Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Efficient Nudged Elastic Band Method using Neural Network Bayesian Algorithm Execution

Pranav Kakhandiki

pkakhand@stanford.edu

Stanford University https://orcid.org/0000-0001-6617-8061

Sathya Chitturi

PostEra

Daniel Ratner

SLAC National Accelerator Laboratory

Sean Gasiorowski

SLAC National Accelerator Laboratory https://orcid.org/0000-0002-4067-2472

Article

Keywords: Nudged Elastic Band, Bayesian Experimental Design, Minimum Energy Pathways, Active Learning, Foundation Models, Machine Learning Potentials

Posted Date: March 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8389167/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Efficient Nudged Elastic Band Method using Neural Network Bayesian Algorithm Execution

Pranav Kakhandiki[1*], Sathya Chitturi[2], Daniel Ratner[2], Sean Gasiorowski[2*]

[1*]Department of Applied Physics, Stanford University, Stanford, 94305, CA, USA.

[2]SLAC National Accelerator Laboratory, Menlo Park, 94025, CA, USA.

*Corresponding author(s). E-mail(s): pkakhand@stanford.edu; sgaz@slac.stanford.edu;

Contributing authors: chitturi@alumni.stanford.edu; dratner@slac.stanford.edu;

Abstract

The discovery of a minimum energy pathway (MEP) between metastable states is crucial for scientific tasks including catalyst and biomolecular design. However, the standard nudged elastic band (NEB) algorithm requires hundreds to tens of thousands of compute-intensive simulations, making applications to complex systems prohibitively expensive. We introduce Neural Network Bayesian Algorithm Execution (NN-BAX), a framework that jointly learns the energy landscape and the MEP. NN-BAX sequentially fine-tunes a foundation model by actively selecting samples targeted at improving the MEP. Tested on Lennard-Jones and Embedded Atom Method systems, our approach achieves a one to two order of magnitude reduction in energy and force evaluations with negligible loss in MEP accuracy and demonstrates scalability to >100-dimensional systems. This work is therefore a promising step towards removing the computational barrier for MEP discovery in scientifically relevant systems, suggesting that weeks-long calculations may be achieved in hours or days with minimal loss in accuracy.

Keywords: Nudged Elastic Band, Bayesian Experimental Design, Minimum Energy Pathways, Active Learning, Foundation Models, Machine Learning Potentials

Introduction

Understanding how materials and chemical systems change from one metastable state to another is crucial in predicting and understanding their behavior. Many such processes in physics, chemistry, materials science, and biology depend on the transition-state pathways (the sequence of atomic configurations connecting two metastable states) and the energy barriers (the energetic cost required to move along that path) that connect different states [1]. These barriers dictate reaction rates, diffusion properties, and activation energies, which in turn affect macroscopic properties such as catalytic efficiency [2]. An example of a system governed by such energy barriers is conformational pathways in proteins [3], where structural changes are fundamental to biological function. Understanding these transitions lays the groundwork for controlling protein activity and developing synthetic biomolecules with tailored functions. Minimum energy pathways (MEPs) describe the most probable paths that atomic systems follow when transitioning between metastable states; critically, the lowest-energy pathway reveals the transition mechanism and corresponding energy barriers [4]. Nudged elastic band (NEB) is a popular method used for finding MEPs in atomic systems [5–7]. In its basic form [5], NEB begins with a set of "moving images" in atomic space between the initial and final state being studied, often using a linear interpolation as an initial guess. Force and energy calculations from physical simulations of the system are then used to nudge the band of images until a convergence criterion is met—commonly the maximum perpendicular force on any atom in any image, f_{max}, falling below a threshold. Each intermediate moving image is connected by an elastic spring force to ensure the band does not collapse to a single point. Since its development, various versions of NEB have been created to scale to more difficult systems or to reduce computational cost [8–12]. However, a major drawback of these "classical" NEB methods is that they require many evaluations of potential energy functions, namely one per image for each NEB iteration. Such simulations may involve computationally expensive density functional theory (DFT) or molecular dynamics (MD) calculations. For example, a recent NEB study of catalysts [13] with DFT took 52 GPU years for 19,000 NEB simulations. This time-consuming nature of NEB limits the possibilities of studying complex realistic systems, such as proteins and catalysts.

Numerous efforts have been made to reduce the computational cost of NEB with machine learning [2, 13–24]. In one paradigm, neural network surrogate models are trained on large, diverse datasets. These models then serve as direct calculators for downstream applications, such as NEB calculations on complex systems. An alternative approach involves sequentially updating a simulation surrogate for a system of interest using techniques from active learning or Bayesian experimental design. Bayesian experimental design techniques can be used to infer properties of black-box functions with limited function queries. Common approaches fall into the Bayesian optimization (BO) family [25], where the aim is to maximize or minimize the unknown function. A useful tool in Bayesian experimental design is Bayesian algorithm execution (BAX) [26], which generalizes the optimization objective to any computable function property (called the BAX "algorithm"). However, these two existing methodologies, namely generic foundation models and Bayesian optimization, have limitations. Foundation models are fast, but sometimes not accurate enough on

