
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Koistinen, Olli-Pekka; Dagbjartsdóttir, Freyja B.; Ásgeirsson, Vilhjálmur; Vehtari, Aki; Jonsson, Hannes

Nudged elastic band calculations accelerated with Gaussian process regression

Published in:
Journal of Chemical Physics

DOI:
[10.1063/1.4986787](https://doi.org/10.1063/1.4986787)

Published: 21/10/2017

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Koistinen, O-P., Dagbjartsdóttir, F. B., Ásgeirsson, V., Vehtari, A., & Jonsson, H. (2017). Nudged elastic band calculations accelerated with Gaussian process regression. *Journal of Chemical Physics*, 147(15), 1-14. [152720]. <https://doi.org/10.1063/1.4986787>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Nudged elastic band calculations accelerated with Gaussian process regression

Olli-Pekka Koistinen, Freyja B. Dagbjartsdóttir, Vilhjálmur Ásgeirsson, Aki Vehtari, and Hannes Jónsson

Citation: *The Journal of Chemical Physics* **147**, 152720 (2017); doi: 10.1063/1.4986787

View online: <https://doi.org/10.1063/1.4986787>

View Table of Contents: <http://aip.scitation.org/toc/jcp/147/15>

Published by the [American Institute of Physics](#)

Articles you may be interested in

[Rock climbing: A local-global algorithm to compute minimum energy and minimum free energy pathways](#)

The Journal of Chemical Physics **147**, 152718 (2017); 10.1063/1.4986298

[A climbing image nudged elastic band method for finding saddle points and minimum energy paths](#)

The Journal of Chemical Physics **113**, 9901 (2000); 10.1063/1.1329672

[An automated nudged elastic band method](#)

The Journal of Chemical Physics **145**, 094107 (2016); 10.1063/1.4961868

[Ab initio sampling of transition paths by conditioned Langevin dynamics](#)

The Journal of Chemical Physics **147**, 152703 (2017); 10.1063/1.4985651

[Predicting reaction coordinates in energy landscapes with diffusion anisotropy](#)

The Journal of Chemical Physics **147**, 152701 (2017); 10.1063/1.4983727

[Acceleration of saddle-point searches with machine learning](#)

The Journal of Chemical Physics **145**, 074106 (2016); 10.1063/1.4960708

PHYSICS TODAY

WHITEPAPERS

ADVANCED LIGHT CURE ADHESIVES

Take a closer look at what these environmentally friendly adhesive systems can do

READ NOW

PRESENTED BY
 **MASTERBOND**
ADHESIVES | SEALANTS | COATINGS

Nudged elastic band calculations accelerated with Gaussian process regression

Olli-Pekka Koistinen,^{1,2,3} Freyja B. Dagbjartsdóttir,² Vilhjálmur Ásgeirsson,² Aki Vehtari,¹ and Hannes Jónsson^{2,3,a)}

¹*Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University, Espoo, Finland*

²*Science Institute and Faculty of Physical Sciences, University of Iceland, Reykjavík, Iceland*

³*Department of Applied Physics, Aalto University, Espoo, Finland*

(Received 6 June 2017; accepted 29 August 2017; published online 21 September 2017)

Minimum energy paths for transitions such as atomic and/or spin rearrangements in thermalized systems are the transition paths of largest statistical weight. Such paths are frequently calculated using the nudged elastic band method, where an initial path is iteratively shifted to the nearest minimum energy path. The computational effort can be large, especially when *ab initio* or electron density functional calculations are used to evaluate the energy and atomic forces. Here, we show how the number of such evaluations can be reduced by an order of magnitude using a Gaussian process regression approach where an approximate energy surface is generated and refined in each iteration. When the goal is to evaluate the transition rate within harmonic transition state theory, the evaluation of the Hessian matrix at the initial and final state minima can be carried out beforehand and used as input in the minimum energy path calculation, thereby improving stability and reducing the number of iterations needed for convergence. A Gaussian process model also provides an uncertainty estimate for the approximate energy surface, and this can be used to focus the calculations on the lesser-known part of the path, thereby reducing the number of needed energy and force evaluations to a half in the present calculations. The methodology is illustrated using the two-dimensional Müller-Brown potential surface and performance assessed on an established benchmark involving 13 rearrangement transitions of a heptamer island on a solid surface. *Published by AIP Publishing.* <https://doi.org/10.1063/1.4986787>

I. INTRODUCTION

Theoretical studies of the transition mechanism and estimation of the rate of thermally activated events involving displacements of atoms or rotations of magnetic moments often involve finding a minimum energy path (MEP) connecting initial and final state minima on the energy surface characterizing the system. An MEP is a natural choice for a reaction coordinate since it represents a path of maximal statistical weight in a system in thermal equilibrium with a heat bath. Transition state theory^{1–3} calculations can be carried out using this reaction coordinate to parametrize, for example, a hyperplanar representation of the transition state.^{4,5} Even though such a reaction coordinate represents only one particular mechanism for the transition, it is possible to discover a new mechanism corresponding to a lower free energy barrier when full variational optimization of both the location and orientation of the hyperplanar transition state is carried out.^{6,7} In such a case, an MEP is just a convenient tool for the implementation of a full free energy calculation.

Most often, transition rates are, however, estimated from the harmonic approximation to transition state theory, where

the maximum rise in the energy along an MEP gives the activation energy of the transition and the pre-exponential factor in the Arrhenius expression for the rate can be obtained from the Hessian matrix evaluated at the initial state minimum and the energy maximum—a first-order saddle point on the energy surface.⁸ While it is possible to use various methods to converge directly on a saddle point starting from some initial guess, knowledge of the whole MEP is useful because it is important to make sure that the highest first-order saddle point for the full transition has been found. Furthermore, calculations of MEPs often reveal unknown intermediate minima and unexpected transition mechanisms⁹ and therefore play an important role in the studies of the mechanism and rate of thermally activated transitions. Most often, such calculations are carried out for transitions involving rearrangements of atoms, but similar considerations apply to thermally activated transitions where magnetic moments rotate from one magnetic state to another.^{10–13}

The nudged elastic band (NEB) method is a commonly used iterative approach to find MEPs.^{5,9,14} For magnetic transitions, a variant of the method called geodesic NEB has been developed.¹⁵ In the NEB method, the path between two local minima on the energy surface is represented by a discrete set of replicas of the system, referred to as “images,” each of them consisting of values for all degrees of freedom. Starting from

^{a)}Electronic mail: hj@hi.is

some initial path, the locations of the images on the energy surface are iteratively optimized so as to obtain a discrete representation of an MEP.

Each NEB calculation typically involves on the order of a hundred evaluations of the energy and force (the negative gradient of the energy) for each one of the images, and the path is typically represented by five to ten images. The evaluations were initially performed mostly using analytical potential energy functions, but nowadays electronic structure calculations are also used extensively in NEB applications. Since a typical electronic structure calculation takes on the order of tens of CPU minutes or more, the NEB calculations can become computationally demanding. In addition, the calculation may need to be repeated if there are several possible final states for the transition. Thus, it would be valuable to find ways to accelerate NEB calculations. To get the most out of the computationally intensive electronic structure calculations, the information obtained from them should be exploited better to decrease the number of NEB iterations instead of forgetting it after one iteration.

The use of machine learning to accelerate MEP and saddle point calculations has been introduced by Peterson,¹⁶ who applied neural networks to construct an approximate energy surface for which NEB calculations were carried out. After relaxation of the path on the approximate energy surface, the true energy and force were evaluated at the images of the relaxed path to see whether or not the path had converged on an MEP on the true energy surface. If true convergence had not been reached, the new energy and force values calculated at the location of the images were added to the training data set and the model was updated. This procedure was repeated iteratively until the approximate energy surface was accurate enough for converging on the true MEP.

Proof-of-principle results have also been presented where Gaussian process regression (GPR)^{17–20} is applied to accelerate NEB calculations.²¹ Since the calculations are largely based on the gradient vector of the energy surface, straightforward inclusion of derivative observations and prediction of derivatives can be seen as advantages of GPR for this application. It is also easy to encode prior assumptions about the smoothness properties of the energy surface into the covariance function of the Gaussian process (GP) model or learn about these properties from the data. Analytical expressions for the posterior predictions conditional on the hyperparameters of the GP model allow both fast predictions and reliable estimation of uncertainties. The predictive performance of GPR has been shown to be competitive with other machine learning methods especially when the number of observations is small.²²

The GPR approach to MEP calculations is extended here by presenting two algorithms to accelerate climbing image nudged elastic band (CI-NEB) calculations, where one of the images is made to converge to a small tolerance on the highest energy maximum along the MEP.²³ The basic GPR approach is described as the all-images-evaluated (AIE) algorithm, where the energy and force are evaluated at all intermediate images of the CI-NEB before the approximation to the energy surface is updated. In a more advanced algorithm, the energy and force are evaluated at only one image before a new

approximate energy surface is constructed. We refer to the latter as the one-image-evaluated (OIE) algorithm. As a probabilistic model, a GP expresses the energy predictions as probability distributions, which means that the uncertainty of the prediction can also be estimated, e.g., as the variance of the posterior distribution. This uncertainty estimate is used by the OIE algorithm to select the image to be evaluated in such a way as to give maximal improvement of the model. By directing the evaluations to locations where they are most needed, the OIE algorithm skips some of the energy and force evaluations and thus decreases the overall computation time compared to the AIE algorithm. This approach has similarities with Bayesian optimization,²⁴ where the uncertainties of a GP model are used to define an acquisition function that is used to select the locations of new evaluations in a global optimization task.

Another extension of the GPR approach presented here applies when the overall goal is to estimate the forward and backward transition rates using harmonic transition state theory. Then, the Hessian matrix needs to be evaluated at the initial and final state minima, as well as at the highest first-order saddle point along the MEP. The evaluation of the Hessian at the endpoint minima can be carried out before the MEP calculation to provide additional input information into the GP model about the shape of the energy surface in the vicinity of the two ends of the MEP.

The article is organized as follows: In Sec. II, a brief introduction to the NEB method is given, followed by presentation of necessary GP theory for the GPR approach in Sec. III. In Sec. IV, the two implementations, the AIE algorithm and the OIE algorithm, are described and illustrated on the two-dimensional Müller-Brown energy surface.²⁵ In Sec. V, the heptamer island benchmark is described and performance statistics are given as a function of the number of degrees of freedom. The article concludes with a discussion in Sec. VI.

