff achieves better success across all experiments using fewer evaluations than RSS, and amortized optimization via model transfer further reduces the computational budget by more than $2\times$ on average.

The stochasticity of the diffusion process allows GO-Diff to robustly escape local minima and explore high-quality regions of the potential energy surface. Unlike greedy optimizers, it maintains

Table 1: Comparison of methods by total single-point evaluations, success, and mean success iteration. A successful optimization run discovers the Pt-heptamer structure.

Method	# Evaluations	Success	Mean Success Iteration
RSS	10,000	1/8	7,816
GO-Diff	2,560	5/8	1,667
GO-Diff + FFG	2,560	8/8	1,994
GO-Diff + Transfer	1,280	7/8	591

diversity throughout sampling, guided by the annealed Boltzmann-weighted objective—enabling both efficiency and reliability even in complex optimization tasks.

4 Discussion

Extending to multi-objective or compositional design. A unique strength of GO-Diff is that it can be extended to optimize over additional properties, such as the atomic composition, electronic or catalytic properties. These can be incorporated directly into the diffusion model or via conditioning guidance, enabling optimization beyond geometry. This is particularly attractive in materials design, where properties are tightly coupled.

5 Limitations

Samples, temperature profile and training steps. The number of samples per iteration, the annealing schedule and the number of training steps per iteration are critical hyperparameters in GO-Diff. Too few samples reduce buffer diversity and slow convergence, while overly large batches raise computational cost with diminishing returns. Likewise, overly aggressive temperature decay or overfitting each iteration can cause premature collapse around suboptimal minima, whereas slower schedules and fewer training steps enable broader exploration and better generalization. In our experiments, we adopt an exponential decay schedule to balance exploration and exploitation. Future extensions may include dynamic temperature adjustment, adaptive sampling, and tuning the number of training steps per iteration—all of which could improve efficiency and robustness without manual tuning.

Scalability. The scalability of atomistic diffusion models remains limited. In practice, current score-based models have mostly been demonstrated on systems with fewer than 20 atoms.[19, 22] Although recent work such as [21, 29, 30, 31] has extended diffusion-based approaches to larger systems, further research is needed to make GO-Diff scalable to very large and realistic sized systems.

6 Conclusion

We introduced GO-Diff, a generative diffusion-based framework for global structure optimization that avoids relaxation steps and pretraining. Trained directly on the energy landscape via a Boltzmann-weighted loss, GO-Diff efficiently samples low-energy configurations. It outperforms random structure search in both efficiency and success, and enables amortized optimization through transfer across related systems. Our results establish diffusion models as scalable optimization engines for atomistic modeling, with promising extensions to compositional design, surrogate acceleration, and multi-objective optimization.

Acknowledgments

We are thankful for funding by the Pioneer Center for Accelerating P2X Materials Discovery (CAPeX), DNRf grant number P3. A.B. thank the Det Frie Forskningsråd under Project “Data-driven quest for TWh scalable Na-ion battery (TeraBatt)” (Ref. Number 2035-00232B) and “Autonomous agents of Discovery for earth-Abundant Na-ion battery cathodes (ADANA)” (Ref. Number 3164-00297B). The authors also thank the Novo Nordisk Foundation Data Science Research Infrastructure 2022

Grant: A high-performance computing infrastructure for data-driven research on sustainable energy materials, Grant no. NNF22OC0078009.