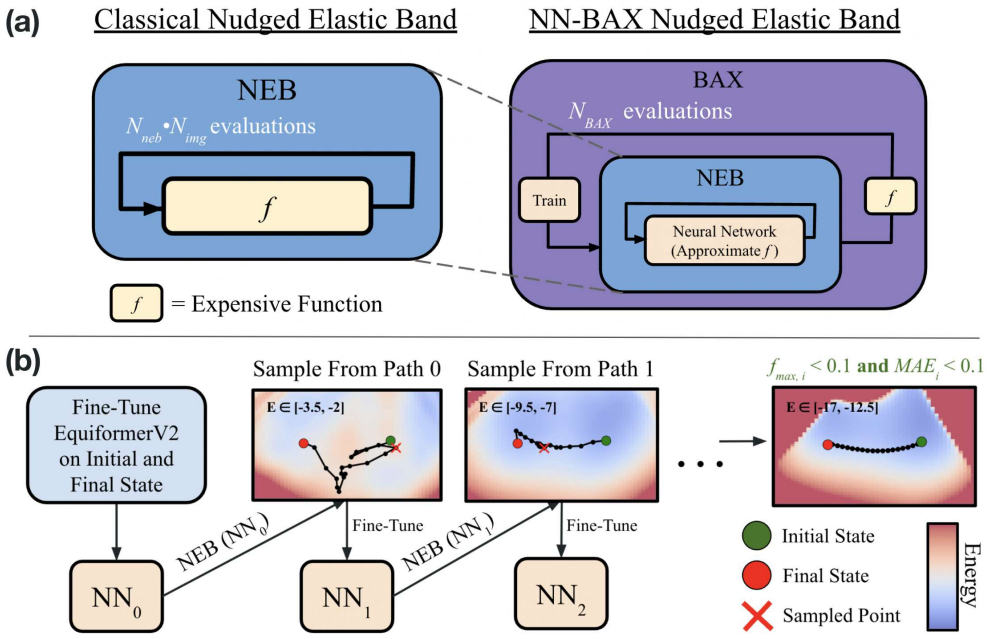


Fig. 1 NN-BAX Pipeline. a) Classical NEB versus NN-BAX NEB, where the expensive simulator f is no longer called in the NEB Loop. The number of function evaluations in classical NEB scales with the number of images times the number of NEB iterations, $N_{neb} \cdot N_{img}$, whereas NN-BAX scales with the number of BAX iterations, N_{BAX}, need to learn the function. b) NN-BAX Loop. The EquiformerV2 model is initialized and trained on the initial and final state. In each loop, NEB runs using the trained model as an approximation for f, a sample is acquired from the resulting path, and the sampled point is added to the training set for the subsequent iteration. Colormap energy ranges (eV) are provided.

new systems and thus may not converge to the correct transition pathway. Bayesian experimental design approaches provide constant updates, but lack sophisticated models needed for complex systems. Specifically, previous work has used Gaussian processes (GPs) as a surrogate model, limiting their expressivity and scalability in high-dimensional spaces [27]. The computational cost of GPs also scales cubically with dataset size [28], thus making GPs unfit for larger datasets. The approach in [14], which fits into the BAX framework and uses a GP, tests on a maximum of 7-atom systems.

In our work, we fuse these approaches, taking advantage of the development of large neural network surrogates for materials systems and applying them within the BAX acquisition framework, allowing for applicability to higher dimensional systems. Specifically, with BAX, we choose function queries that target the output of the NEB optimization on the unknown function defined by the physics simulator. We refer to this neural network BAX methodology as NN-BAX. There exist many foundation models for materials systems [29], and we use EquiformerV2 [30] as our surrogate model, pretrained on the OMat24 [31] dataset. The pre-trained EquiformerV2 model does not perform well for zero-shot NEB calculations on the systems we test. We

therefore treat it as a foundation model and tune it in a few-shot-learning manner using BAX. A visualization of the pipeline is shown in Figure 1.

Results

We evaluate NN-BAX across a hierarchy of atomic transition systems of increasing physical realism and complexity, as described by two well-known empirical potentials. First we study transitions in Lennard–Jones clusters with 7 and 38 atoms, which provide well-characterized, high-dimensional testbeds with known minimum energy pathways. We then extend our analysis to a many-body Embedded Atom Method (EAM) potential for copper surface diffusion, demonstrating applicability beyond pairwise interactions. Finally, we consider complex multi-step transitions in 38-atom Lennard-Jones clusters and introduce Foundation-BAX, which leverages information shared across related paths to further reduce computational cost. Together, these systems allow us to assess the accuracy, efficiency, and scalability of NN-BAX across increasingly challenging regimes.

For each NN-BAX iteration, we initialize the model to a pretrained foundation model. For the cases presented below, we consider the 153 million parameter EquiformerV2 trained on the OMat24 dataset. EquiformerV2 combines the Transformer architecture [32] with consideration of relevant symmetries for materials systems: atomic energies are modeled as invariant under global translations and rotations of the system, while atomic forces are modeled as equivariant, transforming consistently with rotations and translations of the atomic positions. EquiformerV2 has demonstrated strong performance in benchmarks [33] and previous non active learning (zero-shot) neural network NEB studies such as CatTsunami [13]. The inputs to the network are the positions of the atoms, x_i, and the outputs are the forces on each atom and the potential energy of the system.

Lennard-Jones Transitions

We demonstrate NN-BAX on atomic systems governed by the Lennard-Jones (LJ) potential:

$$V_{\mathrm{LJ}}(r) = 4\varepsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right], \tag{1}$$

because of its physical relevance and its well-studied transitions [34–42]. Unlike energy/force computations involving density functional theory, the LJ potential is inexpensive to call. However, it retains much of the complexity of similar chemical systems, making it ideal for testing our method. We analyze systems with 7 and 38 identical atoms, abbreviated as LJ_7 and LJ_{38}. The dimensionality of these systems, particularly the 114-dimensional LJ_{38}, makes Gaussian process surrogates infeasible. We specifically choose N = 7, 38 for our LJ_N systems because these systems exhibit many interesting local minima configurations, thus leading to them also being well-studied. LJ_{38} is especially interesting because the two lowest energy configurations are well separated in atomic configuration space, meaning the transition between the two is nontrivial [8].

