

ACTIVE LEARNING FOR MOLECULAR CONFORMATION OPTIMIZATION WITH A DOMAIN-AGNOSTIC NEURAL SURROGATE ORACLE

Anonymous authors

Paper under double-blind review

ABSTRACT

Molecular conformation optimization is crucial to computer-aided drug discovery and materials design, yet conventional force-based minimization with physics oracles (e.g., DFT) is prohibitively expensive. Neural network potentials (NNPs) are capable of accelerating this process but typically require large quantum chemical datasets for training. To reduce data requirements, active learning (AL) approaches have been designed for this task. The state-of-the-art approach, GOLF, relies on the surrogate oracle to sample new data. However, the surrogate oracle utilizes empirical molecular force fields, which may be absent for a current domain, and thus necessitates careful tuning. We introduce a new AL method for efficient conformation optimization that removes the dependency on empirical force fields. Our approach maintains two NNPs: an online NNP that performs conformation optimization and a target NNP that serves as a trainable surrogate oracle. The target network is an exponential moving average of the online network. During active sampling, the target NNP supplies potential energy estimates that guide data acquisition, while periodic queries to the physics oracle provide ground-truth corrections. Unlike other AL approaches, our method does not require architectural changes to NNP and adds minimal computational overhead compared to the single-model AL pipelines. Across two challenging conformation-optimization benchmarks (based on SPICE2.0 and $\nabla^2 DFT$ datasets) spanning different DFT levels, our method consistently outperforms a baseline NNP trained without AL, achieving substantial improvements with only 1,000 additional conformations.

1 INTRODUCTION

Molecular conformational optimization is a fundamental task in drug discovery and materials design that is used as a preprocessing step in various quantum chemistry (QC) pipelines (Pracht et al., 2020; Bursch et al., 2022). Conformational optimization is traditionally performed by iteratively minimizing the potential energy E using the interatomic forces F as anti-gradients until the local minima of the potential energy surface (PES) is reached. This process involves multiple expensive QC calculations with a physical oracle (in this work, we also call it a genuine oracle or $\mathcal{O}_G$) for energy and interatomic forces evaluation. A popular choice is to use the density functional theory-based genuine oracle. Notably, the computational cost of density functional theory (DFT) calculations scales at least cubically w.r.t. the number of electrons in the system, which limits its applicability to at best systems with several thousand electrons. Recently, deep neural network potentials (NNPs) (Khrabrov et al., 2024; Batatia et al., 2022; Gasteiger et al., 2021) have emerged as a promising alternative that can accurately approximate DFT-level energies and forces orders of magnitude faster.

Despite being trained on large quantities of QC data, NNPs are still prone to optimization problems such as unstable optimization or convergence to poor local minima: in practice, models trained on popular datasets can encounter a distribution shift during iterative structure relaxation (Tsypin et al., 2024; Khrabrov et al., 2024). The straightforward way to deal with this issue is to enrich the training dataset with optimization trajectories obtained with costly physical oracle calculations (Tsypin et al., 2024; Khrabrov et al., 2024; Fu et al., 2025). Although effective, this approach requires a lot of additional computations.

To overcome this issue, active learning methods are applied to efficiently collect additional data. The core idea behind this family of approaches is to identify a small subset of conformations where the model is most uncertain or likely to be wrong and add them to the training set of the NNP after evaluation with a physical oracle. Tsypin et al. (2024) demonstrated that the NNP can match the quality of the DFT-based oracle in the optimization task, given that the training set was enriched with a sufficient number of conformations from the optimization trajectories. They then proposed an active learning approach for efficient training of NNPs called GOLF, whose key idea is to use a cheap but inaccurate surrogate oracle $\mathcal{O}_S$, that is correlated with the genuine oracle $\mathcal{O}_G$, to decide which conformations are evaluated with $\mathcal{O}_G$ and added to the training dataset. Tsypin et al. (2024) considered molecules in vacuum and showed that utilizing the empirical MMFF94 (Halgren, 1996) force field as a surrogate oracle allows for successfully training an NNP that matches the quality of the DFT-based oracle while cutting down the required additional data by 50 times.

Although the active learning learning scheme proposed in GOLF proved to be efficient for the considered domain, a proper surrogate oracle must be carefully selected for each new domain. Moreover, for some complex domains, such as protein-ligand pairs (Faver et al., 2011; Yilmazer & Korth, 2013), organic liquids (Kovács et al., 2025; Kaminski & Jorgensen, 1996), or crystals (Kriz et al., 2023; Al Mamun et al., 2023; Zuo et al., 2020), it is challenging to find an empirical force field that correlates well with DFT-based methods or experimental results. This motivated us to explore active learning approaches that do not rely on empirical force fields. One common approach is to train an ensemble of NNP models and use their prediction variance as an uncertainty measure (Smith et al., 2018; Kulichenko et al., 2023; Carrete et al., 2023; Tan et al., 2023). While effective, ensemble methods significantly increase computational cost since multiple networks must be trained and evaluated in parallel. Alternatively, a single model can be equipped to predict its own uncertainty by modeling a distribution (mean and variance) over the target value instead of a point estimate (Tan et al., 2023; Xu et al., 2024). Such mean-variance estimation (MVE) networks provide uncertainty without multiple models, but they require architectural modifications that necessitate retraining the model to incorporate uncertainty estimation.

In this work, we propose a new approach to address the challenges with previous active learning approaches by introducing Neural Oracle. The core idea is to replace GOLF’s fixed surrogate oracle with a learned surrogate oracle that is updated in tandem with the NNP. Our method draws inspiration from the self-supervised learning BYOL framework (Grill et al., 2020). Similar to BYOL’s use of an online and target network, we maintain a secondary neural network (Neural Oracle) that is updated via a slow-moving average of the main model’s parameters. The Neural Oracle predicts energies and forces in a way that gradually aligns with the online NNP, without requiring direct supervision. This design provides a stable target for the NNP to compare against, effectively quantifying the model’s uncertainty. Unlike GOLF, which relies on an empirical surrogate, our Neural Oracle is adaptable to any chemical domain that can be processed by the NNP. What is more, our approach does not bring additional computational burden caused by training an ensemble of models and does not require any architectural changes or retraining the baseline NNP model.

We demonstrate the efficiency of the proposed GOLF Neural Oracle on a challenging ∇^2 DFT conformational optimization benchmark (Khrabrov et al., 2024), where it surpasses all baselines, including the state-of-the-art GOLF active learning framework. Notably, this superior performance is achieved in a highly data-efficient regime, utilizing only 1000 additional conformations. We then apply the best model to a significantly more challenging SPICE2.0 (Eastman et al., 2024) dataset. We then apply the best model to a significantly more challenging SPICE2.0 (Eastman et al., 2024) dataset. The SPICE2.0 not only contains data from various chemical domains, including dipeptides, solvated molecules, amino-acid ligand pairs, and water clusters, but also features a more advanced level of DFT theory: ω B97M-D3(BJ)/def2-TZVPPD. The triple basis makes the $\mathcal{O}_G$ calculations very costly, so we test our proposed approach in a data-efficient regime. We show that only 1000 additional conformations are enough to improve upon an already strong non-active learning baseline.