II. NUGGED ELASTIC BAND METHOD

The objective of the nudged elastic band (NEB) method is to find a minimum energy path (MEP) connecting two given local minima on an energy surface. An MEP is defined as a path for which the gradient of the energy has zero component perpendicular to the path tangent. In the NEB method, the path is represented in a discretized way as a set of images, which are sets of values of all degrees of freedom in the system (atom coordinates and angles specifying orientation of magnetic vectors, and possibly also simulation box size and shape). The MEP is found iteratively, starting from some initial path between the two minima. Most often, a straight line interpolation between the minima has been used to generate the initial path,¹⁴ but a better approach is to start with a path that interpolates as closely as possible the distances between neighboring atoms, the so-called image dependent pair potential (IDPP) method.²⁶

The key feature of the NEB algorithm is the “nudging,” a projection which is used to separate the force components perpendicular and parallel to the path from each other. If each image is just moved along the force vector (negative gradient

of the energy), they would end up sliding down to the nearest minima. The main idea in the NEB method is to take into account only the force component perpendicular to the path and at the same time control the distribution of the images along the path. The projection of the force requires an estimate of the local tangent to the path at the location of each image. A well-behaved estimate is obtained by defining the tangent based on the vector to the neighboring image of higher energy or, if both of the neighbors are either higher or lower in energy than the current image, a weighted average of the vectors to the two neighboring images.²⁷

To control the distribution of the images along the path, a spring force acting in the direction of the path tangent is typically introduced. The most common choice is to strive for an even distribution, but one can also choose to have, for example, a higher density of images where the energy is larger.²³ The spring force also prevents the path from becoming arbitrarily long in regions of little or no force. This is important, for example, in calculations of adsorption and desorption of molecules at surfaces.

In each iteration, the images are moved along the resultant vector of the spring force and the component of the true force perpendicular to the path, which is here referred to as the NEB force F_{NEB} . The true force is the negative gradient of the energy, and in most cases, an evaluation of the energy also delivers the gradient vector at little or no additional expense. It is, however, typically too expensive to evaluate the second derivatives of the energy, and the iterative algorithms are therefore based solely on the gradient and energy at each image (in addition to the spring force which depends on the distribution of the images). A simple and stable method that has been used extensively in NEB calculations will be used here to control the step size of the movements. It is based on a velocity Verlet dynamics algorithm where only the component of the velocity in the direction of the NEB force is included as long as the inner product with the NEB force is positive.¹⁴ A somewhat higher efficiency can be obtained by using quadratically convergent algorithms such as conjugate gradient or quasi-Newton,²⁸ but those can be less stable especially in the beginning of an NEB calculation.

The most important part of an MEP is the vicinity of the highest energy saddle point, especially in harmonic transition state theory calculations where the highest energy saddle point directly gives an estimate of the activation energy of the transition. It is, therefore, advantageous to let the highest energy image move to the maximum energy along the path. This variant of the NEB method is referred to as the climbing image nudged elastic band (CI-NEB) method.²³ Whereas the component of the true force acting in the direction of the tangent is normally removed from the NEB force, for the climbing image, it is instead flipped around to point towards the direction of higher energy along the path. In the CI-NEB method, the spring force is not applied to the climbing image and the rest of the images are distributed evenly on each side of the climbing image. To keep the intervals reasonably similar on both sides of the climbing image, the regular NEB method is typically conducted first (to some preliminary tolerance) so that the image selected as the climbing image is not too far from the true saddle point. The rest of the MEP is mainly

needed to ensure that the highest saddle point has been identified and to provide an estimate of the tangent to the path in order to carry out the nudging projections of the forces. It is more important to make the climbing image converge rigorously than the other images. It is, therefore, practical to apply a tighter tolerance for the magnitude of the NEB force acting on the climbing image than to the other images in CI-NEB calculations.

In the heptamer island benchmark presented here, the path was represented by seven images, $N_{\text{im}} = 7$, and the initial path was generated using the IDPP method. All spring constants were chosen to be 1 eV/\AA to give an even distribution of the images along the path on each side of the climbing image. The focus here is on calculations where the energy and force are obtained using some *ab initio* or density functional theory calculations. The computational effort in all other parts of the calculation is then insignificant in comparison, and thus the overall computational effort is well characterized by simply the number of times the energy and force need to be evaluated in order to converge on the MEP. Below, we describe various strategies to accelerate CI-NEB calculations with Gaussian process regression.

III. GAUSSIAN PROCESS REGRESSION

A Gaussian process (GP) is a flexible probabilistic model for functions in a continuous domain.^{17–20} It is defined by a mean function $m(\mathbf{x})$ and a covariance function $k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$ so that the joint probability distribution of the function values $\mathbf{f} = [f(\mathbf{x}^{(1)}), f(\mathbf{x}^{(2)}), \dots, f(\mathbf{x}^{(N)})]^T$ at any finite set of input points $\mathbf{X} = [\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}]^T \in \mathbb{R}^{N \times D}$ is a multivariate Gaussian distribution $p(\mathbf{f}) = \mathcal{N}(\mathbf{m}, K(\mathbf{X}, \mathbf{X}))$, where $\mathbf{m} = [m(\mathbf{x}^{(1)}), m(\mathbf{x}^{(2)}), \dots, m(\mathbf{x}^{(N)})]^T$ and the notation $K(\mathbf{X}, \mathbf{X}')$ represents a covariance matrix with entries $K_{ij} = k(\mathbf{x}_i, \mathbf{x}'_j)$. Thus, a GP can be seen as an infinite-dimensional generalization of the multivariate Gaussian distribution, serving as a prior probability distribution for the unknown function f . After evaluating the function at some training data points, the probability model is updated and a posterior probability distribution can be calculated for the function value at any point.

The most important part of the GP model is the covariance function, which defines how the function values at any two input points depend on each other, usually based on the distance between the points. Through selection of the covariance function, different prior assumptions about the properties of the function can be encoded into the model. To favor smooth functions, the infinitely differentiable squared exponential covariance function

$$k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \sigma_c^2 + \sigma_m^2 \exp\left(-\frac{1}{2} \sum_{d=1}^D \frac{(x_d^{(i)} - x_d^{(j)})^2}{l_d^2}\right)$$

is used here. The hyperparameters $\mathbf{l} = \{l_1, \dots, l_D\}$ are length scales that define the range of the covariance in each dimension, and σ_m^2 is a hyperparameter that controls the magnitude of the covariation. The mean function is here set to zero, but the additional constant term σ_c^2 in the covariance function has a similar effect as integration over an unknown constant

intercept term having a Gaussian prior distribution with variance σ_c^2 .

Consider a regression problem $y = f(\mathbf{x}) + \epsilon$, where ϵ is a Gaussian noise term with variance σ^2 , and a training data set $\{\mathbf{X}, \mathbf{y}\}$, where $\mathbf{y} = [y^{(1)}, y^{(2)}, \dots, y^{(N)}]^\top$ includes N noisy output observations from the N input points $\mathbf{X} = [\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}]^\top \in \mathbb{R}^{N \times D}$. The posterior predictive distribution for the function value $f(\mathbf{x}^*)$ at a new point \mathbf{x}^* , conditional on the GP model hyperparameters $\theta = \{\sigma_m^2, \mathbf{I}\}$, is a Gaussian distribution with mean

$$E[f(\mathbf{x}^*) | \mathbf{y}, \mathbf{X}, \theta] = K(\mathbf{x}^*, \mathbf{X}) (K(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I})^{-1} \mathbf{y}$$

and variance

$$\text{Var}[f(\mathbf{x}^*) | \mathbf{y}, \mathbf{X}, \theta]$$

$$= k(\mathbf{x}^*, \mathbf{x}^*) - K(\mathbf{x}^*, \mathbf{X}) (K(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I})^{-1} K(\mathbf{X}, \mathbf{x}^*),$$

where \mathbf{I} is the identity matrix. Here the hyperparameter values $\theta = \{\sigma_m^2, \mathbf{I}\}$ are optimized by defining a prior probability distribution $p(\theta)$ and maximizing the marginal posterior probability density $p(\theta | \mathbf{y}, \mathbf{X}) \propto p(\theta) p(\mathbf{y} | \mathbf{X}, \theta)$, where $p(\mathbf{y} | \mathbf{X}, \theta) = \int_{\mathbf{f}} p(\mathbf{y} | \mathbf{f}) p(\mathbf{f} | \mathbf{X}, \theta) d\mathbf{f}$ is the marginal likelihood of θ in the light of the observed data set $\{\mathbf{X}, \mathbf{y}\}$.

Since differentiation is a linear operation, the derivative of a GP is a GP as well.^{29–33} This makes it straightforward to use derivative information and predict derivatives of the function f . Derivative observations can be included in the model by extending the observation vector \mathbf{y} to include partial derivative observations and by extending the covariance matrix $K(\mathbf{X}, \mathbf{X})$ correspondingly to include covariances between the function values and partial derivatives and covariances between the partial derivatives themselves. In the case of the squared

exponential covariance function, these entries are obtained by

$$\begin{aligned} \text{Cov} \left[\frac{\partial f(\mathbf{x}^{(i)})}{\partial x_d^{(i)}}, f(\mathbf{x}^{(j)}) \right] &= \frac{\partial}{\partial x_d^{(i)}} \text{Cov} [f(\mathbf{x}^{(i)}), f(\mathbf{x}^{(j)})] = \frac{\partial k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})}{\partial x_d^{(i)}} \\ &= -\frac{\sigma_m^2 (x_d^{(i)} - x_d^{(j)})}{l_d^2} \exp \left(-\frac{1}{2} \sum_{g=1}^D \frac{(x_g^{(i)} - x_g^{(j)})^2}{l_g^2} \right) \end{aligned}$$

and

$$\begin{aligned} \text{Cov} \left[\frac{\partial f(\mathbf{x}^{(i)})}{\partial x_{d_1}^{(i)}}, \frac{\partial f(\mathbf{x}^{(j)})}{\partial x_{d_2}^{(j)}} \right] &= \frac{\partial^2}{\partial x_{d_1}^{(i)} \partial x_{d_2}^{(j)}} \text{Cov} [f(\mathbf{x}^{(i)}), f(\mathbf{x}^{(j)})] \\ &= \frac{\partial^2 k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})}{\partial x_{d_1}^{(i)} \partial x_{d_2}^{(j)}} \\ &= \frac{\sigma_m^2}{l_{d_1}^2} \left(\delta_{d_1 d_2} - \frac{(x_{d_1}^{(i)} - x_{d_1}^{(j)})(x_{d_2}^{(i)} - x_{d_2}^{(j)})}{l_{d_2}^2} \right) \\ &\quad \times \exp \left(-\frac{1}{2} \sum_{g=1}^D \frac{(x_g^{(i)} - x_g^{(j)})^2}{l_g^2} \right), \end{aligned}$$

where $\delta_{d_1 d_2} = 1$ if $d_1 = d_2$ and $\delta_{d_1 d_2} = 0$ if $d_1 \neq d_2$. These same expressions are needed when predicting the derivatives. The posterior predictive distribution of the partial derivative of function f with respect to dimension d at a new point \mathbf{x}^* is a Gaussian distribution with mean

$$E \left[\frac{\partial f(\mathbf{x}^*)}{\partial x_d^*} \middle| \mathbf{y}, \mathbf{X}, \theta \right] = \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_d^*} (K(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I})^{-1} \mathbf{y}$$

and variance

$$\text{Var} \left[\frac{\partial f(\mathbf{x}^*)}{\partial x_d^*} \middle| \mathbf{y}, \mathbf{X}, \theta \right] = \frac{\partial^2 k(\mathbf{x}^*, \mathbf{x}^*)}{\partial x_d^* \partial x_d^*} - \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_d^*} (K(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I})^{-1} \frac{\partial K(\mathbf{X}, \mathbf{x}^*)}{\partial x_d^*}.$$