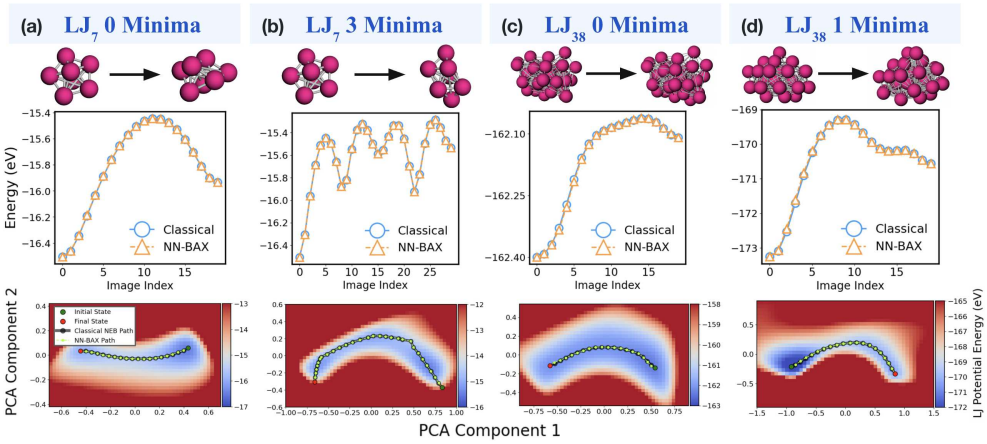


Fig. 2 Results from running NN-BAX on various paths with LJ_7 and LJ_{38} systems. The upper row compares the energy profile predicted by NN-BAX to classical NEB, along with atomic visuals for the transition. The bottom row shows dimension-reduced principal component analysis plots of the energy, with the NN-BAX and classical NEB paths overlaid.

We analyze four transitions of varying difficulty, from LJ_7 and LJ_{38}. For LJ_7 we look at paths with 0 and 3 intermediate minima, and for LJ_{38} we analyze paths with 0 and 1 intermediate minima. Paths with even more minima are incompatible with the NEB method used in this paper, and generally must be subdivided, as detailed in the Foundation-BAX section. These paths converged for classical NEB with f_{max} values of 0.05, 0.15, 0.05, and 0.3eV/Å respectively. We use the string method [9], which ignores spring forces and evenly redistributes images along the path every iteration. Figure 2 compares the converged NN-BAX paths with their classical NEB counterparts. The visuals on top display the atomic structures of the initial and final states. The upper plots show energy versus image index for both classical NEB and NN-BAX, with very close agreement between the two results. Of the four, the worst predicted energy barrier is 0.3% greater than the ground truth. The bottom row shows a two-dimensional principal component analysis (PCA) projection of the MEPs and potential energy landscapes, with projection computed from the ground truth path in atomic coordinate space. We observe that the NN-BAX paths qualitatively line up with the classical NEB paths, indicating that NN-BAX is finding the same structures as classical NEB.

Figure 3 displays the speed-up NN-BAX achieves with respect to classical NEB, in terms of function evaluations. For systems running expensive DFT simulations, we expect the NN-BAX procedural overhead to be negligible, making number of function calls an appropriate figure of merit. We achieve a 1-2 orders of magnitude reduction in function calls across all four paths. Finally, we observe that the LJ_7 path with 3 minima achieves a greater speedup than the other paths. We hypothesize this occurs because while classical NEB time is determined by the complexity of NEB convergence (i.e. number of NEB iterations), NN-BAX is limited by modeling complexity of the potential energy surface. In NN-BAX, the convergence complexity has negligible impact because the NEB loop only calls the surrogate, and in the classical case this

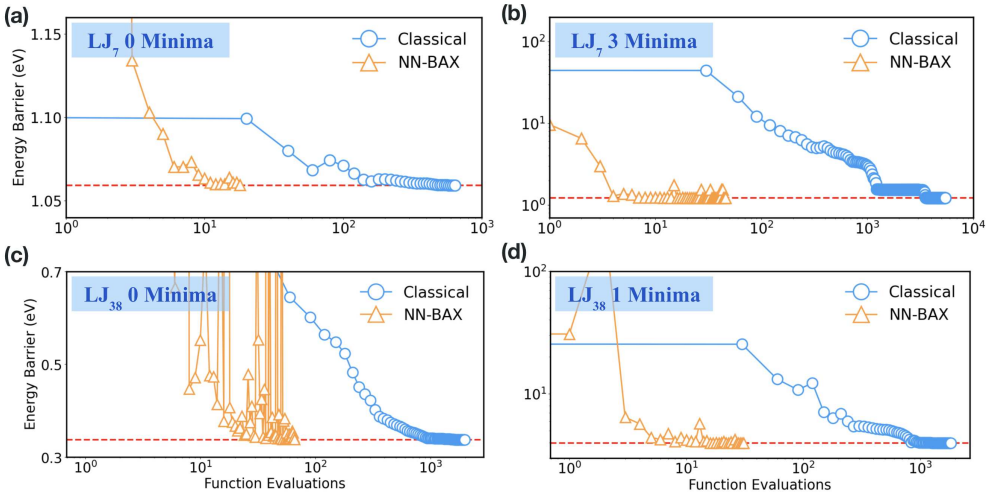


Fig. 3 Energy barrier v. function evaluations, for various LJ_7 and LJ_{38} systems. Speedup factors are 34, 116, 20, and 38 for a), b), c), and d) respectively. The dashed red line denotes the energy barrier achieved by classical NEB. Note that for the NN-BAX curve, each point corresponds to an NEB optimization. We show here how quickly our method matches the ground truth for the energy barrier, the primary physical quantity of interest.

path requires ∼3x as many NEB iterations compared to the others. Thus, because the number of function evaluations in NN-BAX no longer scales with the number of NEB iterations, we observe a greater speedup for this path. We emphasize that NN-BAX has the greatest impact in cases where NEB complexity, specifically the number of NEB iterations required for convergence, is large.

With the current NN-BAX implementation, we estimate the NN-BAX overhead to be equivalent to running classical NEB with simulations that each take approximately 25.4 seconds. For simulation times $\gg$25.4 seconds, as may be expected for DFT calculations on complex systems, the NN-BAX overhead becomes negligible. Details of this estimate, as well as suggestions of implementation optimizations to reduce the NN-BAX overhead, are given in the supplemental material.