2 RELATED WORK

Molecular conformation optimization has been approached using machine learning methods, broadly categorized into models solving position regression tasks, generative methods, and direct optimization strategies.

Conformation generation Generative and position regression models aim to directly produce low-energy 3D conformers. Pioneering generative models often operated on internal coordinates, with early works utilizing normalizing flows (Xu et al., 2021) and torsional diffusion (Jing et al., 2022). Since diffusion models have become the dominant generative paradigm, now often operating directly in Cartesian coordinate space (Lee et al., 2024; Nikitin et al., 2025; Liu et al., 2025). Other frameworks like variational approximations and GFlowNets have also been actively explored (Volokhova et al., 2024). Another class of models frames the problem as a direct, one or few-shot regression task. These methods learn an end-to-end mapping from a given molecular representation to its ground-state 3D coordinates from a 2D molecular graph (Xu et al., 2023; Kim et al., 2025), or by refining a low-quality 3D structure (Lu et al., 2024; Wang et al., 2025). However, a common limitation of these generative and direct prediction models is that their outputs are not guaranteed to reside at a local minimum on the potential energy surface. Consequently, a subsequent relaxation using a physically-informed optimizer is typically necessary, motivating the development of iterative optimization strategies.

Optimization with NNPs NNPs Khrabrov et al. (2024); Batatia et al. (2022); Gasteiger et al. (2021) offer a computationally efficient alternative to traditional quantum mechanical methods for predicting molecular energies and forces. These potentials can then be used in conjunction with standard optimization algorithms, like BFGS, to perform geometry optimizations at a fraction of the computational cost (Liu et al., 2022; Tayfuroglu et al., 2025; Hao et al., 2022). The quality of NNP-based optimizations can be further improved by incorporating additional conformations into the training data, which helps to alleviate issues arising from distribution shifts (Tsy-pin et al., 2024; Khrabrov et al., 2024; Fu et al., 2025). Reinforcement learning has been employed to accelerate the optimization algorithm itself (Ahuja et al., 2021; Zamaraeva et al., 2025). In this work, we focus on improving NNPs optimization accuracy without extensive data generation, a problem that active learning is specifically designed to address.

Active learning for NNPs Active learning strategies, where the model requests additional calculations for uncertain regions of the conformational space, have also been employed to reduce the amount of data required to train accurate NNPs (Zhang et al., 2019; Kahle & Zipoli, 2022; Bilbrey et al., 2025; Mazitov et al., 2025). A “go-to” approach to uncertainty quantification (UQ) is to use the standard deviation of predictions from an ensemble of models. While generally robust, training and running ensembles incur a significant computational overhead (Smith et al., 2018; Kulichenko et al., 2023; Carrete et al., 2023; Bilbrey et al., 2025; Zhang et al., 2020; Schran et al., 2020). Consequently, substantial research has focused on developing cheaper, single-model UQ methods, such as mean-variance estimation (MVE) (Carrete et al., 2023; Xu et al., 2024), evidential regression (Amini et al., 2020), and Gaussian Mixture Models (GMMs) (Zhu et al., 2023). However, comprehensive benchmarks have shown that despite being faster, single-model methods do not consistently outperform the robustness of ensembles (Tan et al., 2023). Other uncertainty metrics, such as latent space distances, have also been explored (Musielewicz et al., 2024).

Furthermore, several works have specifically tailored active learning approaches for explicit search of transition states in reactive systems (Yang et al., 2021; Price et al., 2025) and the direct conformational optimization of both molecules and materials (Hessmann et al., 2025; Tsy-pin et al., 2024; Singh et al., 2024; Wang et al., 2024; Shuaibi et al., 2020). Compared to previous works, our research develops an active learning framework that is domain-agnostic, can be applied without architectural modification to the underlying NNP, and crucially minimizes the computational overhead of both model retraining and new training data generation.

3 NOTATION AND PRELIMINARIES

We define the conformation $s = \{z, \mathbf{X}\}$ of the molecule as a pair of atomic numbers $z = \{z_1, \dots, z_n\}$, $z_i \in \mathbb{N}$ and atomic coordinates $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, $\mathbf{x}_i \in \mathbb{R}^3$, where n is the number of atoms in the molecule. We define the oracle $\mathcal{O}$ as a function that takes conformation s as an input and outputs its potential energy $E_s^{\text{oracle}} \in \mathbb{R}$ and interatomic forces $\mathbf{F}_s^{\text{oracle}} \in \mathbb{R}^{n \times 3}$: $E_s^{\text{oracle}}, \mathbf{F}_s^{\text{oracle}} = \mathcal{O}(s)$. To denote the ground truth interatomic force acting on the i -th atom, we use $F_{s,i}^{\text{oracle}}$. For example, we denote the Psi4-calculated energy as E_s^{DFT} .

3.1 NEURAL NETWORK POTENTIALS

In this work, we use equivariant DimeNet++ (Gasteiger et al., 2020) and GemNet-OC (Gasteiger et al., 2022) due to their strong performance in the ∇^2 DFT optimization benchmark (Khrabrov et al., 2024). We denote the NNP parametrized by weights θ that predicts potential energy as $f_\theta(s) : \{z, \mathbf{X}\} \rightarrow \mathbb{R}$. The forces are derived from the predicted energies by taking a partial derivative:

$$\mathbf{F}_\theta^i(s) = -\frac{\partial f_\theta(s)}{\partial \mathbf{x}_i}, \quad (1)$$

where $\mathbf{F}_\theta^i(s) \in \mathbb{R}^3$ is the force acting on the i -th atom as predicted by the NNP. We follow (Khrabrov et al., 2024) and use slightly different loss functions for DimeNet++ and GemNet-OC:

$$\mathcal{L}^{\text{DimeNet++}}(s, \mathcal{B}; \theta) = \frac{\rho_1}{|\mathcal{B}|} \sum_{s \in \mathcal{B}} \|E_s - f_\theta(s)\|^2 + \frac{\rho_2}{|\mathcal{B}|} \sum_{s \in \mathcal{B}} \frac{1}{n_s} \sum_{i=1}^{n_s} \|\mathbf{F}_s^i - \mathbf{F}_\theta^i(s)\|^2; \quad (2)$$

$$\mathcal{L}^{\text{GemNet-OC}}(s, \mathcal{B}; \theta) = \frac{\rho_1}{|\mathcal{B}|} \sum_{s \in \mathcal{B}} |E_s - f_\theta(s)| + \frac{\rho_2}{\sum_{s \in \mathcal{B}} n_s} \sum_{s \in \mathcal{B}} \sum_{i=1}^{n_s} \|\mathbf{F}_s^i - \mathbf{F}_\theta^i(s)\|^2, \quad (3)$$

where $\mathcal{B}$ is a batch of conformations, E_s and $\mathbf{F}_s$ are the ground truth energies and forces, ρ_1 and ρ_2 are energy and forces scale coefficients respectively, and n_s is the number of atoms in s .