In the present application, the vector \mathbf{x} includes coordinates of the atoms and the function f is the energy of the system. The extended observation vector

$$\mathbf{y}_{\text{ext}} = \left[y^{(1)} \dots y^{(N)}, \frac{\partial f(\mathbf{x}^{(1)})}{\partial x_1^{(1)}} \dots \frac{\partial f(\mathbf{x}^{(N)})}{\partial x_1^{(N)}}, \frac{\partial f(\mathbf{x}^{(1)})}{\partial x_2^{(1)}} \dots \frac{\partial f(\mathbf{x}^{(N)})}{\partial x_2^{(N)}}, \dots, \frac{\partial f(\mathbf{x}^{(1)})}{\partial x_D^{(1)}} \dots \frac{\partial f(\mathbf{x}^{(N)})}{\partial x_D^{(N)}} \right]^\top$$

includes the accurate values of the energy and the partial derivatives of the energy with respect to the coordinates of the atoms (i.e., components of the negative force vector) at the training data points $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}$. The GP model is used to predict the energy $f(\mathbf{x}^*)$ and its gradient vector $[\frac{\partial f(\mathbf{x}^*)}{\partial x_1^*}, \frac{\partial f(\mathbf{x}^*)}{\partial x_2^*}, \dots, \frac{\partial f(\mathbf{x}^*)}{\partial x_D^*}]^\top$ at a new point \mathbf{x}^* , which in this case represents an image on the discrete path representation between the initial and final state minima. Since the training data also include derivative observations, the mean and variance of the posterior predictive distribution of the energy are given as

$$E[f(\mathbf{x}^*) | \mathbf{y}_{\text{ext}}, \mathbf{X}, \theta] = \mathbf{K}_{\text{ext}}^* (\mathbf{K}_{\text{ext}} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_{\text{ext}} \quad (1)$$

and

$$\text{Var}[f(\mathbf{x}^*) | \mathbf{y}_{\text{ext}}, \mathbf{X}, \theta] = k(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{K}_{\text{ext}}^* (\mathbf{K}_{\text{ext}} + \sigma^2 \mathbf{I})^{-1} \mathbf{K}_{\text{ext}}^{* \top}, \quad (2)$$

where

$$\mathbf{K}_{\text{ext}}^* = \begin{bmatrix} K(\mathbf{x}^*, \mathbf{X}) & \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_1} & \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_2} & \dots & \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_D} \end{bmatrix}$$

and

$$\mathbf{K}_{\text{ext}} = \begin{bmatrix} K(\mathbf{X}, \mathbf{X}) & \frac{\partial K(\mathbf{X}, \mathbf{X}')}{\partial x'_1} & \frac{\partial K(\mathbf{X}, \mathbf{X}')}{\partial x'_2} & \dots & \frac{\partial K(\mathbf{X}, \mathbf{X}')}{\partial x'_D} \\ \frac{\partial K(\mathbf{X}, \mathbf{X}')}{\partial x_1} & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_1 \partial x'_1} & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_1 \partial x'_2} & \dots & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_1 \partial x'_D} \\ \frac{\partial K(\mathbf{X}, \mathbf{X}')}{\partial x_2} & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_2 \partial x'_1} & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_2 \partial x'_2} & \dots & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_2 \partial x'_D} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial K(\mathbf{X}, \mathbf{X}')}{\partial x_D} & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_D \partial x'_1} & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_D \partial x'_2} & \dots & \frac{\partial^2 K(\mathbf{X}, \mathbf{X}')}{\partial x_D \partial x'_D} \end{bmatrix}.$$

Correspondingly, the mean and variance of the posterior predictive distribution of the partial derivative of the energy with respect to coordinate d at \mathbf{x}^* are given as

$$\mathbb{E} \left[\frac{\partial f(\mathbf{x}^*)}{\partial x_d^*} \middle| \mathbf{y}_{\text{ext}}, \mathbf{X}, \boldsymbol{\theta} \right] = \frac{\partial \mathbf{K}_{\text{ext}}^*}{\partial x_d^*} (\mathbf{K}_{\text{ext}} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_{\text{ext}} \quad (3)$$

and

$$\begin{aligned} \text{Var} \left[\frac{\partial f(\mathbf{x}^*)}{\partial x_d^*} \middle| \mathbf{y}_{\text{ext}}, \mathbf{X}, \boldsymbol{\theta} \right] \\ = \frac{\partial^2 k(\mathbf{x}^*, \mathbf{x}^{*'})}{\partial x_d^* \partial x_d^{*'}} - \frac{\partial \mathbf{K}_{\text{ext}}^*}{\partial x_d^*} (\mathbf{K}_{\text{ext}} + \sigma^2 \mathbf{I})^{-1} \frac{\partial \mathbf{K}_{\text{ext}}^{*T}}{\partial x_d^{*'}}. \end{aligned}$$

Even if the observations are assumed to be accurate, a small but positive value for the noise variance σ^2 is used to avoid numerical problems when inverting the covariance matrix \mathbf{K}_{ext} .

IV. ALGORITHMS

In this section, the two algorithms using the Gaussian process regression (GPR) approach to accelerate CI-NEB calculations, the all-images-evaluated (AIE) algorithm and the one-image-evaluated (OIE) algorithm, are presented in detail.

A. All-images-evaluated (AIE) algorithm

Input: A GP model, energy (and zero force) at the two minima on the energy surface, coordinates of the N_{im} images on the initial path, a final convergence threshold T_{MEP} for the minimum energy path, an additional final convergence threshold T_{CI} for the climbing image i_{CI} , a preliminary convergence threshold $T_{\text{Clon}}^{\text{GP}}$ for turning climbing image mode on during the relaxation phase, a maximum displacement r_{max} of any image from the nearest observed data point.

Output: A minimum energy path represented by N_{im} images, one of which has climbed to the highest saddle point.

1. Start from the initial path and repeat the following (outer iteration loop):
 - A. Evaluate the true energy and force at the $N_{\text{im}} - 2$ intermediate images of the current path and add them to the training data.

- B. Calculate the accurate NEB force vector $F_{\text{NEB}}(i)$ for each intermediate image $i \in \{2, 3, \dots, N_{\text{im}} - 1\}$.
- C. If $\max_i |F_{\text{NEB}}(i)| < T_{\text{MEP}}$ and $|F_{\text{NEB}}(i_{\text{CI}})| < T_{\text{CI}}$, then stop the algorithm (final convergence reached).
- D. Optimize the hyperparameters of the GP model based on the training data and calculate the matrix inversion in Eq. (1).
- E. Relaxation phase: Start from the initial path, set climbing image mode off, and repeat the following (inner iteration loop):
 - I. Calculate the GP posterior mean energy and gradient at the intermediate images using Eqs. (1) and (3).
 - II. Calculate the approximate NEB force vector $F_{\text{NEB}}^{\text{GP}}(i)$ for each intermediate image using the GP posterior mean gradient.
 - III. If climbing image mode is off and $\max_i |F_{\text{NEB}}^{\text{GP}}(i)| < T_{\text{Clon}}^{\text{GP}}$, then turn climbing image mode on and recalculate the approximate NEB force vectors.
 - IV. If climbing image mode is on and $\max_i |F_{\text{NEB}}^{\text{GP}}(i)| < \frac{1}{10} T_{\text{CI}}$, then stop the relaxation phase (E).
 - V. Move the intermediate images along the approximate NEB force vector $F_{\text{NEB}}^{\text{GP}}(i)$ with a step size defined by the projected velocity Verlet algorithm.
 - VI. If the distance from any current image to the nearest observed data point is larger than r_{max} , then reject the last inner step and stop the relaxation phase (E).

A pseudocode for the AIE algorithm is presented above. Figure 1 shows an illustration of the progression of the algorithm on the two-dimensional Müller-Brown energy surface.²⁵ The energy (and the zero gradient of the energy) at the initial and final state minima is assumed to be provided as input.