Embedded Atom Method Potentials

Although Lennard-Jones structures and transitions host much complexity as a test system, the LJ potential is still pairwise, and thus does not capture some of the many-body nuance seen in DFT. Next, we test NN-BAX on an Embedded Atom Method (EAM) potential. EAM potentials are semi-empirical and designed especially for metals. They take into account both pairwise interactions and the collective electron density from neighboring atoms [43, 44]. This many-body formulation captures the physics of metals, for example surface diffusion [45–48]. Surface diffusion plays a crucial role in processes such as thin-film growth, nanostructure formation, and nanoscale stability of metals [49–52].

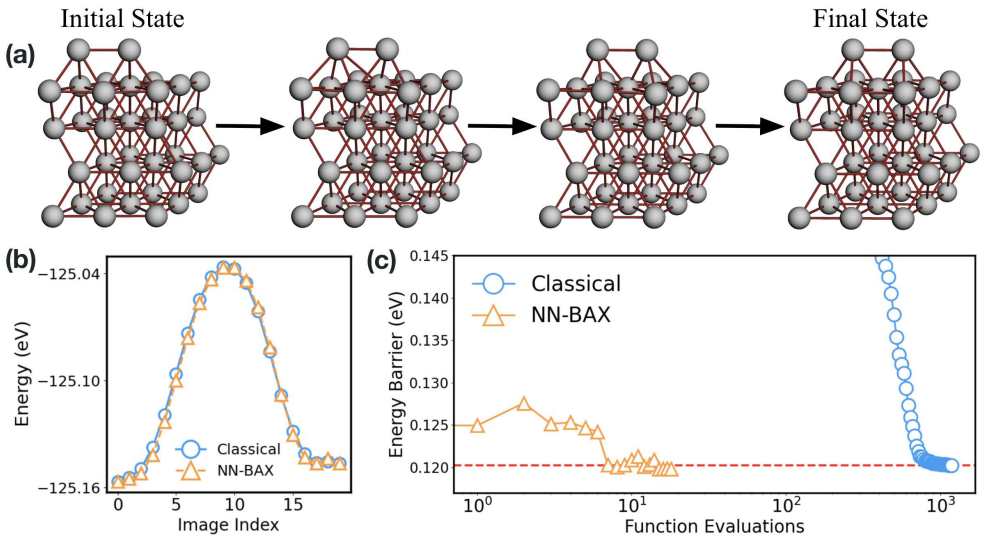


Fig. 4 Embedded Atom Method potential with NN-BAX. a) displays a visual of the copper dimer diffusion transition. b) shows the energy profile from classical NEB and from NN-BAX. c) shows the energy barrier of classical NEB and NN-BAX, as a function of evaluations. The red line denotes the energy barrier achieved by classical NEB.

We analyze a diffusion transition for copper involving a copper dimer moving to an adjacent site. Specifically, we analyze a 3x3x4 Cu(100) slab, with a pair of copper atoms atop. The bottom layer, consisting of nine atoms, is fixed to represent the bulk, and the rest of the atoms are free to move. Thus a total of 29 atoms are dynamic, giving us an 87 dimensional system. Note that while only the two atoms comprising the dimer differ in the initial and final states, the atoms beneath do move non-trivially during the transition. A visual for the transition is shown in Figure 4a. We use the potential described in [53], which is specific to copper. For this transition we use the traditional NEB method with spring forces with $k = 0.1\,\text{eV/Å}^2$, and observe that the path converges for $f_{max} < 0.01\text{eV/Å}$. Shown in Figure 4b is the classical NEB energy profile and the NN-BAX energy profile. The two profiles match well, indicating that NN-BAX has found the correct path. Finally, Figure 4c shows the energy barrier versus function evaluations, and we observe a significant speedup from NN-BAX. Note that the initial energy barrier for the EAM potential is relatively better than those in the LJ systems because this system more closely resembles those in the OMat24 dataset.

Multi-step Transitions with Foundation-BAX

For more complex transitions with many intermediate minima, the NEB algorithm often cannot converge to the correct path. Thus, it is common to use techniques to segment the transition into smaller paths between neighboring local minima. For Lennard-Jones transitions, the `OPTIM` program [54] is used to search for complex multi-step transitions in high dimensional systems such as LJ_{38}. In order to achieve an even

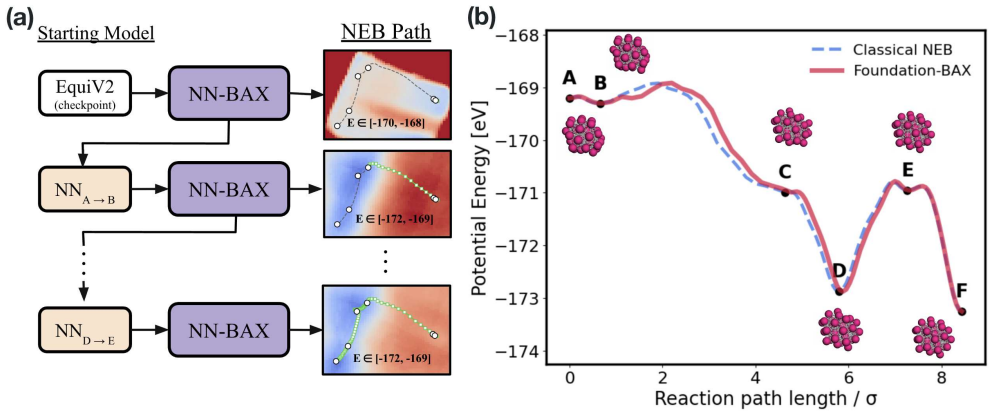


Fig. 5 Foundation-BAX pipeline and path. a) displays the Foundation-BAX pipeline, where the model from the previous path is used as the starting model for the next path b) shows the LJ_{38} path, which consists of 5 sub-transitions, connecting 6 local minima.

greater speedup with respect to NN-BAX for such multi-step transitions, we introduce Foundation-BAX; i.e. we use BAX to assemble and train a foundation model on a shared pool of relevant simulations. One can imagine this technique applied broadly, where researchers who run algorithms on similar systems can collectively pool data to achieve more accurate surrogate models. In our case of multi-step transitions, implementing Foundation-BAX involves using information from previous paths to speed up NN-BAX for subsequent paths. Figure 5a displays the Foundation-BAX pipeline, where we use the final fine-tuned model of the previous path as the initial model for the neighboring path.