To perform geometry optimization with an NNP, the optimizer **Opt** utilizes the forces $\mathbf{F}_\theta(s) \in \mathbb{R}^{n \times 3}$:

$$s_{\text{cur}} = s_{\text{prev}} + \mathbf{Opt}(\mathbf{F}_\theta(s_{\text{prev}})). \quad (4)$$

3.2 EVALUATION METRICS

To evaluate the quality of optimization with NNPs, we optimize conformations from the test dataset $\mathcal{D}_{\text{test}}$ until convergence or the step limit is reached. Each conformation s in the test dataset comes with a ground truth optimal conformation s_{opt} and its energy $E_{s_{\text{opt}}}^{\text{DFT}}$, calculated by performing relaxation with $\mathcal{O}_G$. Following (Tsypin et al., 2024), the quality of the NNP-optimization is evaluated with the averaged percentage of minimized energy:

$$\overline{\text{pct}}_{\text{final}} = 100\% * \frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{s \in \mathcal{D}_{\text{test}}} \frac{E_s^{\text{DFT}} - E_{s_{\text{final}}}^{\text{DFT}}}{E_s^{\text{DFT}} - E_{s_{\text{opt}}}^{\text{DFT}}} = 100\% * \frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{s \in \mathcal{D}_{\text{test}}} \text{pct}(s_t), \quad (5)$$

where s_{final} is the final state of the NNP optimization trajectory. Another metric is the residual energy in state s_{final} : $E^{\text{res}}(s_{\text{final}})$. It is calculated as the delta between $E_{s_{\text{final}}}^{\text{DFT}}$ and the optimal energy:

$$E^{\text{res}}(s_t) = E_{s_t}^{\text{DFT}} - E_{s_{\text{opt}}}^{\text{DFT}}; \quad (6)$$

Similar to $\overline{\text{pct}}_{\text{final}}$, this metric can also be aggregated over the evaluation dataset:

$$\overline{E^{\text{res}}}_{\text{final}} = \frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{s \in \mathcal{D}_{\text{test}}} E^{\text{res}}(s_{\text{final}}). \quad (7)$$

Generally accepted chemical precision is 1 kcal/mol (Helgaker et al., 2004). Thus, another important metric is the percentage of conformations for which the residual energy is less than chemical precision. We consider optimizations with such residual energies successful:

$$\text{pct}_{\text{success}} = \frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{s \in \mathcal{D}_{\text{test}}} I[E^{\text{res}}(s_{\text{final}}) < 1]. \quad (8)$$

Lastly, we track the percentage of “diverged” optimizations pct_{div} . The optimization is considered diverged if the $\mathcal{O}_G$ was not able to calculate the energy of s_{opt} or when the resulting energy $E_{s_{\text{final}}}^{\text{DFT}}$ is larger than the initial energy E_s^{DFT} . We treat the diverged optimization as unsuccessful when calculating the $\text{pct}_{\text{success}}$.

3.3 ACTIVE LEARNING

The core idea of active learning is the iterative refinement of the training set by incorporating new data from regions of configuration space where the model is most uncertain, followed by retraining the model on the augmented dataset. The search for such conformations can be performed using molecular dynamics (Kulichenko et al., 2023)) or biased molecular dynamics (Laio & Gervasio, 2008; Sutton et al., 2012). Two popular approaches for UQ are ensemble- and single model-based (Seung et al., 1992; Nix & Weigend, 1994; Schran et al., 2020; Lakshminarayanan et al., 2017).

An ensemble of models is composed of several NNPs that differ slightly (see Section 5 for details). When the predictions of all members are in close agreement, it indicates that the conformation lies within a region of configuration space that was well represented in the training set. In contrast, large discrepancies among the ensemble predictions signal an out-of-distribution conformation. The uncertainty is calculated as (Tan et al., 2023):

$$\rho_{\text{ensemble}}(s) = \frac{\sigma_E(s)}{\sqrt{n_s}}, \quad \sigma_E^2(s) = \frac{1}{\mathcal{M}-1} \sum_{i=1}^{\mathcal{M}} (\tilde{E}_i(s) - \bar{E}(s))^2. \quad (9)$$

Here, $\tilde{E}_i = f_{\theta_i}(s)$ is the energy predicted by the i -th ensemble member, and $\bar{E}(s) = \sum_{i=1}^{\mathcal{M}} \tilde{E}_i(s)$ is the ensemble-averaged energy. The variance $\sigma_E(s)$ is normalized with the square root of the number of atoms n_s .

One widely adopted single-model approach is the mean variance estimation technique (MVE), where training data is treated as Gaussian random variables, and the NNP is trained to predict mean and variance $\mathcal{N}(\mu, \sigma)$. Training is performed via maximum likelihood estimation (Tan et al., 2023):

$$\mathcal{L}^{\text{MVE}}(s, \mathcal{B}; \theta) = \frac{\rho_1}{|\mathcal{B}|} \sum_{s \in \mathcal{B}} \|E_s - f_{\theta}(s)\|^2 + \frac{\rho_2}{|\mathcal{B}|} \sum_{s \in \mathcal{B}} \frac{1}{2n_s} \sum_{i=1}^{n_s} \log(2\pi\sigma_{F_i}^2) + \frac{\|\mathbf{F}_s^i - \mathbf{F}_{\theta}^i(s)\|^2}{\sigma_{F_i}^2}, \quad (10)$$

where $\sigma_{F_i}^2$ denotes the predicted per-atom force variance. For MVE-based models, the uncertainty estimate is computed as:

$$\rho_{\text{MVE}} = \max_{i \in s} \sigma_{F_i}^2. \quad (11)$$

3.4 GOLF

The GOLF framework (Tsypin et al., 2024) features a unique data sampling scheme, where the conformation selection is not based on the model’s uncertainty. Instead, a conformation is added to the training dataset if the energy (according to a cheap surrogate oracle) has increased after the NNP optimization step (see Equation 4): $\tilde{E}_{\text{prev}} - \tilde{E}_{\text{cur}} < 0$, where $\tilde{E}_{\text{prev}}$, $\tilde{E}_{\text{cur}}$ are the energies predicted by the surrogate oracle in conformations s_{cur} , s_{prev} respectively. If this condition is met, the predicted forces in s_{prev} are considered incorrect and the state s_{prev} is added to the training dataset. Alternatively, small negative energy changes can be encountered near the local minima due to the nature of gradient-based optimization. We discuss this scenario in Section 4.

4 METHOD

Our proposed active learning framework improves on the original GOLF in several aspects. We replace the fixed domain-specific surrogate with a Neural Oracle g_{ϕ} . The oracle parameters are updated after each training epoch by exponential moving averaging of the online NNP weights θ :

$$\phi = \tau\phi + (1 - \tau)\theta. \quad (12)$$

This standard Polyak averaging scheme (Polyak & Juditsky, 1992) is widely used in reinforcement learning (Mnih et al., 2015; Lillicrap et al., 2015; Haarnoja et al., 2018; Guo et al., 2022) and self-supervised learning (Grill et al., 2020) to stabilize the target network. In our work, we utilize this feature to stabilize potential energy estimates and select additional conformations more efficiently.