The algorithm is started by evaluating the energy and force at the $N_{\text{im}} - 2$ intermediate images of an initial path and constructing a GP model for the energy based on the obtained information. The path is then relaxed on the approximate energy surface (GPR iteration 1), which is given as the posterior mean of the GP model, with a regular CI-NEB method using the posterior mean gradient of the GP model to calculate the approximate NEB force at the images. After each relaxation phase, final convergence of the path is checked by evaluating the true energy and force at the images of the relaxed

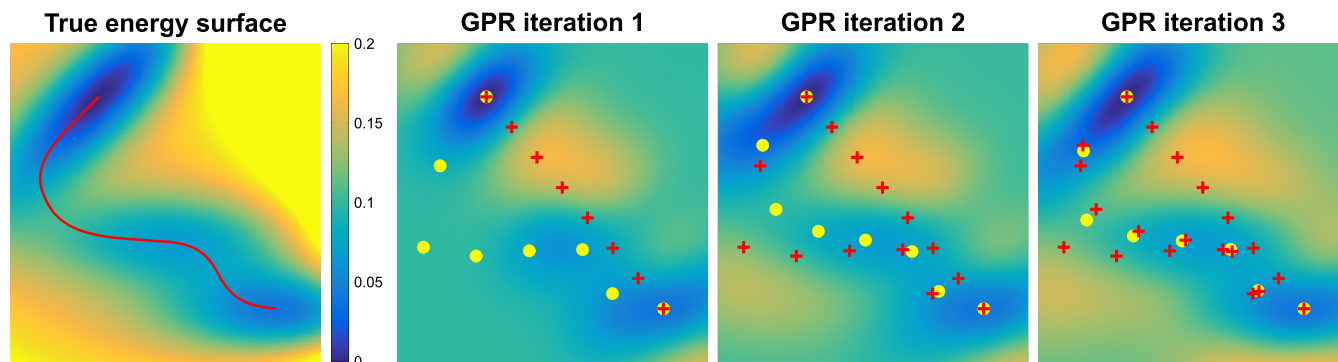


FIG. 1. Far left: The two-dimensional Müller-Brown energy surface, which has three minima, and the minimum energy path (red curve). Three panels to the right: An illustration of the iterative construction of an approximate energy surface in the vicinity of the minimum energy path using the all-images-evaluated algorithm where energy and atomic forces are evaluated at all intermediate images of the nudged elastic band. The initial path is a straight line interpolation between the initial and final state minima. The red + signs show the points at which the energy and atomic forces have been evaluated. The yellow disks show the climbing image nudged elastic band relaxed on the approximate energy surface of each Gaussian process regression iteration. After each GPR iteration, final convergence of the path is checked by energy and force evaluations, which are then added to the training data for the following GPR iteration. After three iterations (and a total of 24 energy and force evaluations), final convergence is confirmed as the magnitude of the NEB force is below the threshold 0.01, both for the climbing image and the other intermediate images.

path, and these observations are then added to the training data to improve the GP model on the following round. As can be seen from Fig. 1, a fairly accurate approximation of the Müller-Brown energy surface is obtained already after three GPR iterations of the AIE algorithm. This corresponds to 18 energy and force evaluations since the path is represented by six movable images in this case. A simplified flowchart of the algorithm is presented in Fig. 2.

The final convergence of the climbing image nudged elastic band is defined by the magnitude of the NEB force (including the spring force parallel to the path tangent) at each image calculated using the energy and force of the true energy surface. Two separate final convergence thresholds are used: $T_{\text{MEP}}^{\text{GP}}$ for the maximum NEB force magnitude $\max_i |F_{\text{NEB}}(i)|$ among the intermediate images i and a tighter T_{CI} for the NEB force magnitude $|F_{\text{NEB}}(i_{\text{CI}})|$ of the climbing image i_{CI} .

To ensure that the incomplete relaxation of the path on the approximate energy surface does not disturb final convergence, the convergence threshold $T_{\text{MEP}}^{\text{GP}}$ for the maximum approximate NEB force magnitude $\max_i |F_{\text{NEB}}^{\text{GP}}(i)|$ during the relaxation phase is defined as $1/10$ of the tighter final threshold T_{CI} . To decrease the amount of inner iterations during the relaxation phase, there is an alternative option for $T_{\text{MEP}}^{\text{GP}}$ to be defined as $1/10$ of the smallest true NEB force magnitude obtained so far on any of the intermediate images (but not

less than $T_{\text{CI}}/10$). If the approximation error is assumed not to decrease more than that during one GPR iteration, there is no need for a tighter convergence on the approximated surface and thus the relaxation phase can be sped up by using a larger tolerance. The divisor 10 can also be replaced by some other suitable number.

To make it more certain that the path converges to the same MEP as the one obtained by the regular CI-NEB method, each relaxation phase on the approximate energy surface is started from the same initial path. The relaxation is first conducted without climbing image mode until a preliminary convergence threshold $T_{\text{Clon}}^{\text{GP}}$ is reached and then continued from the preliminary evenly spaced path with climbing image mode turned on. Starting each relaxation phase from the initial path may possibly improve the stability of the algorithm, but there is also an alternative option which may decrease the number of inner iterations. In this alternative scheme, each relaxation phase would be started from the latest evenly spaced path converged to $T_{\text{Clon}}^{\text{GP}}$, and the climbing image phase would be started from the latest converged CI-NEB path if the climbing image of that path has the highest energy also on the current approximate energy surface. If, instead, the index of the highest energy image has changed, the climbing image phase would be started normally from the preliminarily relaxed evenly spaced path.

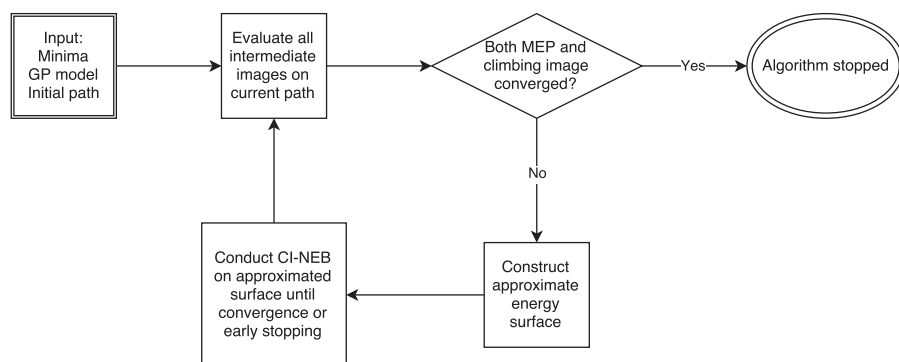


FIG. 2. A flowchart of the all-images-evaluated algorithm, where energy and atomic forces are evaluated at all intermediate images of the climbing image nudged elastic band relaxed on the GP-approximated energy surface.

In the early phases of the algorithm, when little information is available about the energy surface, there is a greater possibility that the path wanders far away from the initial path. To prevent this behavior, it is good to have some early stopping rule for the movement of the path. Thus, if the distance from any current image to the nearest observed data point becomes larger than r_{\max} , then the last inner step is rejected and the relaxation phase stopped. In the heptamer island benchmark described later in Sec. V, r_{\max} is defined as half of the length of the initial path (sum of the distances between adjacent images), but other definitions, e.g., based on the length scale of the GP model, are also possible.

When the final goal is to estimate the transition rates using harmonic transition state theory, the Hessian matrices at the initial and final state minima will need to be calculated. The Hessian is usually calculated with a finite difference method, where energy and force evaluations are made in the neighborhood of the minima. If these calculations, which anyway are needed for the Hessian, are evaluated already in the beginning of the MEP calculation, the calculated values can be added to the initial data set for the GPR calculations to provide information about the shape of the energy surface around the endpoints of the path and improve especially the early phase of the algorithm. To test the effect of the Hessian input in the heptamer island benchmark, a finite difference displacement of 10^{-3} Å is made in the positive direction along each of the atom coordinates, and the values of the energy and force at these points are included as input in the GPR calculation.

Since the GPR calculations using gradient observations require an inversion of an $(1 + D)N \times (1 + D)N$ matrix, the computational effort scales as $\mathcal{O}(((1 + D)N)^3)$, where N is the number of observation points and D is the number of degrees of freedom (here coordinates of movable atoms). As usual, the matrix inversion is computed by forming a Cholesky decomposition and solving a linear system of equations. Since the model stays the same during the relaxation phase on the approximate energy surface, the matrix inversion needs to be computed only once for each GPR iteration and the complexity of one inner iteration on the approximate energy surface is $\mathcal{O}((1 + D)N)$.

B. One-image-evaluated (OIE) algorithm

Input: A GP model, energy (and zero force) at the two minima on the energy surface, coordinates of the N_{im} images on the initial path, a final convergence threshold T_{MEP} for the minimum energy path, an additional final convergence threshold T_{CI} for the climbing image i_{CI} , a preliminary convergence threshold $T_{\text{Clon}}^{\text{GP}}$ for turning climbing image mode on during the relaxation phase, a maximum displacement r_{\max} of any image from the nearest observed data point.

Output: A minimum energy path represented by N_{im} images, one of which has climbed to the highest saddle point.

1. Optimize the hyperparameters of the GP model based on the initial data.
2. Start from the initial path and repeat the following (outer iteration loop):

- A. Calculate the GP posterior variance at the unevaluated images $i \in I_{\text{u}}$ on the current path using Eq. (2).
- B. Evaluate the true energy and force at the image with highest posterior variance and add them to the training data.
- C. Calculate the accurate NEB force vector $F_{\text{NEB}}(i)$ for the evaluated images $i \in I_{\text{e}}$.
- D. If all images on the current path have been evaluated, $\max_i |F_{\text{NEB}}(i)| < T_{\text{MEP}}$ and $|F_{\text{NEB}}(i_{\text{CI}})| < T_{\text{CI}}$, then stop the algorithm (final convergence reached).
- E. Reoptimize the GP hyperparameters, calculate the GP posterior mean energy and gradient at the unevaluated images $i \in I_{\text{u}}$ using Eqs. (1) and (3), and save the matrix inversion for further use.
- F. Calculate the approximate NEB force vector $F_{\text{NEB}}^{\text{GP}}(i)$ for the unevaluated images $i \in I_{\text{u}}$ using the GP posterior mean gradient and set $F_{\text{NEB}}^{\text{GP}}(i) = F_{\text{NEB}}(i)$ for the evaluated images $i \in I_{\text{e}}$.
- G. If $\max_i |F_{\text{NEB}}^{\text{GP}}(i)| < T_{\text{MEP}}$:
 - I. If $i_{\text{CI}} \in I_{\text{u}}$, then evaluate the energy and force at the climbing image i_{CI} , add them to the training data, and go to C.
 - II. If $i_{\text{CI}} \in I_{\text{e}}$ and $|F_{\text{NEB}}(i_{\text{CI}})| < T_{\text{CI}}$, then go to A.
 - III. If $i_{\text{CI}} \in I_{\text{e}}$ and $|F_{\text{NEB}}(i_{\text{CI}})| \geq T_{\text{CI}}$, then execute the relaxation phase (H), evaluate the energy and force at the climbing image, add them to the training data, and go to C.
- H. Relaxation phase: Start from the initial path, set climbing image mode off, and repeat the following (inner iteration loop):
 - I. Calculate the GP posterior mean energy and gradient at the intermediate images using Eqs. (1) and (3).
 - II. Calculate the approximate NEB force vector $F_{\text{NEB}}^{\text{GP}}(i)$ for each intermediate image using the GP posterior mean gradient.
 - III. If climbing image mode is off and $\max_i |F_{\text{NEB}}^{\text{GP}}(i)| < T_{\text{Clon}}^{\text{GP}}$, then turn climbing image mode on and recalculate the approximate NEB force vectors.
 - IV. If climbing image mode is on and $\max_i |F_{\text{NEB}}^{\text{GP}}(i)| < T_{\text{MEP}}^{\text{GP}} = \frac{1}{10} T_{\text{CI}}$, then stop the relaxation phase (H).
 - V. Move the intermediate images along the approximate NEB force vector $F_{\text{NEB}}^{\text{GP}}(i)$ with a step size defined by the projected velocity Verlet algorithm.
 - VI. If the distance from any current image i to the nearest observed data point is larger than r_{\max} , then reject the last inner step, evaluate image i , and go to C.