To test the Foundation-BAX method, we choose a multi-step LJ_{38} transition that connects six local minima, thus containing four intermediate minima. This transition specifically is a subset of a LJ_{38} global minimum to second global minimum transition pathway, taken from example in [55]. Figure 5b displays the potential energy of the true transition pathway, along with the potential energy of the path predicted by Foundation-BAX. We note that the overall shape and energy barrier of each transition are quite similar. Note that there is a slight shift in the predicted energy profile with respect to the true transition, particularly in the transition from B to C. This is a result of Foundation-BAX finding an equally energetically favorable yet marginally different path.

Figure 6 displays the speedup, in terms of function evaluations, achieved by Foundation-BAX, compared to classical NEB and NN-BAX. Figure 6a displays the energy barrier v. function evaluation for Foundation-BAX, NN-BAX, and classical NEB, for the C to D transition. We observe that Foundation-BAX has a 4.5× speedup compared to NN-BAX and a 52× speedup compared to classical NEB. Note the significantly flatter nature of the Foundation-BAX curve, implying that some of the force landscape has already been learned, leading to a quicker convergence. Figure 6b compares the total number of function evaluations for all three methods, for all paths. We start with a standard NN-BAX run for the A to B transition. For subsequent

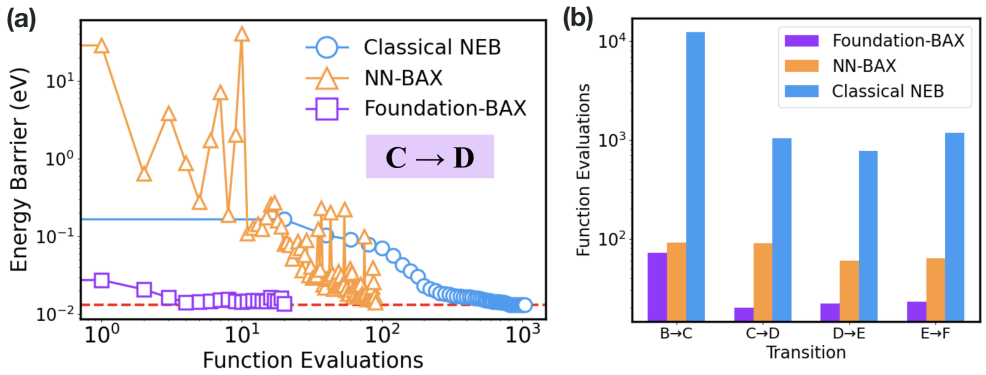


Fig. 6 Foundation-BAX speedup. a) Compares the number of function evaluations for Foundation-BAX, NN-BAX, and Classical NEB for the C to D transition. b) Shows the total number of function evaluations required for all three methods for all paths.

runs, we observe that Foundation-BAX requires fewer function evaluations to converge compared to NN-BAX. Furthermore, the speedup with respect to NN-BAX for the final three transitions are greater than that of the B to C transition, indicating that the model benefited by learning from more paths. Finally we note that Foundation-BAX works optimally with different convergence hyperparameters than NN-BAX, and elaborate more in supplemental materials.

Discussion

Accurately characterizing the transformations between metastable states is key to understanding and predicting the behavior of materials and chemical systems. The nudged elastic band algorithm is critical in finding transition states and energy barriers in such chemical systems. Systems modeled by Lennard-Jones or Embedded Atom Method potentials reproduce key features of complex chemical interactions, providing a controlled framework for studying high-dimensional atomic dynamics. By applying our methods to these high dimensional systems, we unify efforts in active learning and foundation models using BAX. Specifically, we apply NN-BAX to NEB on LJ and EAM potentials and observe a speedup of 1-2 orders of magnitude in terms of number of function evaluations. Furthermore, we observe that it is possible to gain a greater advantage in multi-step transitions with Foundation-BAX, where we fine-tune on previous fine-tuned models. We anticipate that Foundation-BAX can extend beyond multi-step transitions and be used more broadly to accelerate the speed of a suite of jobs governed by the same potential.

Previous use of BAX has been restricted to Gaussian process surrogates, which famously suffer from the curse of dimensionality [27]. The primary advantage of NN-BAX relies on the introduction of a neural network as a surrogate model in the BAX procedure, which allows for scalability to higher dimensional systems. We emphasize that while LJ and EAM do not require expensive *ab initio* calculations, our work with NN-BAX describes a broader system-agnostic procedure. Therefore, we expect

our work to extend to both new surrogate models and NEB applications, including high-dimensional systems that require expensive simulations like DFT. This technique may help us understand systems currently infeasible to study, including complex catalytic reactions and conformational pathways in proteins [3, 13]. Estimating 1000 energy/force evaluations per NEB optimization in [13], we estimate ∼14.4 minutes per DFT calculation for catalysts. This fits the ≫25.4 seconds we present above in order for the NN-BAX overhead to be negligible. With some improvements to our implementation (see supplementary material) we expect that we can reduce the NN-BAX overhead to ∼3.1 seconds, further amplifying applicability.

This work shows that coupling algorithm-aware active learning with expressive, symmetry-aware neural network surrogates can substantially reduce the cost of transition state discovery. Our approach provides a path towards expediting high-dimensional transition-state calculations for complex systems in physics, chemistry, and biology, including those which are currently computationally intractable.