References

- [1] Artem R. Oganov, Chris J. Pickard, Qiang Zhu, and Richard J. Needs. Structure prediction drives materials discovery. *Nature Reviews Materials*, 4(5):331–348, May 2019.
- [2] Chris J Pickard and R J Needs. Ab initio random structure searching. *Journal of Physics: Condensed Matter*, 23(5):053201, January 2011.
- [3] David J. Wales and Jonathan P. K. Doye. Global Optimization by Basin-Hopping and the Lowest Energy Structures of Lennard-Jones Clusters Containing up to 110 Atoms. *The Journal of Physical Chemistry A*, 101(28):5111–5116, July 1997.
- [4] D. M. Deaven and K. M. Ho. Molecular Geometry Optimization with a Genetic Algorithm. *Physical Review Letters*, 75(2):288–291, July 1995.
- [5] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598):671–680, May 1983.
- [6] Albert P. Bartók, Mike C. Payne, Risi Kondor, and Gábor Csányi. Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons. *Physical Review Letters*, 104(13):136403, April 2010.
- [7] K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller. SchNet – A deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24):241722, March 2018.
- [8] Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *Proceedings of the 38th International Conference on Machine Learning*, pages 9377–9388. PMLR, July 2021.
- [9] 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, May 2022.
- [10] Ilyes Batatia, David P. Kovacs, Gregor Simm, Christoph Ortner, and Gabor Csanyi. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields. *Advances in Neural Information Processing Systems*, 35:11423–11436, December 2022.
- [11] Bowen Deng, Peichen Zhong, KyuJung Jun, Janosh Riebesell, Kevin Han, Christopher J. Bartel, and Gerbrand Ceder. CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nature Machine Intelligence*, 5(9):1031–1041, September 2023.
- [12] Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M. Elena, Dávid P. Kovács, Janosh Riebesell, Xavier R. Advincula, Mark Asta, William J. Baldwin, Noam Bernstein, Arghya Bhowmik, Samuel M. Blau, Vlad Cărare, James P. Darby, Sandip De, Flaviano Della Pia, Volker L. Deringer, Rokas Elijošius, Zakariya El-Machachi, Edwin Fako, Andrea C. Ferrari, Annalena Genreith-Schrieffer, Janine George, Rhys E. A. Goodall, Clare P. Grey, Shuang Han, Will Handley, Hendrik H. Heenen, Kersti Hermansson, Christian Holm, Jad Jaafar, Stephan Hofmann, Konstantin S. Jakob, Hyunwook Jung, Venkat Kapil, Aaron D. Kaplan, Nima Karimitari, Namu Kroupa, Jolla Kullgren, Matthew C. Kuner, Domantas Kuryla, Guoda Liepuoniute, Johannes T. Margraf, Ioan-Bogdan Magdău, Angelos Michaelides, J. Harry Moore, Aakash A. Naik, Samuel P. Niblett, Sam Walton Norwood, Niamh O’Neill, Christoph Ortner, Kristin A. Persson, Karsten Reuter, Andrew S. Rosen, Lars L. Schaaf, Christoph Schran, Eric Sivonxay, Tamás K. Stenczel, Viktor Svahn, Christopher Sutton, Cas van der Oord, Eszter Varga-Umbrich, Tejs Vegge, Martin Vondrák, Yangshuai Wang, William C. Witt, Fabian Zills, and Gábor Csányi. A foundation model for atomistic materials chemistry, December 2023.
- [13] Malthe K. Bisbo and Bjørk Hammer. Efficient Global Structure Optimization with a Machine-Learned Surrogate Model. *Physical Review Letters*, 124(8):086102, February 2020.
- [14] Nikolaj Rønne, Mads-Peter V. Christiansen, Andreas Møller Slavensky, Zeyuan Tang, Florian Brix, Mikkel Elkjær Pedersen, Malthe Kjær Bisbo, and Bjørk Hammer. Atomistic structure search using local surrogate model. *The Journal of Chemical Physics*, 157(17), 2022.

