

Minimum energy path calculations with Gaussian process regression

O-P. Koistinen¹, E. Maras², A. Vehtari¹, H. Jónsson^{2,3}

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

²Department of Applied Physics, Aalto University, Finland

³Faculty of Physical Sciences, University of Iceland, 107 Reykjavík, Iceland

hj@hi.is

DOI 10.17586/2220-8054-2016-7-6-925-935

The calculation of minimum energy paths for transitions such as atomic and/or spin rearrangements is an important task in many contexts and can often be used to determine the mechanism and rate of transitions. An important challenge is to reduce the computational effort in such calculations, especially when *ab initio* or electron density functional calculations are used to evaluate the energy since they can require large computational effort. Gaussian process regression is used here to reduce significantly the number of energy evaluations needed to find minimum energy paths of atomic rearrangements. By using results of previous calculations to construct an approximate energy surface and then converge to the minimum energy path on that surface in each Gaussian process iteration, the number of energy evaluations is reduced significantly as compared with regular nudged elastic band calculations. For a test problem involving rearrangements of a heptamer island on a crystal surface, the number of energy evaluations is reduced to less than a fifth. The scaling of the computational effort with the number of degrees of freedom as well as various possible further improvements to this approach are discussed.

Keywords: minimum energy path, machine learning, Gaussian process, transition mechanism, saddle point.

Received: 2 December 2016

1. Introduction

The task of predicting the rate and identifying the mechanism of transitions involving some rearrangements of atoms in or on the surface of solids shows up in many different applications, for example diffusion, crystal growth, chemical catalysis, nanotechnology, etc. At a finite temperature, the thermal fluctuations in the dynamics of atoms can lead to rearrangements from one stable configuration to another, but these are rare events on the time scale of atomic vibrations, so direct dynamics simulations cannot in most cases be used for these types of studies. The separation of time scales typically amounts to several orders of magnitude and a direct simulation would take impossibly long time. Instead, algorithms based on statistical mechanics as well as classical dynamics and focusing on the relevant rare events need to be applied [1–3]. Typical transitions involve not just one or a few atoms but rather a large number of atoms so the challenge is also to deal with multiple degrees of freedom. One way of looking at the problem is to characterise the motion of the system on a high dimensional energy surface where the number of degrees of freedom is easily more than a hundred. A key concept is the reaction coordinate which usually is taken to be a minimum energy path (MEP) on the energy surface connecting one minimum to another. The rate of transitions in solids is usually evaluated within harmonic transition state theory which is based on a quadratic expansion of the energy surface at the initial state minimum and at the highest maximum along the MEP, which is a first order saddle point on the energy surface [4]. For given initial and final states, the task is to determine the MEP and identify the saddle point(s) as well as possible unknown, intermediate minima [5]. The discussion here has been in terms of rearrangements of atoms, but similar considerations apply to reorientations of magnetic moments [6–9].

The nudged elastic band (NEB) method is commonly used to find MEPs for atomic rearrangements [5, 10, 11]. An analogous method, referred to as the geodesic NEB, has been developed for magnetic transitions [12]. In NEB calculations, some initial path is constructed between two local minima on the energy surface and the path is represented by a discrete set of replicas of the system. The replicas are referred to as images of the system. They consist of some set of values for all degrees of freedom in the system. The NEB algorithm then optimises iteratively the location of the images that are between the endpoint minima so as to obtain a discrete representation of the MEP. Initially, the method was mainly used in combination with analytical potential energy functions, but today the method is used extensively in combination with electronic structure calculations. A large amount of computer time is used in these calculations. Each calculation typically involves 100 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 5 to 10 images. Since a typical electronic structure calculation takes on the order of tens of CPU minutes or more,

these calculations can be heavy. Also, several different possible final states usually need to be tested and the NEB calculation therefore repeated. In light of the widespread use and large amount of CPU time used in NEB calculations, it is of great practical importance to find ways to accelerate the calculations. The goal should be to use the information coming from all the computationally intensive electronic structure calculations in an optimal way so as to reduce as much as possible the number of iterations needed to reach the MEP.

It has recently been shown that a machine learning algorithm based on neural networks can be used to significantly reduce the computational effort in NEB calculations [13]. An approximate representation of the energy surface is constructed from the calculations using a machine learning approach and the MEP calculated using the NEB method on this approximate surface. Then, additional evaluations are made of the true energy surface, the approximate model surface refined, etc., until convergence on the MEP of the true energy surface has been reached. The number of function evaluations was shown to drop dramatically by applying such an approach [13].

We present here an initial step in the development of a similar approach to accelerated MEP calculations based on Gaussian process regression [14–17]. This approach could have some advantages over neural networks for such applications. Neural networks have a large number of weights which can have multimodal distributions making the search for global optimum difficult and leading to possible dependence on the initial values of the parameters [13]. Also, the handling of uncertainties in GP theory is easier than in neural networks since the prediction equations are analytical and integration over the parameter space can be carried out more easily. It is, therefore, of interest to test the efficiency of the GP approach in MEP calculations. We report in this article initial feasibility studies. More extensive testing and comparison with other approaches such as neural networks is left for future work.

The article is organized in the following way: The methodology is presented in the next section, followed by a section on applications, both a simple two-dimensional system and a larger test problem involving rearrangements of a heptamer island on a crystal surface. The article concludes with a discussion section.