Methods

NN-BAX Algorithm

There are a variety of forms of BAX acquisition. InfoBAX [26], an information-based acquisition function, aims to maximize the expected information gain (EIG) about an algorithm output. However, it is costly to compute, involving several model retrainings to estimate the associated entropies. Heuristic or approximate BAX algorithms have been proposed that are much more efficient, such as MeanBAX [56] and PS-BAX [57] which run the algorithm on the posterior mean or a posterior sample and then take the output point with the highest predicted uncertainty, avoiding the cost of computing the full EIG. In this work, we take a similar, heuristic approach. Our BAX acquisition involves running NEB on a deterministic model of the potential energy/forces. We then randomly sample from the resulting NEB path to choose the next point to acquire and update the model for the next loop. Our design also leaves open a clear path toward incorporating more rigorous acquisition strategies that involve uncertainty quantification, such as MeanBAX and InfoBAX, which may yield additional gains.

System Setup and Training

To create a baseline, we first find local minima using basin hopping [34]. We then run classical NEB between pairs of local minima, varying f_{max} as needed to achieve convergence. To check stability of the resulting MEPs, we randomly perturb the positions of the images from the converged path by $\sigma_x \in [0, 0.1]$ Å and re-run the NEB minimization. If after 100 iterations the path does not move and no lower energy path is discovered, the MEP is considered a "ground truth" target for our NN-BAX optimization. For both classical and NN-BAX NEB calculation, we use the Atomic Simulation Environment (ASE) library [58]. All NEB runs in this paper use the FIRE optimizer [59] with parameters $dt = 0.1$ and $dt_{max} = 1.0$.

One problem with running NEB on a surrogate is that for less accurate models in earlier iterations, NEB often does not converge below the given f_{max} criterion. Thus

we must set some finite number of NEB iterations, N_{neb} otherwise the process would never converge. However for NN-BAX, unlike classical NEB, the number of function evaluations does not scale with N_{neb}. Thus, we introduce a hyperparameter $N_{neb,max}$, where every BAX iteration runs NEB for $N_{neb,max}$ iterations, and the path with the lowest f_{max} is selected. We find that $N_{neb,max} = 200$ works well for the initial suite of LJ transitions. For the EAM transition we use $N_{neb,max} = 100$. Finally, for the Foundation-BAX pathways, we used $N_{neb,max} = 100$ for all segments except from B to C, where we use a larger $N_{neb,max} = 650$.

In this paper we introduce a new convergence hyperparameter, MAE_i, which corresponds to the mean absolute error (MAE) in the prediction of model $i-1$ on the acquired image in iteration i. Reduction in MAE_i indicates that the model is performing better on points not in the training set and, once NN-BAX begins predicting similar final paths, the sampled points grow increasingly similar, making for an easier modeling task. For NN-BAX to converge, $\text{MAE}_i < m$ and $f^i_{max,BAX} < t$ must be satisfied simultaneously, where m and t are the respective convergence thresholds. For MAE_i convergence, $m = 0.1$ proved to generally work. For force convergence, we set $f_{\text{max,BAX}} = n\, f_{\text{max,classical}}$, where $n > 1$ accounts for small amounts of noise in the modeled predictions with respect to true potential. In practice, we found $n = 2$ to work well for our paths. For further confidence in convergence we introduce a patience metric p, where the convergence criteria must be achieved p times for NN-BAX to converge. $p = 2$ was found to be generally robust.

In each BAX iteration, the network was trained for 50 epochs with a batch size of 2 using the AdamW optimizer [60] (weight decay 0.001, gradient clipping at 100) and an initial learning rate of 2×10^{-4}. A cosine learning rate schedule was employed, with a warmup phase lasting 1 epoch (warmup factor 0.2) and a minimum learning rate set to 1% of the initial value. Exponential moving average (EMA) updates with a decay of 0.999 were applied throughout training. We use the mean absolute error (MAE) loss, giving forces four times the weight of energy since our convergence criterion is force-based. All training hyperparameters were the same those fairchem used to train the EquiformerV2 checkpoint [61], except for batch size and epochs which were modified to accommodate smaller dataset sizes. All models are trained on a single NVIDIA Tesla A100 GPU at the SLAC Shared Science Data Facility.

Acknowledgements. We acknowledge Guanzhi Li, Yu Lin, Minkyung Han, Yu Zhang, Samuel Klein, and David Wales for helpful insights and manuscript feedback. SG and PK were supported by the Department of Energy, Laboratory Directed Research and Development program at SLAC National Accelerator Laboratory, under contract DE-AC02-76SF00515.

Code availability. We provide a user-friendly implementation of NN-BAX with the nudged elastic band algorithm at: https://github.com/pranavkakhandiki/nn-bax-neb. This repository also contains example notebooks to aid in analysis.

Competing interests. The authors declare no competing interests.

References

[1] Piskulich, Z.A., Mesele, O.O., Thompson, W.H.: Activation energies and beyond. The Journal of Physical Chemistry A **123**(33), 7185–7194 (2019)

[2] Schaaf, L.L., Fako, E., De, S., Schäfer, A., Csányi, G.: Accurate energy barriers for catalytic reaction pathways: an automatic training protocol for machine learning force fields. npj Computational Materials **9**(1), 180 (2023)

[3] Mathews, D.H., Case, D.A.: Nudged elastic band calculation of minimal energy paths for the conformational change of a gg non-canonical pair. Journal of molecular biology **357**(5), 1683–1693 (2006)

[4] Eyring, H.: The activated complex in chemical reactions. The Journal of chemical physics **3**(2), 107–115 (1935)

[5] Jónsson, H., Mills, G., Jacobsen, K.W.: Nudged elastic band method for finding minimum energy paths of transitions. In: Classical and Quantum Dynamics in Condensed Phase Simulations, pp. 385–404. World Scientific, ??? (1998)

[6] Henkelman, G., Jónsson, H.: Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. The Journal of chemical physics **113**(22), 9978–9985 (2000)

[7] Henkelman, G., Uberuaga, B.P., Jónsson, H.: A climbing image nudged elastic band method for finding saddle points and minimum energy paths. The Journal of chemical physics **113**(22), 9901–9904 (2000)