Moreover, to further optimize active data acquisition, we do not add final conformations of converged (or finished by reaching step limit) trajectories to the training dataset, as the online

Algorithm 1 GOLF with Neural Oracle

Require: baseline dataset $\mathcal{D}_0$; baseline dataset subsample size D ; online NNP f_θ ; neural oracle g_ϕ ; genuine oracle $\mathcal{O}_G$; optimizer Opt ; optimization step limit T ; maximum number of negative energy changes M ; per-cycle additional conformations k ; total additional conformations K ; per-cycle epochs C ; batch size B ; replay buffer mix ratio $\eta \in [0, 1]$; EMA coefficient $\tau \in [0, 1]$

- 1: Initialize replay buffer $\mathcal{R} \leftarrow \emptyset$; initialize $\phi \leftarrow \theta$
- 2: **for** $i \in 1, \dots, \lceil K/k \rceil$ **do** ▷ Runs for $\lceil K/k \rceil$ cycles
- 3: $S \leftarrow \emptyset$
- 4: **while** $|S| < k$ **do** ▷ Additional conformations collection
- 5: Sample $s \sim \mathcal{D}$; $t \leftarrow 0$; $m \leftarrow 0$; $\tilde{E}_{\text{prev}} \leftarrow g_\phi(s)$; $s_{\text{prev}} \leftarrow s$
- 6: **repeat**
- 7: $s' \leftarrow s + \text{Opt}(\mathbf{F}_{f_\theta}(s))$ ▷ Forces predicted with online NNP
- 8: $\tilde{E}_{\text{cur}} \leftarrow g_\phi(s')$
- 9: **if** $\tilde{E}_{\text{cur}} - \tilde{E}_{\text{prev}} > 0$ **then** ▷ Mistake: negative energy change
- 10: $m \leftarrow m + 1$
- 11: **end if**
- 12: $s_{\text{prev}} \leftarrow s$; $s \leftarrow s'$; $\tilde{E}_{\text{prev}} \leftarrow \tilde{E}_{\text{cur}}$; $t \leftarrow t + 1$
- 13: **until** $\text{Converged}(s)$ **or** $t \leq T$ **or** $m < P$ ▷ Optimize until convergence or TL reached
- 14: **if** $m = M$ **then**
- 15: $S \xleftarrow{\text{add}} s_{\text{prev}}$ ▷ Encountered M mistakes! Add previous conformation to $\mathcal{R}$
- 16: **else**
- 17: **discard** and resample
- 18: **end if**
- 19: **end while**
- 20: **for all** $s \in S$ **do** ▷ Evaluate collected conformations with DFT
- 21: $E_{\text{DFT}}(s), \mathbf{F}_{\text{DFT}}(s) \leftarrow \mathcal{O}_G(s)$
- 22: $\mathcal{R} \xleftarrow{\text{add}} \{s, E_{\text{DFT}}(s), \mathbf{F}_{\text{DFT}}(s)\}$
- 23: **end for**
- 24: Sample shard $\mathcal{D}_i \subset \mathcal{D}_0$ with $|\mathcal{D}_i| = D$ ▷ Training
- 25: **for** $e = 1$ to C **do**
- 26: **for** batch $\in \mathcal{D}_i$ **do**
- 27: Sample $(1 - \eta)B$ from $\mathcal{D}_j$ and ηB from $\mathcal{R}$ to form batch $\mathcal{B}$
- 28: Update f_θ on $\mathcal{B}$ using Eq. 2,3
- 29: **end for**
- 30: $\phi \leftarrow \tau \phi + (1 - \tau) \theta$ ▷ EMA of neural oracle
- 31: **end for**
- 32: **end for**

NNP already performs well for such conformations. The GOLF method uses the energy change $\tilde{E}_{\text{prev}} - \tilde{E}_{\text{cur}}$ after the optimization step to select additional conformations. However, due to the nature of the selected optimization algorithm, the energy can sometimes increase near the local minima. This does not necessarily mean that the forces are poorly predicted, but with GOLF’s data acquisition scheme, such conformation will still be added to the training dataset, hindering the efficiency of the approach. To minimize the amount of near-optimal non-informative conformations in the training set, we introduce a “mistake budget” M . The conformation is only added to the training dataset if there have been M or more negative energy changes in the trajectory.

The training of GOLF Neural Oracle alternates between data collection and supervised updates of the online NNP. In each collection phase, we build a batch S of k conformations. Starting from $s_0 \sim \mathcal{D}$, we iteratively optimize the conformation using predicted forces $\mathbf{F}_\theta(s)$ and estimate potential energy with a Neural Oracle:

$$s_{\text{cur}} = s_{\text{prev}} + \text{Opt}(\mathbf{F}_\theta(s_{\text{prev}})), \quad \tilde{E}_{\text{cur}} = g_\phi(s_{\text{cur}}). \quad (13)$$

We maintain a *mistake counter* m that increments whenever the energy change according to the neural oracle is negative,

$$\tilde{E}_{\text{prev}} - \tilde{E}_{\text{cur}} < 0 \implies m \leftarrow m + 1, \quad (14)$$

Table 1: Evaluation on the test set $\mathcal{D}_{\text{test}}$. Columns report the averaged percentage of minimized energy $\overline{\text{pct}}_{\text{final}}$, the mean residual energy $\overline{E}_{\text{final}}^{\text{res}}$ in kcal/mol, the success rate $\text{pct}_{\text{success}}$ (fraction with $E^{\text{res}} < 1$ kcal/mol), and the diverged conformations percentage pct_{div} . Arrows indicate preferred direction ($\uparrow$ higher is better; $\downarrow$ lower is better). Best results are in bold.

Model	$\overline{\text{pct}}_{\text{final}} \uparrow$ (%)	$\overline{E}_{\text{final}}^{\text{res}} \downarrow$ (kcal/mol)	$\text{pct}_{\text{success}} \uparrow$ (%)	$\text{pct}_{\text{div}} \downarrow$ (%)
DimeNet++ baseline	96.21	1.67	41.45	0.85
<i>Active Learning Models</i>				
Ensemble	99.99	0.001	87.6	0.25
MVE	99.86	0.096	87.4	0.85
GOLF-RDKit	100.13	-0.04	89.05	0.55
GOLF-Neural Oracle (ours)	100.24	-0.092	90.5	0.15

and we terminate a trajectory on optimizer convergence, a step limit T , or when $m = M$. If $m = M$, the previous conformation s_{prev} is added to S . When enough data is collected ($|S| = k$), conformations are evaluated with $\mathcal{O}_G$ and added to the replay buffer $\mathcal{R}$.