- [15] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations, February 2021.
- [16] Emiel Hoogeboom, Víctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant Diffusion for Molecule Generation in 3D. In *Proceedings of the 39th International Conference on Machine Learning*, pages 8867–8887. PMLR, June 2022.
- [17] Minkai Xu, Alexander S. Powers, Ron O. Dror, Stefano Ermon, and Jure Leskovec. Geometric Latent Diffusion Models for 3D Molecule Generation. In *Proceedings of the 40th International Conference on Machine Learning*, pages 38592–38610. PMLR, July 2023.
- [18] François Cornet, Grigory Bartosh, Mikkel N. Schmidt, and Christian A. Naesseth. Equivariant Neural Diffusion for Molecule Generation. *Advances in Neural Information Processing Systems*, 37:49429–49460, December 2024.
- [19] Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal Diffusion Variational Autoencoder for Periodic Material Generation, March 2022.
- [20] Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal Structure Prediction by Joint Equivariant Diffusion on Lattices and Fractional Coordinates. In *Workshop on "Machine Learning for Materials" ICLR 2023*, April 2023.
- [21] Nikolaj Rønne, Alán Aspuru-Guzik, and Bjørk Hammer. Generative diffusion model for surface structure discovery. *Physical Review B*, 110(23):235427, December 2024.
- [22] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, Roberto Sordillo, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jieliang Li, Chunlei Yang, Wenjie Li, Ryota Tomioka, and Tian Xie. A generative model for inorganic materials design. *Nature*, 639(8055):624–632, March 2025.
- [23] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, September 2019.
- [24] Yang Song and Stefano Ermon. Generative Modeling by Estimating Gradients of the Data Distribution, October 2020.
- [25] Tim Hesterberg. Weighted Average Importance Sampling and Defensive Mixture Distributions. 37(2), 1995.
- [26] S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso, and A. M. Stuart. Importance Sampling: Intrinsic Dimension and Computational Cost. *Statistical Science*, 32(3):405–431, 2017.
- [27] Nikolaj Rønne and Bjørk Hammer. Atomistic Generative Diffusion for Materials Modeling, July 2025.
- [28] José A. Garrido Torres, Paul C. Jennings, Martin H. Hansen, Jacob R. Boes, and Thomas Bligaard. Low-Scaling Algorithm for Nudged Elastic Band Calculations Using a Surrogate Machine Learning Model. *Physical Review Letters*, 122(15):156001, April 2019.
- [29] François Cornet, Bardi Benediktsson, Bjarke Hastrup, Mikkel N. Schmidt, and Arghya Bhowmik. OM-Diff: Inverse-design of organometallic catalysts with guided equivariant denoising diffusion. *Digital Discovery*, 3(9):1793–1811, 2024.
- [30] Chaitanya K. Joshi, Xiang Fu, Yi-Lun Liao, Vahe Gharakhanyan, Benjamin Kurt Miller, Anuroop Sriram, and Zachary Ward Ulissi. All-atom Diffusion Transformers: Unified generative modelling of molecules and materials. In *Forty-Second International Conference on Machine Learning*, June 2025.
- [31] François Cornet, Pratham Deshmukh, Bardi Benediktsson, Mikkel N Schmidt, and Arghya Bhowmik. Improving generative inverse design of molecular catalysts in small data regime. *Machine Learning: Science and Technology*, 6(2):025057, June 2025.
- [32] Siddarth Krishnamoorthy, Satvik Mehul Mashkaria, and Aditya Grover. Diffusion Models for Black-Box Optimization. In *Proceedings of the 40th International Conference on Machine Learning*, pages 17842–17857. PMLR, July 2023.

- [33] Zihao Li, Hui Yuan, Kaixuan Huang, Chengzhuo Ni, Yinyu Ye, Minshuo Chen, and Mengdi Wang. Diffusion Model for Data-Driven Black-Box Optimization, March 2024.
- [34] Tara Akhound-Sadegh, Jarrid Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, and Alexander Tong. Iterated Denoising Energy Matching for Sampling from Boltzmann Densities, June 2024.
- [35] RuiKang OuYang, Bo Qiang, Zixing Song, and José Miguel Hernández-Lobato. BNEM: A Boltzmann Sampler Based on Bootstrapped Noised Energy Matching, February 2025.
- [36] Aaron Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Brandon Wood, Daniel Levine, Bin Hu, Brandon Amos, Brian Karrer, Xiang Fu, Guan-Hong Liu, and Ricky T. Q. Chen. Adjoint Sampling: Highly Scalable Diffusion Samplers via Adjoint Matching, April 2025.
- [37] Pavlos S. Efraimidis and Paul G. Spirakis. Weighted random sampling with a reservoir. *Information Processing Letters*, 97(5):181–185, March 2006.
- [38] Mads-Peter V. Christiansen, Nikolaj Rønne, and Bjørk Hammer. Atomistic global optimization X: A Python package for optimization of atomistic structures. *The Journal of Chemical Physics*, 157(5):054701, August 2022.