[8] Trygubenko, S.A., Wales, D.J.: A doubly nudged elastic band method for finding transition states. The Journal of chemical physics **120**(5), 2082–2094 (2004)

[9] Weinan, E., Ren, W., Vanden-Eijnden, E.: String method for the study of rare events. Physical Review B **66**(5), 052301 (2002)

[10] Maragakis, P., Andreev, S.A., Brumer, Y., Reichman, D.R., Kaxiras, E.: Adaptive nudged elastic band approach for transition state calculation. The Journal of chemical physics **117**(10), 4651–4658 (2002)

[11] Sheppard, D., Xiao, P., Chemelewski, W., Johnson, D.D., Henkelman, G.: A generalized solid-state nudged elastic band method. The Journal of chemical physics **136**(7) (2012)

[12] Kolsbjerg, E.L., Groves, M.N., Hammer, B.: An automated nudged elastic band method. The Journal of chemical physics **145**(9) (2016)

[13] Wander, B., Shuaibi, M., Kitchin, J.R., Ulissi, Z.W., Zitnick, C.L.: Cattsunami: Accelerating transition state energy calculations with pretrained graph neural networks. ACS Catalysis **15**(7), 5283–5294 (2025)

[14] Garrido Torres, J.A., Jennings, P.C., Hansen, M.H., Boes, J.R., Bligaard, T.: Low-scaling algorithm for nudged elastic band calculations using a surrogate machine learning model. Physical review letters **122**(15), 156001 (2019)

[15] Koistinen, O.-P., Ásgeirsson, V., Vehtari, A., Jónsson, H.: Nudged elastic band calculations accelerated with gaussian process regression based on inverse interatomic distances. Journal of chemical theory and computation **15**(12), 6738–6751 (2019)

[16] Fu, X., Wood, B.M., Barroso-Luque, L., Levine, D.S., Gao, M., Dzamba, M., Zitnick, C.L.: Learning smooth and expressive interatomic potentials for physical property prediction. arXiv preprint arXiv:2502.12147 (2025)

[17] Luo, J., Fan, T., Zhang, J., Qiu, P., Shi, X., Chen, L.: Automatic identification of slip pathways in ductile inorganic materials by combining the active learning strategy and neb method. npj Computational Materials **11**(1), 41 (2025)

[18] Deng, B., Choi, Y., Zhong, P., Riebesell, J., Anand, S., Li, Z., Jun, K., Persson, K.A., Ceder, G.: Systematic softening in universal machine learning interatomic potentials. npj Computational Materials **11**(1), 9 (2025)

[19] Li, B., Xiao, J., Gao, Y., Zhang, J.Z., Zhu, T.: Transition state searching accelerated by neural network potential. Journal of Chemical Information and Modeling **65**(5), 2297–2303 (2025)

[20] Perego, S., Bonati, L.: Data efficient machine learning potentials for modeling catalytic reactivity via active learning and enhanced sampling. npj Computational Materials **10**(1), 291 (2024)

[21] Meyer, R., Schmuck, K.S., Hauser, A.W.: Machine learning in computational chemistry: An evaluation of method performance for nudged elastic band calculations. Journal of Chemical Theory and Computation **15**(11), 6513–6523 (2019)

[22] Koistinen, O.-P., Dagbjartsdóttir, F.B., Ásgeirsson, V., Vehtari, A., Jónsson, H.: Nudged elastic band calculations accelerated with gaussian process regression. The Journal of chemical physics **147**(15) (2017)

[23] Schreiner, M., Bhowmik, A., Vegge, T., Jørgensen, P.B., Winther, O.: Neural-neb—neural networks can find reaction paths fast. Machine Learning: Science and Technology **3**(4), 045022 (2022)

[24] Teng, C., Wang, Y., Bao, J.L.: Physical prior mean function-driven gaussian processes search for minimum-energy reaction paths with a climbing-image nudged elastic band: a general method for gas-phase, interfacial, and bulk-phase reactions. Journal of Chemical Theory and Computation **20**(10), 4308–4324 (2024)

[25] Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. Journal of Global optimization **13**(4), 455–492 (1998)

[26] Neiswanger, W., Wang, K.A., Ermon, S.: Bayesian algorithm execution: Estimating computable properties of black-box functions using mutual information. In: International Conference on Machine Learning, pp. 8005–8015 (2021). PMLR

[27] Binois, M., Wycoff, N.: A survey on high-dimensional gaussian process modeling with application to bayesian optimization. ACM Transactions on Evolutionary Learning and Optimization **2**(2), 1–26 (2022)

[28] Williams, C.K.I., Rasmussen, C.E.: Gaussian Processes for Machine Learning. MIT Press, Cambridge, MA (2006)

[29] Riebesell, J., Goodall, R.E., Benner, P., Chiang, Y., Deng, B., Ceder, G., Asta, M., Lee, A.A., Jain, A., Persson, K.A.: A framework to evaluate machine learning crystal stability predictions. Nature Machine Intelligence **7**(6), 836–847 (2025)

[30] Liao, Y.-L., Wood, B., Das, A., Smidt, T.: Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations. arXiv preprint arXiv:2306.12059 (2023)

[31] Barroso-Luque, L., Shuaibi, M., Fu, X., Wood, B.M., Dzamba, M., Gao, M., Rizvi, A., Zitnick, C.L., Ulissi, Z.W.: Open materials 2024 (omat24) inorganic materials dataset and models. arXiv preprint arXiv:2410.12771 (2024)

[32] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)

[33] Riebesell, J., Goodall, R.E., Benner, P., Chiang, Y., Deng, B., Lee, A.A., Jain, A., Persson, K.A.: Matbench discovery–a framework to evaluate machine learning crystal stability predictions. arXiv preprint arXiv:2308.14920 (2023)

[34] Wales, D.J., Doye, J.P.: 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 (1997)