A pseudocode for the OIE algorithm is presented above. Figure 3 shows an illustration of the progression of the algorithm on the Müller-Brown energy surface. Both the GP approximation to the energy surface and the estimated uncertainty after one, two, three, and eleven GPR iterations are shown. The energy (and the zero gradient of the energy) at the initial and final state minima is assumed to be given as input.

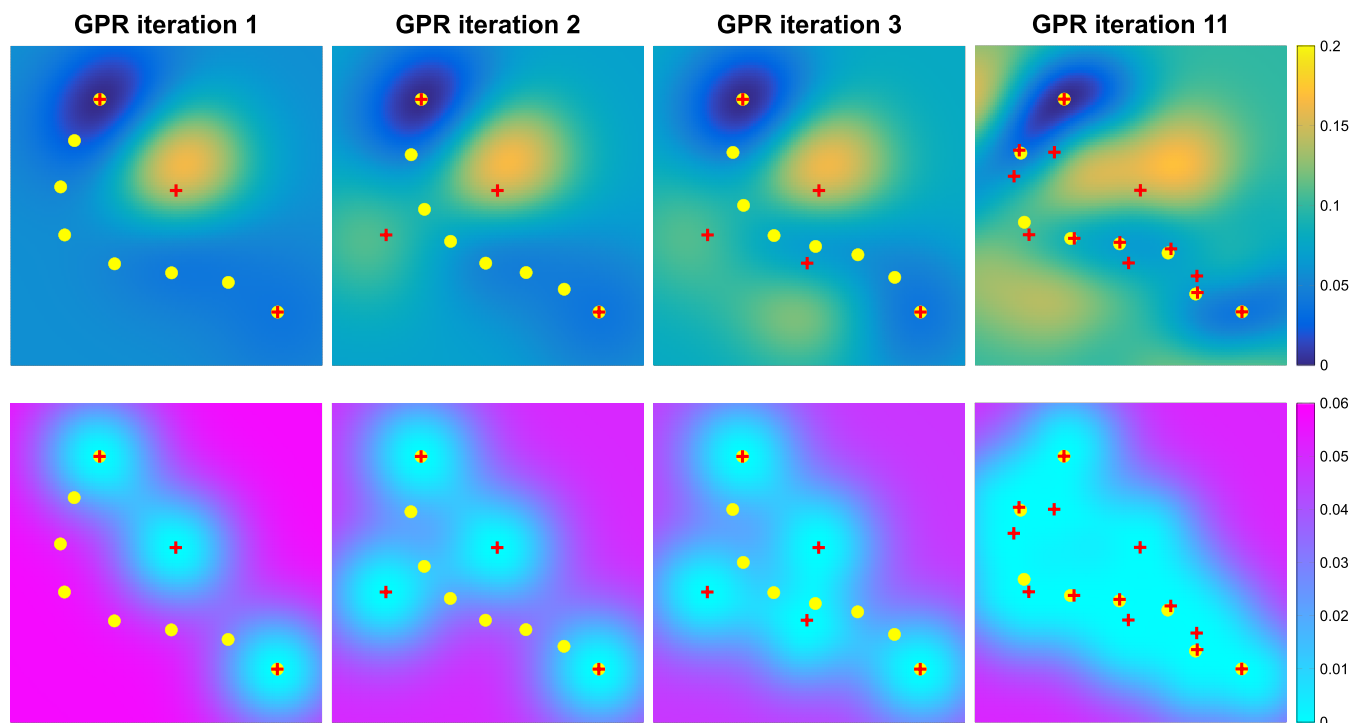


FIG. 3. An illustration of the iterative construction of an approximate energy surface to the two-dimensional Müller-Brown energy surface (shown in Fig. 1) in the vicinity of the minimum energy path using the one-image-evaluated algorithm where the energy and atomic forces are evaluated only at one image of the nudged elastic band. The initial path is a straight line interpolation between the initial and final state minima. Upper panel: The red + signs show the points at which the energy and atomic forces have been evaluated. The yellow disks show the climbing image nudged elastic band relaxed on the approximate energy surface of each Gaussian process regression iteration. After GPR iterations 1, 2, and 3, the energy and force are calculated at the image where the estimated uncertainty is largest and the observed data are then added to the training data set for the following GPR iteration. Lower panel: The estimated uncertainty (standard deviation) of the energy approximation shown directly above. After eleven iterations, the path is not displaced further but the final convergence is checked by evaluating the energy and force at each intermediate image one by one. After 17 evaluations, final convergence is confirmed as the magnitude of the NEB force is below the threshold 0.01 both for the climbing image and the other intermediate images.

The algorithm is started by evaluating the true energy and force at the image located in the middle of the initial path where the uncertainty of the initial GP model is largest. The GP model is then updated based on the obtained information, and the path is relaxed on the approximate energy surface (GPR iteration 1). After each relaxation phase, the true energy and force are evaluated at only one image of the relaxed path before updating the GP model. According to the main rule, the image with the highest uncertainty estimate is selected for evaluation, and the information obtained is then used to improve the GP model on the following round. As can be seen from Fig. 3, a fairly accurate approximation of the Müller-Brown energy surface is obtained already after eleven GPR iterations of the OIE algorithm. This corresponds to only eleven energy and force evaluations, quite a bit fewer than the 18 included in the three GPR iterations of the AIE algorithm shown in Fig. 1.

The details of the OIE algorithm are otherwise similar to the AIE algorithm, but since only one image is evaluated between the GPR iterations, relaxing the path between each evaluation would mean that the accurate NEB force is only known for one image at a time. Thus, it would not be known for sure whether the path has converged on a true MEP. Approximations for the NEB forces can of course be calculated at the unevaluated images based on the updated GP model, but since the NEB forces have been relaxed to zero based on the previous

GP approximation and since the largest changes to the approximation usually emerge near the new observation point, it is most likely that these approximations underestimate the NEB force magnitudes. The approximated NEB forces, however, together with the accurate ones, at least indicate if there is a possibility that the path may have converged. The general idea for the convergence check of the OIE algorithm is that when the maximum magnitude of both the accurate and approximated NEB forces is below the final convergence threshold T_{MEP} , more images are evaluated without moving the path until some of the magnitudes rise above the threshold or all images have been evaluated. Since the images with the highest uncertainty, which are the most likely ones to violate the convergence criterion, are evaluated first, it is likely that the convergence check will be interrupted early if the path has not yet truly converged.

The special role of the climbing image makes the evaluation rules a bit more complicated. Since the climbing image has a tighter convergence threshold and since the position of the climbing image also affects the distribution of the other images, it is desirable to favor evaluations of the climbing image during the convergence check. As long as the maximum magnitude of the accurate and approximated NEB forces is above T_{MEP} , the GP relaxation phase is executed normally and the image with the highest uncertainty is evaluated. After the maximum magnitude has reduced below T_{MEP} , the

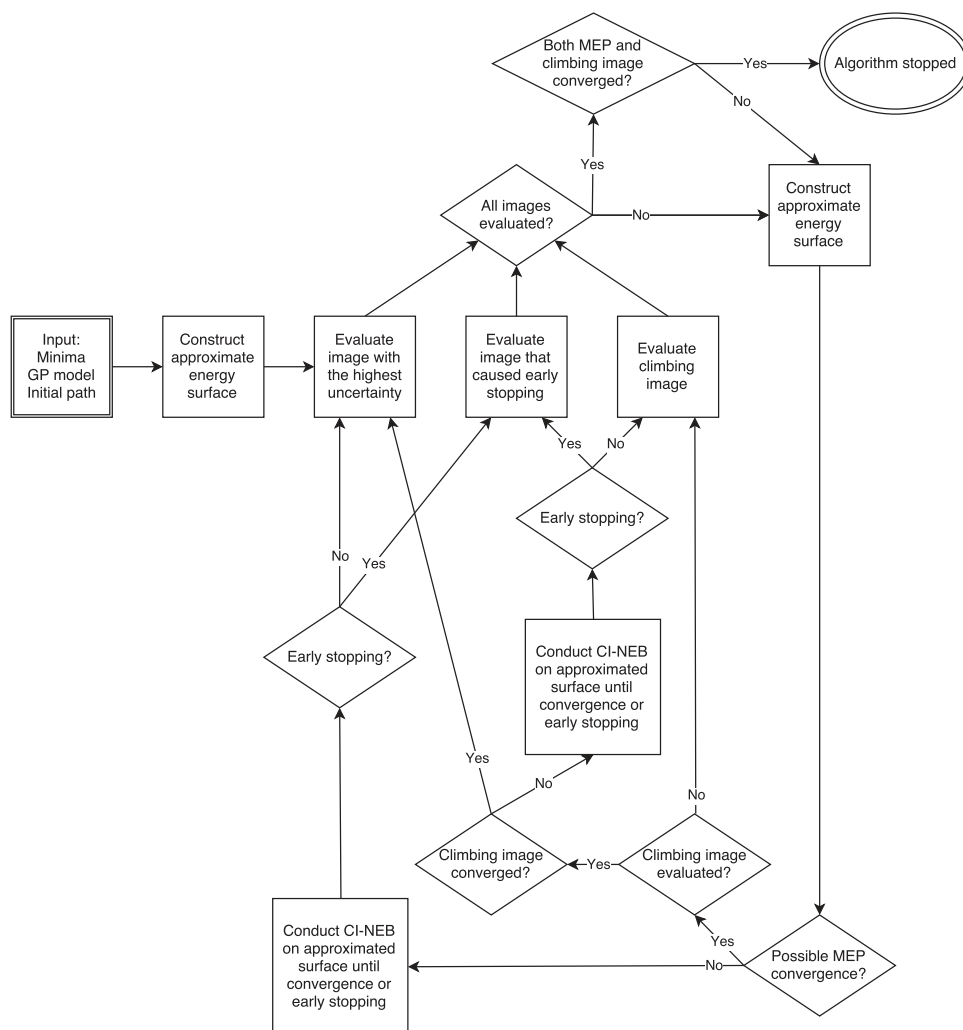


FIG. 4. A flowchart of the one-image-evaluated algorithm, where the energy and atomic forces are evaluated at only one image before a new Gaussian process regression iteration.

climbing image is evaluated without moving the path (if not already evaluated). As long as the maximum NEB force magnitude stays below T_{MEP} but the NEB force magnitude $|F_{\text{NEB}}(i_{\text{CI}})|$ on the climbing image is above T_{CI} , the path is relaxed and the climbing image is re-evaluated. Finally, if the maximum magnitude of the accurate and approximated NEB forces is below T_{MEP} , the climbing image has been evaluated and $|F_{\text{NEB}}(i_{\text{CI}})| < T_{\text{CI}}$, then more images are evaluated without moving the path, starting from the image with the highest uncertainty.