A Appendix

A.1 Related work

Diffusion Models for Black-Box Optimization. Krishnamoorthy et al.[32] introduce Denoising Diffusion Optimization Models (DDOM), an inverse approach for black-box optimization that learns a conditional diffusion-based generative model mapping target function values to input configurations. It employs dataset reweighing and classifier-free guidance alongside a two-stage training approach.

Diffusion Model for Data-Driven Black-Box Optimization. Li et al.[33] propose a reward-directed conditional diffusion model trained on mixed unlabeled and labeled data. By conditioning on high predicted reward, they cast design optimization as conditional sampling.

Iterated Denoising Energy Matching for Sampling from Boltzmann Densities. Akhound-Sadegh et al.[34] introduce iDEM, a novel diffusion-based Boltzmann sampler trained using an energy-matching loss using Monte-Carlo samples to estimate the score. iDEM alternates between drawing samples from its current model and updating via an energy-matching loss. iDEM does not require prior data, and purely relies on energy and forces of the potential.

BNEM: A Boltzmann Sampler Based on Bootstrapped Noised Energy Matching. This recent follow-up[35] extends iDEM by leveraging bootstrapped energy estimates to improve sampling robustness and improve performance.

Adjoint Sampling: Highly Scalable Diffusion Samplers via Adjoint Matching. Havens et al.[36] propose an adjoint-matching based sampling scheme to train diffusion-based Boltzmann samplers. The key contribution alongside the reciprocal adjoint matching is the possibility to get many gradient updates with few potential energy evaluations and model samples. They show state-of-the-art performance on synthetic energy functions and difficult conformational sampling.

Unlike prior Boltzmann samplers such as iDEM[34] and BNEM[35], which estimate training targets from Monte Carlo generated samples, or Adjoint Sampling[36], which relies on adjoint matching to improve update efficiency, GO-Diff uses a direct Boltzmann-weighted score-matching loss requiring only energy evaluations — avoiding force labels or MC estimation. Combined with an annealed temperature schedule and replay-buffer self-sampling, this yields a simpler and very sample-efficient training loop. Furthermore, while previous samplers have not demonstrated transfer across systems, we show that pretrained GO-Diff models can be reused for amortized global optimization achieving faster optimization.

A.2 Score-Based Diffusion Models

Score-based diffusion models[15] are a class of generative models that learn to reverse a diffusion process that progressively adds noise to data. These models are grounded in stochastic differential equations (SDEs) and denoising score matching.

The forward process is defined as a stochastic differential equation (e.g., a variance-preserving or variance-exploding SDE) that transforms a data sample $\mathbf{x}_0 \sim p_{\text{data}}$ into a noisy sample $\mathbf{x}_t$ over time $t \in [0, 1]$. The generative goal is to model the reverse-time dynamics using a parameterized score function $s_\theta(\mathbf{x}_t, t) \approx \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$.

The model is trained using score matching, which minimizes the expected squared difference between added noise to samples and the prediction hereof. The learning objective is expressed as:

$$\mathcal{L}_\theta = \mathbb{E}_{t \sim \mathcal{U}(0,1)} \left[\lambda(t) \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}, \mathbf{x}_t \sim p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} \left\| s_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) \right\|_2^2 \right], \quad (8)$$

where $\lambda(t)$ is a weighting function depending on the noise schedule, and $p_{t|0}$ denotes the marginal of the forward SDE conditioned on the initial data point.