[35] Doye, J.P., Wales, D.J.: Saddle points and dynamics of lennard-jones clusters, solids, and supercooled liquids. The Journal of Chemical Physics **116**(9), 3777–3788 (2002)

[36] Schwerdtfeger, P., Wales, D.J.: 100 years of the lennard-jones potential. Journal of Chemical Theory and Computation **20**(9), 3379–3405 (2024)

[37] Wales, D.J.: Rearrangements of 55-atom lennard-jones and (c60) 55 clusters. The Journal of chemical physics **101**(5), 3750–3762 (1994)

[38] Doye, J.P., Wales, D.J., Miller, M.A.: Thermodynamics and the global optimization of lennard-jones clusters. The Journal of Chemical Physics **109**(19), 8143–8153 (1998)

[39] Doye, J.P., Miller, M.A., Wales, D.J.: The double-funnel energy landscape of the 38-atom lennard-jones cluster. The Journal of Chemical Physics **110**(14), 6896–6906 (1999)

[40] Sehgal, R.M., Maroudas, D., Ford, D.M.: Phase behavior of the 38-atom lennard-jones cluster. The Journal of Chemical Physics **140**(10) (2014)

[41] Wales, D.J., Miller, M.A., Walsh, T.R.: Archetypal energy landscapes. Nature **394**(6695), 758–760 (1998)

[42] Bohner, M.U., Meisner, J., Kästner, J.: A quadratically-converging nudged elastic band optimizer. Journal of Chemical Theory and Computation **9**(8), 3498–3504 (2013)

[43] Daw, M.S., Baskes, M.I.: Embedded-atom method: Derivation and application to impurities, surfaces, and other defects in metals. Physical review B **29**(12), 6443 (1984)

[44] Foiles, S., Baskes, M., Daw, M.S.: Embedded-atom-method functions for the fcc metals cu, ag, au, ni, pd, pt, and their alloys. Physical review B **33**(12), 7983 (1986)

[45] Marinica, M.-C., Barreteau, C., Spanjaard, D., Desjonquères, M.-C.: Diffusion rates of cu adatoms on cu (111) in the presence of an adisland nucleated at fcc or hcp sites. Physical Review B—Condensed Matter and Materials Physics **72**(11), 115402 (2005)

[46] Morgenstern, K., Braun, K.-F., Rieder, K.-H.: Direct imaging of cu dimer formation, motion, and interaction with cu atoms on ag (111). Physical review letters **93**(5), 056102 (2004)

[47] Sbiaai, K., Eddiai, A., Boughaleb, Y., Mazroui, M., Raty, J.-Y., Meddad, M., Kara, A.: Diffusion processes of trimers on missing row surfaces: Cu 3/ag (110) and ag 3/cu (110). Optical and Quantum Electronics **46**(1), 15–22 (2014)

[48] Mińkowski, M., Załuska-Kotur, M.A.: Diffusion of cu adatoms and dimers on cu (111) and ag (111) surfaces. Surface Science **642**, 22–32 (2015)

[49] Xia, X., Xie, S., Liu, M., Peng, H.-C., Lu, N., Wang, J., Kim, M.J., Xia, Y.: On the role of surface diffusion in determining the shape or morphology of noble-metal nanocrystals. Proceedings of the National Academy of Sciences **110**(17), 6669–6673 (2013)

[50] Bonzel, H.P.: In: Laskar, A.L., Bocquet, J.L., Brebec, G., Monty, C. (eds.) Surface Diffusion on Metals, pp. 307–308. Springer, Dordrecht (1990)

[51] Wang, W., Wang, M., Ambrosi, E., Bricalli, A., Laudato, M., Sun, Z., Chen, X., Ielmini, D.: Surface diffusion-limited lifetime of silver and copper nanofilaments in resistive switching devices. Nature communications **10**(1), 81 (2019)

[52] Forgerini, F.L., Marchiori, R.: A brief review of mathematical models of thin film growth and surfaces: a possible route to avoid defects in stents. Biomatter **4**(1), 28871 (2014)

[53] Mishin, Y., Mehl, M.J., Papaconstantopoulos, D.A., Voter, A.F., Kress, J.D.: Structural stability and lattice defects in copper: Ab initio, tight-binding, and embedded-atom calculations. Physical Review B **63**(22), 224106 (2001)

[54] Wales, D.: Optim: A program for optimizing geometries and calculating reaction pathways. URL http://www-wales. ch. cam. ac. uk/software. html (2003)

[55] Group, W.: Examples from the Wales Group Repository. University of Cambridge. Accessed: 2025-11-03. https://github.com/wales-group/examples

[56] Chitturi, S.R., Ramdas, A., Wu, Y., Rohr, B., Ermon, S., Dionne, J., Jornada, F.H.d., Dunne, M., Tassone, C., Neiswanger, W., *et al.*: Targeted materials discovery using bayesian algorithm execution. npj Computational Materials **10**(1), 156 (2024)

[57] Cheng, C.X., Astudillo, R., Desautels, T.A., Yue, Y.: Practical bayesian algorithm execution via posterior sampling. Advances in Neural Information Processing Systems **37**, 135186–135207 (2024)

[58] Larsen, A.H., Mortensen, J.J., Blomqvist, J., Castelli, I.E., Christensen, R., Dułak, M., Friis, J., Groves, M.N., Hammer, B., Hargus, C., *et al.*: The atomic simulation environment—a python library for working with atoms. Journal of Physics: Condensed Matter **29**(27), 273002 (2017)

[59] Bitzek, E., Koskinen, P., Gähler, F., Moseler, M., Gumbsch, P.: Structural relaxation made simple. Physical review letters **97**(17), 170201 (2006)

[60] Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)

[61] team, F.: FAIRChem: FAIR Chemistry — library of machine learning methods for chemistry. https://github.com/facebookresearch/fairchem. commit hash as of accessed version, accessed 2025-12-09 (2025)

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- NNBAXSupp.pdf