Another exception to the selection of the image to be evaluated is caused by the early stopping rule during the relaxation phase. If the distance from any current image i to the nearest observed data point becomes larger than r_{max} , then the last inner step is rejected, the relaxation phase stopped, and image i evaluated next. A simplified flowchart of the OIE algorithm is presented in Fig. 4.

V. APPLICATION TO THE HEPTAMER ISLAND BENCHMARK

A test problem that has been used in several studies of algorithms for finding MEPs and saddle points involves an island of seven atoms on the (111) surface of a face-centered

cubic (FCC) crystal.^{34,35} The interaction between the atoms is described with a simple Morse potential to make the implementation of the benchmark easy. The initial, saddle point and final configurations of the atoms for the 13 lowest activation energy transitions, labeled from *A* to *M*, are shown in Fig. 5. In the initial state, the seven atoms sit at FCC surface sites and form a compact island. In transitions *A* and *B*, the whole island shifts to HCP sites on the surface. In some of the other transitions, a pair of edge atoms slides to adjacent FCC sites, an atom half way dissociates from the island, or a pair of edge atoms moves in such a way that one of the atoms is displaced away from the island while the other atom takes its place.

The calculations were carried out using five intermediate images ($N_{\text{im}} = 7$) in the CI-NEB calculations starting with an IDPP path, and the images were moved iteratively to an MEP using the projected velocity Verlet algorithm¹⁴ with a time step of 0.1 fs. A time step of 1 fs is too large and leads to overshooting, but a time step of 0.01 fs requires a significantly larger number of iterations. The algorithms were continued until the magnitude of the true NEB force acting on the climbing image had dropped below $T_{\text{CI}} = 0.01$ eV/Å. A larger tolerance, $T_{\text{MEP}} = 0.3$ eV/Å, was used for the magnitude of the NEB force acting on the other images in the CI-NEB calculation. During each relaxation phase on the GP-approximated surface, the

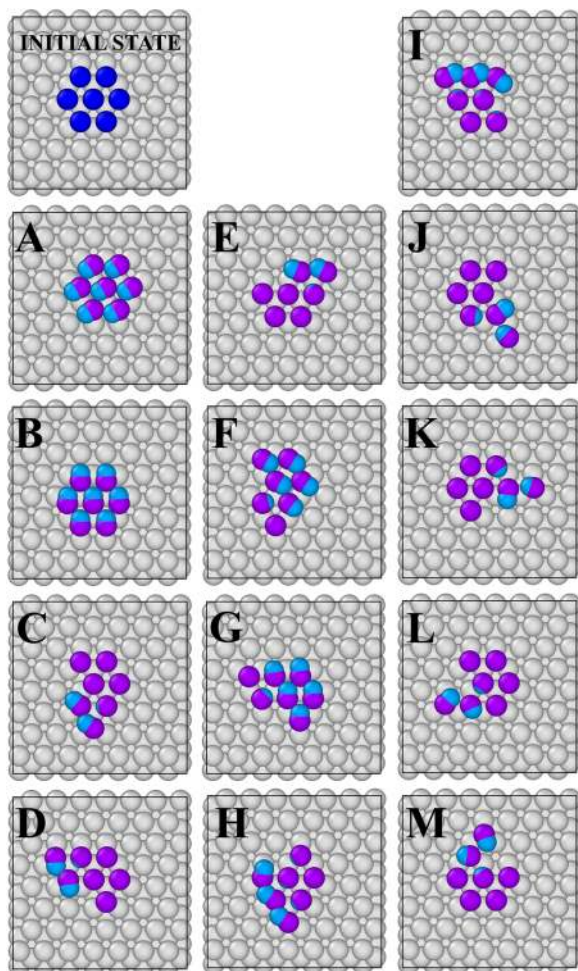


FIG. 5. An illustration of the atomic transitions of the heptamer island benchmark problem. The initial configuration involves a seven-atom island (dark blue disks in the uppermost left column) adsorbed at FCC sites of an FCC(111) surface. Saddle point configuration (light blue disks) and final configuration (purple disks) are shown together for each of the transitions, labeled A–M.

preliminary convergence threshold $T_{\text{Clon}}^{\text{GP}}$ for turning climbing image mode on was $1 \text{ eV}/\text{\AA}$. The GPR calculations were carried out using the GPstuff toolbox.³⁶ The fixed parameters of the GP model were chosen to be $\sigma^2 = 10^{-8} \text{ eV}^2$ and $\sigma_c^2 = 100 \text{ eV}^2$. A common length scale $l_d = l$ was used for all dimensions $d = 1, \dots, D$, and a zero mean Student's t -distribution (restricted to positive values) with 1 \AA^2 scale and four degrees of freedom was used as a prior distribution for l and a log-uniform distribution for σ_m^2 (i.e., the default priors of GPstuff) in the optimization of the hyperparameters. The prior distributions stabilize the point estimates of the hyperparameters especially in the beginning, when there is little data available. The optimization was performed using the scaled conjugate gradient algorithm.³⁷

By using input from the Hessian at the initial and final state minima, the path relaxed on the GP-approximated energy surface can become qualitatively similar to the true MEP with fewer GPR iterations. This is illustrated for transition *I* in Fig. 6. It shows the true energy evaluated at the location of the images of the initial path, the converged MEP, and the path after one GPR iteration in the AIE algorithm with and

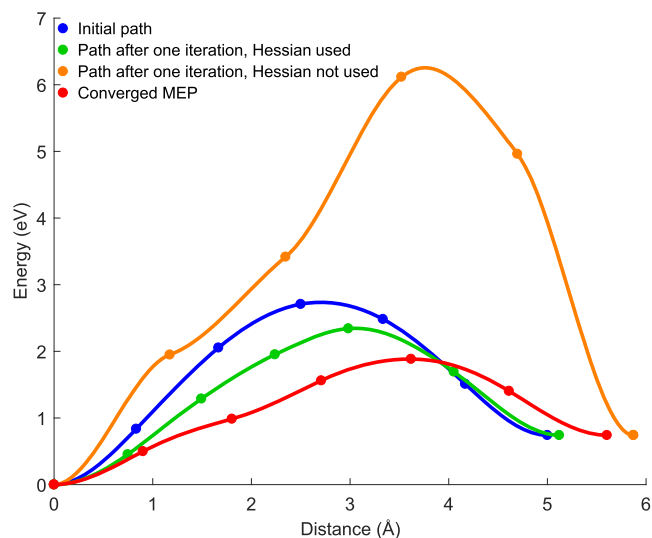


FIG. 6. An illustration of the improvement that can be obtained by using the Hessian at the initial and final state minima. The calculations are for transition *I* when the substrate atoms are frozen. The green and orange points show the true energy evaluated at the location of images of a climbing image nudged elastic band relaxed on the first approximate energy surface of the all-images-evaluated algorithm, i.e., at this point, the energy and force had been evaluated at all the images of the initial path. The path obtained when the Hessian is not used (orange) is far from the converged minimum energy path (red), quite a bit farther away than the initial path (blue), and the climbing image has moved to the final state minimum. The path obtained when the Hessian is used (green), however, represents a clear improvement to the initial path, having moved significantly closer to the converged MEP. Eventually, the path converges after 14 Gaussian process regression iterations (a total of 75 energy and force evaluations required to confirm final convergence) when the Hessian is used, and after 17 GPR iterations when it is not used.

without input from the Hessian. The estimate without the Hessian input is quite poor at this point, the path reaching an area of high energy and the climbing image moving to the final state minimum.

The reason for this behavior can be seen from Fig. 7, where the GP-approximated energy at the location of the images after one GPR iteration (five energy and force evaluations) is shown. The maximum of the approximated energy along the path is indeed at the final state of the path. However, with the input from the Hessian, a qualitatively correct path is obtained already after one GPR iteration. In this case, the information coming from the Hessian about the curvature at the endpoints ensures that the GP-approximated surface has minima at those points. Without the Hessian input, the three subsequent GPR iterations still show qualitatively wrong variation of the energy along the path, and it is only after the fifth iteration that enough input has been obtained for the GP model to be reasonably accurate to produce a path qualitatively similar to the true MEP.