After collection, the online NNP f_{θ} is trained for a small number of epochs using minibatches that mix baseline data from $\mathcal{D}_0$ and additional data $\mathcal{R}$. With mixture coefficient $\eta \in [0, 1]$, each minibatch $\mathcal{B}$ contains a fraction $(1 - \eta)$ drawn from $\mathcal{D}_0$ and a fraction η drawn from $\mathcal{R}$. The training objective is defined in Equations 2 and 3. At the end of each epoch, the Neural Oracle is updated via the EMA rule in Equation 12, and the loop returns to the batched collection phase. Over $\lceil K/M \rceil$ cycles, this procedure acquires K DFT-labeled conformations. The full training procedure is detailed in Algorithm 1.

5 EXPERIMENTS

5.1 ∇^2 DFT OPTIMIZATION BENCHMARK

We benchmark the proposed GOLF-Neural Oracle on ∇^2 DFT (Khrabrov et al., 2024) benchmark. This benchmark evaluates the optimization performance of NNPs on a subset of 2000 molecules from the ∇^2 DFT and utilizes ω B97X-D/def2-SVP level of DFT theory. This benchmark only contains molecules in vacuum, which makes direct comparison with GOLF-RDKit possible.

As baselines, we consider two classical active learning approaches: ensemble and MVE (Tan et al., 2023); the GOLF-RDKit method; and the baseline NNP trained on a fixed dataset without active learning finetuning. The DimeNet++ is used as the NNP for all iterative optimization methods. We use the $\mathcal{D}^{\text{medium}}$ subset of ∇^2 DFT to train the baseline NNP and as $\mathcal{D}_0$ for active learning approaches.

For all active learning approaches, we use the same data collection scheme and training hyperparameters except for the conformation selection criterion. The number of additional conformations K equals to 10000. The number of additional conformations per-cycle is $k = 200$, and the number of epochs per cycle is $C = 5$, resulting in a total of 250 training epochs. The initial conformations are sampled from $\mathcal{D}_0$ and then optimized with the L-BFGS (Liu & Nocedal, 1989) algorithm (as implemented in Pytorch Paszke et al. (2019)) using the forces predicted by the NNP (see Equation 13). The training step limit is $T = 100$. We set $\eta = 0.5$, so the batch consists of conformation from $\mathcal{D}_0$ and $\mathcal{R}$ in equal parts.

In the case of the ensemble, the forces are averaged over ensemble members. For the ensemble, we used $\mathcal{M} = 4$ DimeNet++ NNPs. Following (Tan et al., 2023), each NNP was trained on $\mathcal{D}^{\text{medium}}$ using identical hyperparameters, but with different random seeds. The criterion for conformation selection is $\rho_{\text{ensemble}} > t_{\text{ensemble}} = 2.8 \times 10^{-4}$. We found this threshold to strike a good balance between the model’s performance and data collection time. For MVE, we modified the final layers of DimeNet++ so that the network outputs the conformational energy along with the per-atom force variance $\sigma_{F_i}^2$. To ensure the positivity of variance values, a soft-plus activation function was applied

Table 2: Ablation on the mistake budget M (maximum negative energy changes per trajectory) and the EMA coefficient τ used to update the Neural Oracle. Arrows indicate preferred direction.

Method	τ	M	$\overline{\text{pct}}_{\text{final}} \uparrow$ (%)	$\overline{E}_{\text{final}}^{\text{res}} \downarrow$ (kcal/mol)	$\text{pct}_{\text{success}} \uparrow$ (%)	$\text{pct}_{\text{div}} \downarrow$ (%)
GOLF-RDKit	–	1	100.13	-0.04	89.05	0.55
	–	3	100.2	-0.073	89.35	0.75
GOLF-Neural Oracle	0.9	1	100.11	-0.035	89.0	0.3
	0.9	3	100.24	-0.092	90.5	0.15
	0.7	3	100.18	-0.061	89.5	0.4
	0.5	3	100.18	-0.061	89.55	0.35

to the corresponding outputs. The DimeNet++ for MVE was pretrained on $\mathcal{D}_{\text{medium}}$ using loss defined in Equation 10. The same DimeNet++ modification and loss function were used during the active learning phase. The criterion for conformation selection is $\rho_{\text{MVE}} > t_{\text{MVE}} = 1 \times 10^{-5}$.

All iterative optimization baselines featured the same evaluation procedure: the conformations from the $\mathcal{D}^{\text{test}}$ were optimized with L-BFGS until convergence or step limit $T_{\text{eval}} = 200$. The performance on the $\nabla^2\text{DFT}$ optimization benchmark is in Table 1. The active learning approaches surpass the direct optimization model and the baseline DimeNet++. Moreover, the GOLF approaches outperform the ensemble and MVE. We hypothesize this is because the conformation selection criterion in GOLF merely depends on the model’s uncertainty, but instead is specifically tailored for the optimization.

5.2 GOLF ABLATIONS

In this section, we study how hyperparameters τ and M affect the models. The results are provided in Table 2. The “mistake budget” $M > 1$ improves the performance of the GOLF-Neural Oracle, while not significantly affecting the GOLF-RDKit. Despite the fact that energies predicted by MMFF94 (used as a surrogate oracle in GOLF-RDKit) correlate with the DFT energies, the local minima of these functionals differ significantly. Therefore, the MMFF94-predicted energies near the DFT local minima can behave arbitrarily, so the increased M does not help to better select additional training conformations.

For the GOLF-Neural Oracle, τ regulates the rate of the EMA update (see Equation 12). The lower the update rate, the more Neural Oracle resembles the online NNP, and, conversely, a higher update rate leads to a more stable Neural Oracle. We found that reasonably high values of τ lead to stable energy estimation and better additional conformation selection.

Table 3: Data-efficient regime. We report the number of additional conformations K used for training and evaluate on $\mathcal{D}_{\text{test}}$ with the same metrics as before. Arrows indicate preferred direction

Model	K (add. confs)	$\overline{\text{pct}}_{\text{final}} \uparrow$ (%)	$\overline{E}_{\text{final}}^{\text{res}} \downarrow$ (kcal/mol)	$\text{pct}_{\text{success}} \uparrow$ (%)	$\text{pct}_{\text{div}} \downarrow$ (%)
MVE	1000	99.03	0.437	76.7	0.55
GOLF-Neural Oracle (ours)	1000	99.19	0.376	81.45	0.45

5.3 DATA EFFICIENT REGIME

Additionally, we test our proposed method in a “data-efficient” regime, where only 1000 additional conformations are collected during the active learning finetuning. This is especially useful when training with an expensive $\mathcal{O}_G$. For the data efficient regime, we use $K = 1000$, $k = 20$, $C = 1$ and keep other hyperparameters unchanged. The results in Table 3 indicate that our proposed approach remains effective in a data-efficient regime and outperforms other classical active learning approaches.

Table 4: Evaluation on the test set $\mathcal{D}_{\text{SPICE}}^{\text{test}}$. Columns report the averaged percentage of minimized energy $\overline{\text{pct}}_{\text{final}}$, the mean residual energy $\overline{E}_{\text{final}}^{\text{res}}$ in kcal/mol, the success rate $\text{pct}_{\text{success}}$ (fraction with $E^{\text{res}} < 1$ kcal/mol), and the diverged conformations percentage pct_{div} . Arrows indicate preferred direction ($\uparrow$ higher is better; $\downarrow$ lower is better). Best results are in bold.