Once trained, new samples are generated by solving the reverse-time SDE using numerical solvers such as the Euler-Maruyama samplers.

For a complete description of applying diffusion models to the materials domain see Ref. [27] and specifically for the application to surface-supported systems see Ref. [21].

A.3 Buffer update

We maintain the buffer by applying weighted reservoir sampling [37] over all evaluated structures, using Boltzmann weights (at the current temperature) as sampling probabilities. This yields a dynamic buffer that gradually shifts focus toward lower-energy configurations as training progresses.

A.4 Training algorithm

Below we present pseudocode for the full GO-Diff optimization loop:

Algorithm 1 GO-Diff Optimization Procedure

```

1: Initialize: buffer  $\mathcal{B} \leftarrow \emptyset$ , score model  $s_\theta$ , temperature schedule  $\{T_k\}_{k=1}^K$ 
2: for iteration  $k = 1, \dots, K$  do
3:   Set temperature  $T \leftarrow T_k$ 
4:   Sample  $N$  structures  $\{X^{(j)}\}_{j=1}^N$  from  $p_\theta(\mathbf{x}_0)$  via reverse SDE
5:   for each structure  $X^{(j)}$  do
6:     Evaluate energy  $E^{(j)} \leftarrow U(X^{(j)})$ 
7:   end for
8:   Update buffer  $\mathcal{B}$  with  $\{(X^{(j)}, E^{(j)})\}$ 
9:   Train  $s_\theta$  using  $\mathcal{L}_\theta^{\text{Boltzmann}}$  on  $\mathcal{B}$ 
10: end for

```

A.5 Hyperparameters

We provide common and specific hyperparameters for all experiments.

Table 2: Common hyperparameters used in GO-Diff experiments.

Parameter	Value
Diffusion sampling steps	500
Noise schedule	Linear (VE-SDE)
Score model architecture	PaiNN GNN (4 blocks); 6Å cutoff
Final temperature T_K	0.02
Learning rate	10^{-4}
Optimizer	AdamW

Table 3: Hyperparameters used in **Pt-addatom on stepped Pt surface** experiment.

Parameter	Value
Buffer size $ \mathcal{B} $	16
Samples per iteration N	32
Initial temperature T_1	5.0
Annealing schedule	Exponential decay over 10 iterations
Batch size	8
Training epochs per iteration	1000

The **Pt-heptamer on Pt(111) + FFG** experiment uses same settings as without FFG and with $\zeta = 3$. The step size α is found using L-BFGS with a scaling factor of 0.2. The number of reverse diffusion steps is kept at 500, but this effectively doubles the number of model predictions during sampling due to the direct force predictions for the corrector step.

Table 4: Hyperparameters used in **Pt-heptamer on Pt(111)** experiment.

Parameter	Value
Buffer size $ \mathcal{B} $	64
Samples per iteration N	128
Initial temperature T_1	5.0
Annealing schedule	Exponential decay over 20 iterations
Batch size	16
Training epochs per iteration	2000

Table 5: Hyperparameters used in **Pt-heptamer on Pt(111) + transfer** experiment.

Parameter	Value
Buffer size $ \mathcal{B} $	64
Samples per iteration N	64
Initial temperature T_1	3.0
Annealing schedule	Exponential decay over 20 iterations
Batch size	16
Training epochs per iteration	4000

A.6 RSS

RSS is performed using the software package AGOX[38] following their documentation.

A.7 Compute resources

All experiments were run on single NVIDIA SM3090 GPU with 24GB of memory.

A.8 Code availability

All details to reproduce the experiments are provided at <https://github.com/nronne/go-diff> including specific implementation details and hyperparameters. The diffusion model is implemented in the AGeDi software package[27].