The overall reduction in the number of energy and force evaluations corresponds essentially to the first three GPR iterations that can be skipped when the Hessian is provided. It takes 17 GPR iterations without and 14 iterations with the Hessian input to reach convergence in this case. The energy and force evaluations needed to construct the Hessian are not counted in the cost of finding the MEP since they need to be carried out anyway if the transition rate is estimated using harmonic transition state theory. While the reduction in the number of

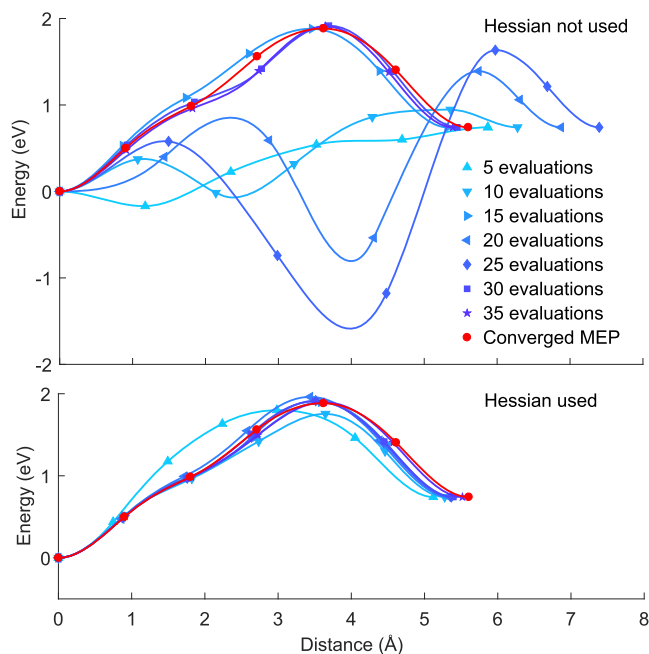


FIG. 7. Comparison of the performance of the all-images-evaluated algorithm with and without input from the Hessian at the initial state minima. The markers show the GP-approximated energy at the images of the climbing image nudged elastic band relaxed on the approximate energy surface after one to seven Gaussian process regression iterations (5 to 35 energy and force evaluations, blue) and after final convergence (red) for transition I when the substrate atoms are frozen. Without the Hessian input (upper graph), the approximate energy surface does not have an intermediate barrier and the climbing image moves to the final state minimum. The true energy evaluated at each image along this path is shown in Fig. 6 and shows a large energy barrier. It takes six GPR iterations (30 energy and force evaluations) before the path relaxed on the approximated surface starts to look qualitatively similar to the converged minimum energy path. With the Hessian input (lower graph), the energy along the path relaxed on the approximate energy surface is qualitatively similar to that of the converged MEP already after one iteration.

energy and force evaluations (here 15) corresponds to savings of about 20%, the importance of the Hessian input can be greater in more challenging systems when avoiding exploration of regions far away from the MEP where the atomic forces tend to be large.

Even larger savings are obtained by using the uncertainty estimate provided by the GP model to evaluate the energy and force only at the image where the true energy is the most poorly known instead of all images, i.e., to move to the OIE from the AIE algorithm. This is illustrated for transition F in Fig. 8, where the GP-approximated energy at the CI-NEB images is shown after a certain number of energy and force evaluations for the AIE and OIE algorithms. In the AIE case, the path is still qualitatively incorrect after 20 evaluations (four GPR iterations), while this suffices for the OIE algorithm (20 GPR iterations) to nearly reach convergence. For the AIE algorithm, it takes a total of 75 energy and force evaluations to confirm final convergence, while the OIE algorithm requires 39. Similar reduction in the number of energy and force evaluations was found for the other transitions.

The number of energy and force evaluations required to confirm final convergence for each of the 13 transitions using the AIE and OIE algorithms is given in Table I as a

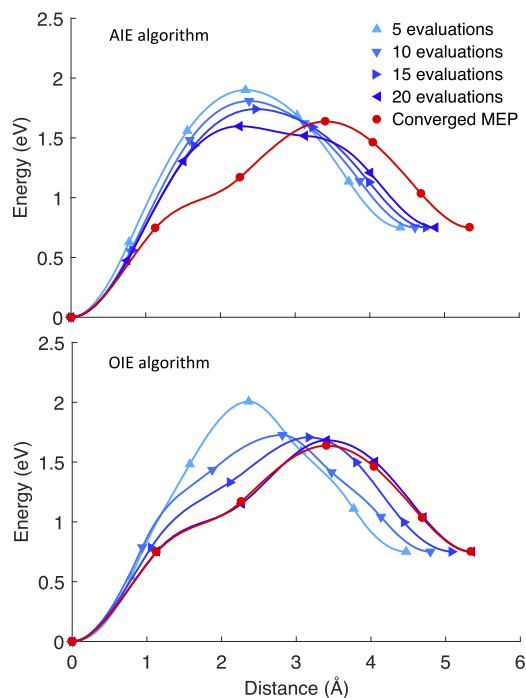


FIG. 8. Comparison of the performance of the all-images-evaluated (AIE) and the one-image-evaluated (OIE) algorithms (Hessian input used) for transition F when the substrate atoms are frozen. The markers show the GP-approximated energy at the images of the climbing image nudged elastic band relaxed on the approximate energy surface after 5, 10, 15, and 20 energy and force evaluations (blue) and after final convergence (red). With the AIE algorithm (upper figure), the variation of the energy is significantly different for the path on the approximate surface compared with the true minimum energy path even after 20 evaluations (four Gaussian process regression iterations). With the OIE algorithm (lower figure), the variation of the energy along the path on the approximate surface is close to that of the converged MEP after 20 evaluations (corresponding to 20 GPR iterations in this case).

fraction of the number of evaluations required by a regular CI-NEB method. Also, the effect of using the Hessian input for the AIE algorithm is shown. These numbers correspond to the case where six nearest substrate atoms can move in addition to the seven island atoms. The relative number of evaluations compared to the regular CI-NEB varies between the transitions. A clear trend is that the more complex the transition and the more the iterations required by the regular CI-NEB, the larger is the relative effect of the GPR approach.

The average number of energy and force evaluations for transitions C – M as a function of the number of degrees of freedom is shown in Fig. 9.³⁸ For the smallest number of degrees of freedom, 21, only the seven island atoms are allowed to move while all the substrate atoms are frozen. For the larger numbers of degrees of freedom, some of the surface atoms are also allowed to move. Starting with the AIE algorithm, the use of the Hessian input reduces the number of evaluations by about 20%, but the transition to the OIE algorithm has an even larger effect, a reduction to a half.

The OIE results represent savings of an order of magnitude with respect to the regular CI-NEB calculation. The number of energy and force evaluations reported here for the CI-NEB method is similar to what has been reported earlier for this test problem.^{34,35} It is possible to use a more efficient minimization

TABLE I. The number of energy and force evaluations needed to converge the regular climbing image nudged elastic band (CI-NEB) calculations of the heptamer island benchmark transitions (shown in Fig. 5) when 39 degrees of freedom are included, and the reduction in the number of evaluations obtained with the Gaussian process regression approach using the all-images-evaluated algorithm without the Hessian input (AIE), all-images-evaluated algorithm with the Hessian input (AIE-H), and one-image-evaluated algorithm without the Hessian input (OIE).

Transition	Number of evaluations		Number of evaluations as a fraction of evaluations needed for CI-NEB	
	CI-NEB	AIE	AIE-H	OIE
A	120	0.42	0.42	0.13
B	120	0.42	0.42	0.23
C	285	0.25	0.21	0.13
D	265	0.26	0.23	0.14
E	290	0.24	0.19	0.13
F	855	0.12	0.12	0.05
G	840	0.13	0.11	0.05
H	1480	0.08	0.07	0.04
I	1480	0.07	0.07	0.04
J	605	0.15	0.12	0.07
K	610	0.14	0.12	0.07
L	565	0.17	0.12	0.06
M	570	0.17	0.11	0.06

scheme to relax the images in CI-NEB calculations,²⁸ but the difference is not so large. The test results presented here, therefore, show that the use of GPR can significantly reduce the computational effort in, for example, calculations of MEPs for surface processes.

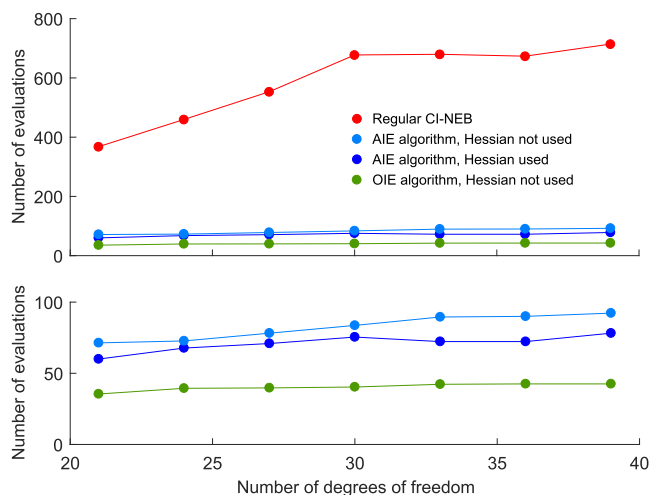


FIG. 9. Average number of energy and force evaluations needed to converge to a minimum energy path in climbing image nudged elastic band calculations of the heptamer island benchmark as a function of the number of degrees of freedom, increased by allowing a larger number of substrate atoms to move. The lower figure shows the same results as the upper figure on a different scale to better distinguish between the various implementations of the Gaussian process regression approach. For the larger numbers of degrees of freedom, the one-image-evaluated (OIE) algorithm provides about 1/20 reduction in the number of energy and force evaluations as compared with a regular CI-NEB calculation. The all-images-evaluated (AIE) algorithm requires about twice as many evaluations as the OIE algorithm, but the use of the Hessian at the initial and final state minima can reduce that by about 20%. The use of the Hessian has less effect when the OIE algorithm is used (not shown).

VI. DISCUSSION

The results presented here show that the GPR approach can reduce the number of energy and force evaluations needed in CI-NEB calculations of MEPs by an order of magnitude. This is important since a large amount of computer time is used in such calculations, especially when *ab initio* or density functional theory calculations are used. Compared with the previous proof-of-principle calculations,²¹ three major improvements to the algorithm have been presented here. First of all, the CI-NEB algorithm was used where one of the images is pushed up to the maximum along the MEP. This provides stability and accelerates convergence because it focuses more evaluations of the energy and force in the vicinity of the first-order saddle point, the most important part of the energy surface.

Second, the benefit of using the finite difference estimate of the Hessian at the endpoints was demonstrated and found to result in a 20% reduction in the number of energy and function evaluations needed to converge on an MEP with the AIE algorithm. This estimation of the Hessian does not represent any additional energy and force evaluations in cases where the goal is to calculate the transition rates using harmonic transition state theory. The usual ordering of the calculations is simply changed; the Hessian at the endpoint minima is evaluated before the MEP calculation rather than afterwards. If a higher level of rate theory, such as optimal hyperplanar transition state theory,⁶ is used, then analogous information about the initial and final states can be obtained from dynamical trajectories.

While the Hessian input reduced the number of energy and function evaluations only by 20%, a significant advantage is likely in greater stability and lower probability of the path escaping into irrelevant regions of the energy surface in the first few GPR iterations. This we have seen in preliminary studies of dissociative adsorption of H₂ on a metal surface and water molecule diffusion on an ice Ih(0001) surface.³⁹ The addition of the finite difference data points for the Hessian adds significantly to the memory requirements of the GPR calculation. Those data points could, however, be dropped after just a few GPR iterations since they are not needed when the GP approximation of the energy surface becomes reasonably accurate around the MEP. A more elegant and efficient way of incorporating the information from the Hessian into the covariance calculation could also be developed. The implementation described here represents just an initial test to see how important such information can be.

