Generative Pseudo-Force Fields for Structure Generation

S. S. P. Hessmann[1,2], K. Kahouli[1,2], S. Gugler[1,2], M. Plainer[1,3,4], N. W. A. Gebauer[1,2]
[1]Machine Learning Group, Technische Universität Berlin, 10587 Berlin, Germany
[2]Berlin Institute for the Foundations of Learning and Data (BIFOLD), 10587 Berlin, Germany
[3]Freie Universität Berlin, 14195 Berlin, Germany
[4]Konrad Zuse School of Excellence in Learning and Intelligent Systems (ELIZA), Berlin, Germany
stefaan.hessmann@tu-berlin.de

The discovery and design of molecules and materials is a fundamental challenge in computational chemistry, with applications ranging from drug design to materials science, including the search for novel battery materials, catalysts, and functional nanomaterials. Most design methods target equilibrium structures, i.e. local minima on the Born–Oppenheimer potential energy surface (PES). Traditionally, these are found via structure relaxation, which iteratively adjusts atomic positions along interatomic forces. However, each relaxation trajectory requires many expensive quantum mechanical force calculations, making large-scale exploration infeasible. Machine learning force fields (MLFFs) [1] reduce this cost by orders of magnitude, but require training data with accurate force labels for non-equilibrium geometries, which remains challenging to obtain. Generative models offer an alternative by learning distributions of equilibrium structures without needing force labels. Among these, diffusion-based models [2] have achieved state-of-the-art results in molecular structure generation, but they require many sampling steps and are restricted to fixed noise schedules that depend on explicit time-step information. While recent work has shown that the noise level can be estimated with a neural network [3], allowing generation from close to equilibrium structures, GPFF goes further by enabling sampling from arbitrary, user-designed priors that were not used during training using a direct equilibrium predictor. In this work, we propose the generative pseudo-force field (GPFF), a generative model that bridges the concepts of traditional structure relaxation and diffusion-based generation. We define a simple quadratic pseudo-PES, $E(X|X_0) = \|X - X_0\|^2$, whose minima coincide with known equilibrium structures X_0, and derive corresponding pseudo-forces $F(X|X_0) = 2(X_0 - X)$. Training data is generated by perturbing equilibrium geometries via the one-step forward noising kernel of variance exploding (VE) diffusion, while pseudo-force labels are computed on-the-fly at negligible cost. Crucially, GPFF is trained without the diffusion time step as input, making it time-step-agnostic. The quadratic form of the pseudo-PES further yields a closed-form equilibrium predictor from any point, which is applied iteratively during sampling with optional stochastic noise injection. We show that the pseudo-forces relate exactly to the VE diffusion score function via $s(X,t) = F(X) / (2\sigma^2)$, establishing GPFF as a time-step-independent variant of VE diffusion models. Moreover, the one-step forward noising kernel corresponds to the Boltzmann distribution of our pseudo-PES, where the noise variance maps directly to physical temperature, providing a thermodynamically motivated interpretation of the diffusion process. The time-step independence of GPFF unlocks several practical advantages over standard diffusion models. First, it allows the design of informed prior distributions that incorporate domain knowledge, rather than being restricted to isotropic Gaussian noise. For instance, geometric properties such as covariance of the generated molecules can be enforced directly through the prior. Second, GPFF naturally supports scaffold sampling, where part of a molecular structure is predefined and only the remaining degrees of freedom are generated. We demonstrate that equilibrium prediction with GPFF achieves state-of-the-art validity of generated molecules. Notably, the predicted equilibrium structures remain close to the prior geometries, which implies that priors can be tailored to obey specific target properties, effectively steering generation towards desired regions of chemical space. On the other hand, when sampling with established diffusion samplers such as the reverse-SDE of Song et al. [4] or Heun's second-order method [5], the same trained GPFF model recovers the proper target distribution. This reveals a dual character of GPFF: a single model supports both a relaxation mode, where generated structures stay close to designed priors, and a diffusion mode, where the full equilibrium distribution is sampled. Looking ahead, these capabilities potentially make GPFF a compelling framework for materials design under periodic boundary conditions, where informed priors constructed from known structural motifs could guide the generation of stable crystal structures.

References

[1] Schütt, Kristof T., et al. "Schnet–a deep learning architecture for molecules and materials." The Journal of chemical physics 148.24 (2018).
[2] Hoogeboom, Emiel, et al. "Equivariant diffusion for molecule generation in 3d." International conference on machine learning. PMLR, 2022.
[3] Kahouli, Khaled, et al. "Molecular relaxation by reverse diffusion with time step prediction." Machine Learning: Science and Technology 5.3 (2024): 035038.

[4] Song, Yang, et al. "Score-based generative modeling through stochastic differential equations." arXiv preprint arXiv:2011.13456 (2020).
[5] Karras, Tero, et al. "Elucidating the design space of diffusion-based generative models." Advances in neural information processing systems 35 (2022): 26565-26577.

Figures

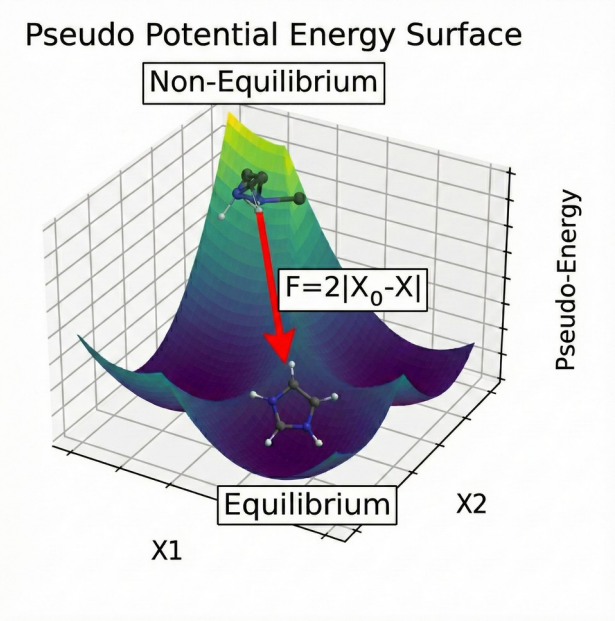

Figure 1. Quadratic pseudo-PES with minima at the equilibrium geometries X_0. The pseudo-force $F = 2(X_0 - X)$ drives non-equilibrium structures towards the minimum.