Model	$\overline{\text{pct}}_{\text{final}} \uparrow$ (%)	$\overline{E}_{\text{final}}^{\text{res}} \downarrow$ (kcal/mol)	$\text{pct}_{\text{success}} \uparrow$ (%)	$\text{pct}_{\text{div}} \downarrow$ (%)
Nutmeg–Large	76.60	11.17	3.67	5.45
MACE–OFF23 (Large)	98.27	1.04	64.11	2.79
GOLF–Neural Oracle (ours)	99.88	0.66	94.2	0.96

5.4 SPICE

To test both transferability to a more advanced level of DFT theory and to multi-domain chemical data, we benchmark the proposed GOLF-Neural Oracle on the SPICE2.0 dataset. To get the $\mathcal{D}_{\text{SPICE}}^{\text{test}}$, we optimized a subset of test set described in Eastman et al. (2024) (see Appendix A.1 for details). This resulted in $|\mathcal{D}_{\text{SPICE}}^{\text{test}}| = 216$ conformations optimized at $\omega\text{B97M-D3(BJ)}/\text{def2-TZVPPD}$ level of theory. As baselines, we selected the Nutmeg-large model (Eastman et al., 2024) and the MACE-OFF23 (Large) Kovács et al. (2025) as they have already been pretrained on the full SPICE2.0 dataset. We then trained GemNet-OC on the full SPICE dataset and finetuned it with GOLF-Neural Oracle.

As the $\mathcal{D}_{\text{SPICE}}^{\text{test}}$ consists of large ligand molecules, pentapeptides and amino-acid ligand dimers, we filter out molecular systems with less than 40 atoms from the SPICE2.0 dataset and call this filtered dataset $\mathcal{D}_0^{\text{filtered}}$. Hyperparameters specified in Section 5.3 were used to train the GOLF-Neural Oracle on $\mathcal{D}_0^{\text{filtered}}$. The results in Table 4 demonstrate that GOLF-Neural Oracle successfully optimizes almost all molecules from $\mathcal{D}_{\text{SPICE}}^{\text{test}}$, significantly surpassing MACE-OFF23 which was trained without the active learning.

6 CONCLUSION

In this work, we have presented a new framework called GOLF with Neural Oracle for molecular conformation optimization learning. We show that a trained neural Oracle can successfully replace a cheap physical simulator, and help the final model achieve a quality comparable with an expensive physical simulator. We thoroughly compare our approach with several baselines, including recent conformation generation models and an adaptation of other active learning schemes. A primary direction for future work is to apply our framework to larger systems and systems dominated by intermolecular interactions, such as optimizing ligand conformations within protein binding pockets or relaxing adsorbates on surfaces. Furthermore, we plan to test the scalability of the Neural Oracle for condensed-phase systems, including liquids and solids under periodic boundary conditions. A particularly promising extension will be the integration of our active learning approach into a hybrid ML/MM (Machine Learning / Molecular Mechanics) framework. This would allow for the data-efficient optimization of a high-accuracy NNP for a reactive site while treating the surrounding environment with a classical force field, enabling the study of chemical processes in large and complex biological systems.

REPRODUCIBILITY STATEMENT

We will soon release the code used to train all active learning approaches in this study.

REFERENCES

Kabir Ahuja, William H Green, and Yi-Pei Li. Learning to optimize molecular geometries using reinforcement learning. *Journal of Chemical Theory and Computation*, 17(2):818–825, 2021.

- Abdullah Al Mamun, Shuozhi Xu, Xiang-Guo Li, and Yanqing Su. Comparing interatomic potentials in calculating basic structural parameters and peierls stress in tungsten-based random binary alloys. *Physica Scripta*, 98(10):105923, 2023.
- Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. *Advances in neural information processing systems*, 33:14927–14937, 2020.
- Ilyes Batatia, David P Kovacs, Gregor Simm, Christoph Ortner, and Gábor Csányi. Mace: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in neural information processing systems*, 35:11423–11436, 2022.
- Jenna A Bilbrey, Jesun S Firoz, Mal-Soon Lee, and Sutanay Choudhury. Uncertainty quantification for neural network potential foundation models. *npj Computational Materials*, 11(1):109, 2025.
- Markus Bursch, Jan-Michael Mewes, Andreas Hansen, and Stefan Grimme. Best-practice dft protocols for basic molecular computational chemistry. *Angewandte Chemie*, 134(42):e202205735, 2022.
- Jesús Carrete, Hadrián Montes-Campos, Ralf Wanzenböck, Esther Heid, and Georg KH Madsen. Deep ensembles vs committees for uncertainty estimation in neural-network force fields: Comparison and application to active learning. *The Journal of Chemical Physics*, 158(20), 2023.
- Peter Eastman, Benjamin P Pritchard, John D Chodera, and Thomas E Markland. Nutmeg and spice: models and data for biomolecular machine learning. *Journal of chemical theory and computation*, 20(19):8583–8593, 2024.
- John C Faver, Mark L Benson, Xiao He, Benjamin P Roberts, Bing Wang, Michael S Marshall, Matthew R Kennedy, C David Sherrill, and Kenneth M Merz Jr. Formal estimation of errors in computed absolute interaction energies of protein- ligand complexes. *Journal of chemical theory and computation*, 7(3):790–797, 2011.
- Cong Fu, Yuchao Lin, Zachary Krueger, Wendi Yu, Xiaoning Qian, Byung-Jun Yoon, Raymundo Arróyave, Xiaofeng Qian, Toshiyuki Maeda, Maho Nakata, et al. A benchmark for quantum chemistry relaxations via machine learning interatomic potentials. *arXiv preprint arXiv:2506.23008*, 2025.
- Johannes Gasteiger, Shankari Giri, Johannes T Margraf, and Stephan Günnemann. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. *arXiv preprint arXiv:2011.14115*, 2020.
- Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. *Advances in Neural Information Processing Systems*, 34:6790–6802, 2021.
- Johannes Gasteiger, Muhammed Shuaibi, Anuroop Sriram, Stephan Günnemann, Zachary Ward Ulissi, C. Lawrence Zitnick, and Abhishek Das. Gemnet-OC: Developing graph neural networks for large and diverse molecular simulation datasets. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=u8tvSxm4Bs>.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- Zhaohan Guo, Shantanu Thakoor, Miruna Pîslar, Bernardo Avila Pires, Florent Altché, Corentin Tallec, Alaa Saade, Daniele Calandriello, Jean-Bastien Grill, Yunhao Tang, et al. Byol-explore: Exploration by bootstrapped prediction. *Advances in neural information processing systems*, 35: 31855–31870, 2022.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. Pmlr, 2018.