As a third improvement, a significant reduction in the computational effort was shown to be possible by using the one-image-evaluated algorithm instead of the earlier all-images-evaluated approach. The number of energy and force evaluations was reduced to a half in the heptamer benchmark. In the OIE algorithm, the true energy and force are evaluated only at one image, rather than all images along the path, before a new GPR iteration. A calculated choice of the location of each evaluation can be made based on the uncertainty estimate provided by the GP model. Interestingly, the use of the Hessian did not provide significant reduction in the number of energy and force evaluations required by the OIE algorithm.

This apparently stems from the fact that the OIE algorithm involves fewer evaluations in the early phase of the iterative GPR process where the approximation to the energy surface is poor.

The heptamer island benchmark studied here is a relatively simple example, and it will be important to test the GPR approach on more complex systems to be able to fully assess its utility and to develop the methodology further. On complex energy surfaces, there may exist multiple MEPs connecting the two endpoint minima, which would require some kind of sampling of MEPs.⁴⁰ Also, some systems may have multiple local minima and highly curved MEPs, which can lead to convergence problems unless a large number of images are included in the calculation. In systems where various types of molecular interactions are involved, the optimal length scale may vary depending on the location in the coordinate space. In such cases, it may be advantageous to use a GP model that allows different length scales in different parts of the space.

In order to tackle large systems, the scaling of the GPR calculations will need to be improved. A more efficient implementation could be obtained, e.g., by using a compactly supported covariance function to produce a sparse covariance matrix where data points far away from each other become independent.⁴¹ It may also be possible to reduce the dimensionality by using partially additive models, where the interaction term in the energy function for far away atoms is ignored. There is, however, a large set of important problems, such as calculations of catalytic processes often involving rather small molecules adsorbed on surfaces, where the complexity is comparable to the heptamer island benchmark and where the GPR approach is clearly going to offer a significant reduction in computational effort in NEB calculations of MEPs.

At a low enough temperature, quantum mechanical tunneling becomes the dominant transition mechanism, and the task is then to find a minimal action path.^{9,42,43} The effect of the GPR approach in tunneling path calculations could be even larger than for MEP calculations since each iteration involves more energy and force calculations (Feynman paths rather than points in configuration space) and thereby more data for the modeling. In addition to atomic rearrangements, it will be valuable to apply the GPR approach also to magnetic transitions where the magnetic properties of the system are evaluated by computationally intensive *ab initio* or density functional theory calculations. The relevant degrees of freedom in magnetic transitions are the angles defining the orientation of the magnetic vectors, and the task is again to find MEPs on the energy surface with respect to those angles.^{11–13}

ACKNOWLEDGMENTS

We thank Professor Andrew Peterson for helpful discussions. This work was supported by the Icelandic Research Fund and the Academy of Finland (Grant No. 278260). We acknowledge the computational resources provided by the Aalto Science-IT project and the Computer Services of the University of Iceland.

- ¹E. Wigner, *Trans. Faraday Soc.* **34**, 29 (1938).
- ²H. A. Kramers, *Physica* **7**, 284 (1940).
- ³J. C. Keck, “Variational theory of reaction rates,” in *Advance in Chemical Physics*, edited by I. Prigogine (John Wiley & Sons, 1967), Vol. 13, pp. 85–121.
- ⁴G. K. Schenter, G. Mills, and H. Jónsson, *J. Chem. Phys.* **101**, 8964 (1994).
- ⁵G. Mills, H. Jónsson, and G. K. Schenter, *Surf. Sci.* **324**, 305 (1995).
- ⁶G. H. Jóhannesson and H. Jónsson, *J. Chem. Phys.* **115**, 9644 (2001).
- ⁷T. Bligaard and H. Jónsson, *Comput. Phys. Commun.* **169**, 284 (2005).
- ⁸G. H. Vineyard, *J. Phys. Chem. Solids* **3**, 121 (1957).
- ⁹H. Jónsson, *Proc. Natl. Acad. Sci. U. S. A.* **108**, 944 (2011).
- ¹⁰P. F. Bessarab, V. M. Uzdin, and H. Jónsson, *Phys. Rev. B* **85**, 184409 (2012).
- ¹¹P. F. Bessarab, V. M. Uzdin, and H. Jónsson, *Z. Phys. Chem.* **227**, 1543 (2013).
- ¹²P. F. Bessarab, V. M. Uzdin, and H. Jónsson, *Phys. Rev. B* **89**, 214424 (2014).
- ¹³P. F. Bessarab, A. Skorodumov, V. M. Uzdin, and H. Jónsson, *Nanosyst.: Phys., Chem., Math.* **5**, 757 (2014).
- ¹⁴H. Jónsson, G. Mills, and K. W. Jacobsen, “Nudged elastic band method for finding minimum energy paths of transitions,” in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by B. J. Berne, G. Cicciotti, and D. F. Coker (World Scientific, 1998), pp. 385–404.
- ¹⁵P. F. Bessarab, V. M. Uzdin, and H. Jónsson, *Comput. Phys. Commun.* **196**, 335 (2015).
- ¹⁶A. A. Peterson, *J. Chem. Phys.* **145**, 074106 (2016).
- ¹⁷A. O’Hagan and J. F. C. Kingman, *J. R. Stat. Soc. B* **40**, 1 (1978).
- ¹⁸D. J. C. MacKay, “Introduction to Gaussian processes,” in *Neural Networks and Machine Learning*, edited by C. M. Bishop (Springer Verlag, 1998), pp. 133–166.
- ¹⁹R. M. Neal, “Regression and classification using Gaussian process priors (with discussion),” in *Bayesian Statistics*, edited by J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith (Oxford University Press, 1999), Vol. 6, pp. 475–501.
- ²⁰C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, 2006).
- ²¹O.-P. Koistinen, E. Maras, A. Vehtari, and H. Jónsson, *Nanosyst.: Phys., Chem., Math.* **7**, 925 (2016); a slightly corrected version is available as e-print [arXiv:1703.10423](https://arxiv.org/abs/1703.10423).
- ²²J. Lampinen and A. Vehtari, *Neural Networks* **14**, 257 (2001).
- ²³G. Henkelman, B. P. Uberuaga, and H. Jónsson, *J. Chem. Phys.* **113**, 9901 (2000).
- ²⁴B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, *Proc. IEEE* **104**, 148 (2016).
- ²⁵K. Müller and L. D. Brown, *Theor. Chim. Acta* **53**, 75 (1979).
- ²⁶S. Smidstrup, A. Pedersen, K. Stokbro, and H. Jónsson, *J. Chem. Phys.* **140**, 214106 (2014).
- ²⁷G. Henkelman and H. Jónsson, *J. Chem. Phys.* **113**, 9978 (2000).
- ²⁸D. Sheppard, R. Terrell, and G. Henkelman, *J. Chem. Phys.* **128**, 134106 (2008).
- ²⁹A. O’Hagan, “Some Bayesian numerical analysis,” in *Bayesian Statistics*, edited by J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith (Oxford University Press, 1992), Vol. 4, pp. 345–363.
- ³⁰C. E. Rasmussen, “Gaussian processes to speed up hybrid Monte Carlo for expensive Bayesian integrals,” in *Bayesian Statistics*, edited by J. M. Bernardo, A. P. Dawid, J. O. Berger, M. West, D. Heckerman, M. J. Bayarri, and A. F. M. Smith (Oxford University Press, 2003), Vol. 7, pp. 651–659.
- ³¹E. Solak, R. Murray-Smith, W. E. Leithead, D. J. Leith, and C. E. Rasmussen, “Derivative observations in Gaussian process models of dynamic systems,” in *Advances in Neural Information Processing Systems*, edited by S. Becker, S. Thrun, and K. Obermayer (MIT Press, 2003), Vol. 15, pp. 1057–1064.
- ³²J. Riihimäki and A. Vehtari, “Gaussian processes with monotonicity information,” in *Proceedings of Machine Learning Research*, edited by Y. W. Teh and M. Titterton (PMLR, 2010), Vol. 9, pp. 645–652.
- ³³A. P. Bartók and G. Csányi, *Int. J. Quantum Chem.* **115**, 1051 (2015).
- ³⁴G. Henkelman, G. H. Jóhannesson, and H. Jónsson, “Methods for finding saddle points and minimum energy paths,” in *Progress in Theoretical Chemistry and Physics*, edited by S. D. Schwartz (Kluwer Academic, 2000), Vol. 5, pp. 269–300.
- ³⁵S. T. Chill, J. Stevenson, V. Ruhle, C. Shang, P. Xiao, J. D. Farrell, D. J. Wales, and G. Henkelman, *J. Chem. Theory Comput.* **10**, 5476 (2014).
- ³⁶J. Vanhatalo, J. Riihimäki, J. Hartikainen, P. Jylänki, V. Tolvanen, and A. Vehtari, *J. Mach. Learn. Res.* **14**, 1175 (2013).

- ³⁷C. M. Bishop, “Scaled conjugate gradients,” in *Neural Networks for Pattern Recognition* (Clarendon Press, 1995), pp. 282–285.
- ³⁸Transitions *A* and *B* are not included in the averages shown in Fig. 9 because the regular CI-NEB required an anomalously large number of iterations for some of the intermediate numbers of degrees of freedom. The results of the GPR algorithms were, however, similar for all numbers of degrees of freedom tested here.
- ³⁹E. R. Batista and H. Jónsson, *Comput. Mater. Sci.* **20**, 325 (2001).
- ⁴⁰E. Maras, O. Trushin, A. Stukowski, T. Ala-Nissilä, and H. Jónsson, *Comput. Phys. Commun.* **205**, 13 (2016).
- ⁴¹J. Vanhatalo and A. Vehtari, “Speeding up the binary Gaussian process classification,” in *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI, 2010)*, edited by P. Gründwald and P. Spirtes (AUAI Press, 2010), pp. 623–632.
- ⁴²G. Mills, G. K. Schenter, D. E. Makarov, and H. Jónsson, *Chem. Phys. Lett.* **278**, 91 (1997).
- ⁴³G. Mills, G. K. Schenter, D. E. Makarov, and H. Jónsson, “RAW quantum transition state theory,” in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by B. J. Berne, G. Ciccotti, and D. F. Coker (World Scientific, 1998), pp. 405–421.