- Thomas A Halgren. Merck molecular force field. i. basis, form, scope, parameterization, and performance of mmff94. *Journal of computational chemistry*, 17(5-6):490–519, 1996.
- Dongxiao Hao, Xibing He, Adrian E Roitberg, Shengli Zhang, and Junmei Wang. Development and evaluation of geometry optimization algorithms in conjunction with ani potentials. *Journal of Chemical Theory and Computation*, 18(2):978–991, 2022.
- Trygve Helgaker, Torgeir A Ruden, Poul Jørgensen, Jeppe Olsen, and Wim Klopper. A priori calculation of molecular properties to chemical accuracy. *Journal of physical organic chemistry*, 17(11):913–933, 2004.
- Stefaan SP Hessmann, Kristof T Schütt, Niklas WA Gebauer, Michael Gastegger, Tamio Oguchi, and Tomoki Yamashita. Accelerating crystal structure search through active learning with neural networks for rapid relaxations. *npj Computational Materials*, 11(1):44, 2025.
- Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional diffusion for molecular conformer generation. *Advances in neural information processing systems*, 35:24240–24253, 2022.
- Leonid Kahle and Federico Zipoli. Quality of uncertainty estimates from neural network potential ensembles. *Physical Review E*, 105(1):015311, 2022.
- George Kaminski and William L Jorgensen. Performance of the amber94, mmff94, and opl-aa force fields for modeling organic liquids. *The Journal of Physical Chemistry*, 100(46):18010–18013, 1996.
- Kuzma Khrabrov, Anton Ber, Artem Tsylin, Konstantin Ushenin, Egor Rumiantsev, Alexander Telepov, Dmitry Protasov, Ilya Shenbin, Anton M. Alekseev, Mikhail Shirokikh, Sergey Nikolenko, Elena Tutubalina, and Artur Kadurin. ∇^2 SDFT: A universal quantum chemistry dataset of drug-like molecules and a benchmark for neural network potentials. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL <https://openreview.net/forum?id=ElUrNM9U8c>.
- Taewon Kim, Hyunjin Seo, Sungsoo Ahn, and Eunho Yang. Rebind: Enhancing ground-state molecular conformation prediction via force-based graph rewiring. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Dávid Péter Kovács, J Harry Moore, Nicholas J Browning, Ilyes Batatia, Joshua T Horton, Yixuan Pu, Venkat Kapil, William C Witt, Ioan-Bogdan Magdau, Daniel J Cole, et al. Mace-off: Short-range transferable machine learning force fields for organic molecules. *Journal of the American Chemical Society*, 147(21):17598–17611, 2025.
- Kristian Kriz, Lisa Schmidt, Alfred T Andersson, Marie-Madeleine Walz, and David van der Spoel. An imbalance in the force: the need for standardized benchmarks for molecular simulation. *Journal of Chemical Information and Modeling*, 63(2):412–431, 2023.
- Maksim Kulichenko, Kipton Barros, Nicholas Lubbers, Ying Wai Li, Richard Messerly, Sergei Tretiak, Justin S Smith, and Benjamin Nebgen. Uncertainty-driven dynamics for active learning of interatomic potentials. *Nature computational science*, 3(3):230–239, 2023.
- Alessandro Laio and Francesco L Gervasio. Metadynamics: a method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science. *Reports on Progress in Physics*, 71(12):126601, 2008.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Danyeong Lee, Dohoon Lee, Dongmin Bang, and Sun Kim. Disco: Diffusion schrödinger bridge for molecular conformer optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 13365–13373, 2024.

- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.
- Zhen Liu, Tetiana Zubatiuk, Adrian Roitberg, and Olexandr Isayev. Auto3d: Automatic generation of the low-energy 3d structures with ani neural network potentials. *Journal of Chemical Information and Modeling*, 62(22):5373–5382, 2022.
- Zhiyuan Liu, Yanchen Luo, Han Huang, Enzhi Zhang, Sihang Li, Junfeng Fang, Yaorui Shi, Xiang Wang, Kenji Kawaguchi, and Tat-Seng Chua. Next-mol: 3d diffusion meets 1d language modeling for 3d molecule generation. *arXiv preprint arXiv:2502.12638*, 2025.
- Shuqi Lu, Zhifeng Gao, Di He, Linfeng Zhang, and Guolin Ke. Data-driven quantum chemical property prediction leveraging 3d conformations with uni-mol+. *Nature communications*, 15(1):7104, 2024.
- Narbe Mardirossian and Martin Head-Gordon. ω b97m-v: A combinatorially optimized, range-separated hybrid, meta-gga density functional with vv10 nonlocal correlation. *The Journal of chemical physics*, 144(21), 2016.
- Arslan Mazitov, Ivan Kruglov, Alexey V Yanilkin, Aleksey V Arsenin, Valentyn S Volkov, Dmitry G Kvashnin, Artem R Oganov, and Kostya S Novoselov. Substrate-aware computational design of two-dimensional materials. *npj Computational Materials*, 11(1):270, 2025.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Joseph Musielewicz, Janice Lan, Matt Uyttendaele, and John R Kitchin. Improved uncertainty estimation of graph neural network potentials using engineered latent space distances. *The Journal of Physical Chemistry C*, 128(49):20799–20810, 2024.
- Filipp Nikitin, Dylan M Anstine, Roman Zubatyuk, Saeed Gopal Paliwal, and Olexandr Isayev. Scalable low-energy molecular conformer generation with quantum mechanical accuracy. *ChemRxiv preprint ChemRxiv:10.26434/chemrxiv-2025-k4h7v*, 2025.
- David A Nix and Andreas S Weigend. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE international conference on neural networks (ICNN'94)*, volume 1, pp. 55–60. IEEE, 1994.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.
- Philipp Pracht, Fabian Bohle, and Stefan Grimme. Automated exploration of the low-energy chemical space with fast quantum chemical methods. *Physical Chemistry Chemical Physics*, 22(14):7169–7192, 2020.
- Trevor Price, Saurabh Sivakumar, Matthew S Johnson, Judit Zádor, and Ambarish Kulkarni. Automated pynta-based curriculum for ml-accelerated calculation of transition states. *The Journal of Physical Chemistry C*, 129(16):7751–7761, 2025.
- Christoph Schran, Krystof Brezina, and Ondrej Marsalek. Committee neural network potentials control generalization errors and enable active learning. *The Journal of Chemical Physics*, 153(10), 2020.
- H Sebastian Seung, Manfred Oppen, and Haim Sompolinsky. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pp. 287–294, 1992.

- Muhammed Shuaibi, Saurabh Sivakumar, Rui Qi Chen, and Zachary W Ulissi. Enabling robust offline active learning for machine learning potentials using simple physics-based priors. *Machine Learning: Science and Technology*, 2(2):025007, 2020.
- Akksay Singh, Jiaqi Wang, Graeme Henkelman, and Lei Li. Uncertainty based machine learning-dft hybrid framework for accelerating geometry optimization. *Journal of Chemical Theory and Computation*, 20(22):10022–10033, 2024.
- Daniel GA Smith, Lori A Burns, Andrew C Simmonett, Robert M Parrish, Matthew C Schieber, Raimondas Galvelis, Peter Kraus, Holger Kruse, Roberto Di Remigio, Asem Alenaizan, et al. Psi4 1.4: Open-source software for high-throughput quantum chemistry. *The Journal of chemical physics*, 152(18), 2020.
- Justin S Smith, Ben Nebgen, Nicholas Lubbers, Olexandr Isayev, and Adrian E Roitberg. Less is more: Sampling chemical space with active learning. *The Journal of chemical physics*, 148(24), 2018.
- Ludovico Sutto, Simone Marsili, and Francesco Luigi Gervasio. New advances in metadynamics. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 2(5):771–779, 2012.
- Aik Rui Tan, Shingo Urata, Samuel Goldman, Johannes CB Dietschreit, and Rafael Gómez-Bombarelli. Single-model uncertainty quantification in neural network potentials does not consistently outperform model ensembles. *npj Computational Materials*, 9(1):225, 2023.
- Omer Tayfuroglu, Irem N Zengin, M Serdar Koca, and Abdulkadir Kocak. Deepconf: Leveraging ani-ml potentials for exploring local minima with application to bioactive conformations. *Journal of Chemical Information and Modeling*, 65(6):2818–2833, 2025.
- Artem Tsy-pin, Leonid Anatolievich Ugadiarov, Kuzma Khrabrov, Alexander Telepov, Egor Rumiantsev, Alexey Skrynnik, Aleksandr Panov, Dmitry P. Vetrov, Elena Tutubalina, and Artur Kadurin. Gradual optimization learning for conformational energy minimization. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=FMMFla9iFl>.
- Alexandra Volokhova, Michał Koziarski, Alex Hernández-García, Cheng-Hao Liu, Santiago Miret, Pablo Lemos, Luca Thiede, Zichao Yan, Alán Aspuru-Guzik, and Yoshua Bengio. Towards equilibrium molecular conformation generation with gflownets. *Digital Discovery*, 3(5):1038–1047, 2024.
- Fanmeng Wang, Minjie Cheng, and Hongteng Xu. Wgformer: An se (3)-transformer driven by wasserstein gradient flows for molecular ground-state conformation prediction. In *Forty-second International Conference on Machine Learning*, 2025.
- Xiaoxiao Wang, Joseph Musielewicz, Richard Tran, Sudheesh Kumar Ethirajan, Xiaoyan Fu, Hilda Mera, John R Kitchin, Rachel C Kurchin, and Zachary W Ulissi. Generalization of graph-based active learning relaxation strategies across materials. *Machine Learning: Science and Technology*, 5(2):025018, 2024.
- Florian Weigend and Reinhart Ahlrichs. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for h to rn: Design and assessment of accuracy. *Physical Chemistry Chemical Physics*, 7(18):3297–3305, 2005.
- Guikun Xu, Yongquan Jiang, PengChuan Lei, Yan Yang, and Jim Chen. Gtmgc: Using graph transformer to predict molecule’s ground-state conformation. In *The Twelfth International Conference on Learning Representations*, 2023.
- Han Xu, Taoyong Cui, Chenyu Tang, Jinzhe Ma, Dongzhan Zhou, Yuqiang Li, Xiang Gao, Xingao Gong, Wanli Ouyang, Shufei Zhang, et al. Evidential deep learning for interatomic potentials. *arXiv preprint arXiv:2407.13994*, 2024.
- Minkai Xu, Shitong Luo, Yoshua Bengio, Jian Peng, and Jian Tang. Learning neural generative dynamics for molecular conformation generation. *arXiv preprint arXiv:2102.10240*, 2021.

- Yilin Yang, Omar A Jiménez-Negrón, and John R Kitchin. Machine-learning accelerated geometry optimization in molecular simulation. *The Journal of Chemical Physics*, 154(23), 2021.
- Nusret Duygu Yilmazer and Martin Korth. Comparison of molecular mechanics, semi-empirical quantum mechanical, and density functional theory methods for scoring protein–ligand interactions. *The Journal of Physical Chemistry B*, 117(27):8075–8084, 2013.
- Elena Zamaraeva, Christopher M Collins, George R Darling, Matthew S Dyer, Bei Peng, Rahul Savani, Dmytro Antypov, Vladimir V Gusev, Judith Clymo, Paul G Spirakis, et al. Macs: Multi-agent reinforcement learning for optimization of crystal structures. *arXiv preprint arXiv:2506.04195*, 2025.
- Linfeng Zhang, De-Ye Lin, Han Wang, Roberto Car, and Weinan E. Active learning of uniformly accurate interatomic potentials for materials simulation. *Physical Review Materials*, 3(2):023804, 2019.
- Yuzhi Zhang, Haidi Wang, Weijie Chen, Jinzhe Zeng, Linfeng Zhang, Han Wang, et al. Dp-gen: A concurrent learning platform for the generation of reliable deep learning based potential energy models. *Computer Physics Communications*, 253:107206, 2020.
- Albert Zhu, Simon Batzner, Albert Musaelian, and Boris Kozinsky. Fast uncertainty estimates in deep learning interatomic potentials. *The Journal of Chemical Physics*, 158(16), 2023.
- Yunxing Zuo, Chi Chen, Xiangguo Li, Zhi Deng, Yiming Chen, Jörg Behler, Gábor Csányi, Alexander V Shapeev, Aidan P Thompson, Mitchell A Wood, et al. Performance and cost assessment of machine learning interatomic potentials. *The Journal of Physical Chemistry A*, 124(4):731–745, 2020.

A APPENDIX

A.1 OPTIMIZATION OF SPICE CONFORMATIONS

To generate the ground truth optimized conformations, we performed full geometry relaxations for each initial conformation from the test set of the SPICE dataset. These calculations were carried out using the Psi4 computational chemistry package (Smith et al., 2020). The optimizations were driven by the optking optimizer, employing the ω B97M-D3BJ (Mardirossian & Head-Gordon, 2016) functional with the def2-TZVPPD (Weigend & Ahlrichs, 2005) basis set to remain consistent with the original level of theory used in the SPICE methodology (Eastman et al., 2024). Each optimization was run until standard convergence criteria were met, defined as the maximum force component on any atom falling below 3×10^{-4} Hartree/Bohr, energy change less than 3×10^{-6} Hartree, and maximum atom displacement less than 1.2×10^{-3} Bohr. Throughout each relaxation, the geometry, potential energy, and interatomic forces were recorded at every step, yielding a complete optimization trajectory for each initial structure.

However, we faced significant challenges during this data generation process, primarily due to the optking optimizer’s reliance on an internal coordinate system. For larger and more flexible molecules, the back-transformation from the optimized internal coordinates to Cartesian coordinates frequently failed, leading to a substantial number of unsuccessful optimizations. Consequently, we were only able to generate complete and successful optimization trajectories for approximately one-third of the molecules from the test set. For this successfully optimized subset, a single conformation was relaxed for each molecule to construct our final ground truth dataset.