2. Methods

The method presented here for finding the minimum energy paths can be viewed as an acceleration of a NEB calculation by making use of Gaussian process theory. Previously calculated data points are used to construct an approximate model of the energy surface and the MEP is found for this approximate surface before additional calculations of the true energy are carried out. This gives an interpolation between the calculated points and also provides an extrapolation that can be used to explore the energy surface with larger moves. The savings in computational effort are based on the fact that several computationally light iterations can be made for the approximate surface in between the computationally demanding evaluations of the true energy function. A brief review of the NEB method is first given, then a description of the Gaussian process regression, and finally a detailed algorithm describing how the calculations were carried out in the present case.

2.1. Nudged elastic band method

Given two local minima on the energy surface, the task is to find an MEP connecting the two. The definition of an MEP is that the gradient has zero component perpendicular to the path tangent at each point along the path. The NEB method needs to be started with some initial path between the two minima that is represented by a set of images. Most often, a straight line interpolation between the minima is used to generate the initial path [11], but a better approach is to start with a path that interpolates as closely as possible the changes distances between atoms [18].

The key aspect of the NEB algorithm is the nudging, a force projection which is used to decouple the displacements of the images perpendicular to the path towards the MEP from the displacements that affect their distribution along the path. In order to make this projection, an estimate of the local tangent to the path at each of the images is needed. A numerically stable choice involves finding the line segment from the current image to the adjacent image of higher energy [19].

Given this decoupling, there are several different options for distributing the images along the path. Some constraint is needed to prevent the images from sliding down to the minima at the two ends. In most cases an even distribution is chosen, but one can also choose to have, for example, higher density of images where the energy is larger [20]. An attractive spring force is typically introduced between adjacent images to control the spacing between images and this also prevents the path from becoming arbitrarily long in regions of little or no force. The latter is important, for example, in calculations of adsorption and desorption of molecules at surfaces. For systems that can freely translate and rotate, such as nano-clusters in free space, it is important to remove the translational and rotational degrees of freedom. This is non-trivial because the system cannot be treated as a rigid body. A method for doing this efficiently based on quaternions has recently been presented [21].

The component of the force acting on each image perpendicular to the path is used to iteratively move the images from the initial path to the MEP. The force is the negative of the gradient and in most cases an evaluation of the energy delivers also the gradient vector at little or no extra expense. The largest amount of information from an evaluation of a point on the energy surface is, therefore, represented by the gradient. It is however typically too expensive to evaluate second derivatives of the energy and iterative algorithms for moving the images towards the MEP are therefore based solely on the gradient and the energy at each point. A simple and numerically stable method that has been used extensively in NEB calculations will be used here. It is based on the velocity Verlet method where only the component of the velocity in the direction of the force is included and the velocity is zeroed if its dot product with the force becomes negative [11]. A somewhat higher efficiency can be obtained by using a quadratically convergent algorithm such as conjugate gradients or quasi-Newton [22] but those can be less stable especially in the beginning of an NEB calculation. A linear interpolation between the initial state minima was used in all the calculations presented here and the number of images, N_p , chosen to be either 5 or 8. An equal distribution of the images along the path was chosen.

The focus here is on calculations where the energy and the gradient are obtained using some ab initio or density functional theory calculation. The computational effort in all other parts of the calculation is then insignificant in comparison and the computational effort is well characterised by simply the number of times the energy and force need to be evaluated in order to converge on the MEP. Below, we introduce a strategy to accelerate the MEP search with Gaussian process regression.

2.2. Gaussian processes regression

The general idea behind the strategy is similar to the one introduced by Peterson [13]. The idea is to use the calculations carried out so far to train an approximate model of the energy surface, and find the MEP with the conventional methods using the approximations of the energy and gradient based on this model. After converging to the MEP on the approximate energy surface, the true energy and force are evaluated again, showing whether or not the path has converged to the true MEP. If not, the model is updated with the new values of the true energy and force to get a more accurate approximation, and this is continued iteratively, until the true MEP has been found. Since the number of true energy and force evaluations is the measure of computational effort, basically any method can be used to optimise the path on the approximate energy surface, as long as it converges to an MEP.

Here, a Gaussian process (GP) is used as a probabilistic model for the energy surface. GPs provide a flexible framework for modelling multidimensional functions. Through the selection of the covariance function and its hyperparameters, smoothness properties of the function can easily be defined and those properties can also be learned from the data. It is also straightforward to both include derivative observations into the model and to predict the derivative of the modelled function. Analytical expressions for the posterior predictions conditional on the hyperparameters allow both fast predictions and reliable estimation of uncertainties. In cases where only a small number of observations are available, Gaussian processes have been shown to have good predictive performance compared to other machine learning methods [23].

A GP can be seen as a probability distribution over functions in a continuous domain, see, e.g., [14–17]. In a GP, the joint probability distribution of the function values $f(\mathbf{x}^{(1)}), f(\mathbf{x}^{(2)}), \dots, f(\mathbf{x}^{(N)})$ at any finite set of input points $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^D$ is a multivariate Gaussian distribution. A GP is defined by a mean function $m(\mathbf{x})$ and a covariance function $k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$, which determines the covariance between $f(\mathbf{x}^{(i)})$ and $f(\mathbf{x}^{(j)})$, e.g., based on the distance between $\mathbf{x}^{(i)}$ and $\mathbf{x}^{(j)}$.

Consider a regression problem $y = f(\mathbf{x}) + \epsilon$, where ϵ is Gaussian noise with variance σ^2 , and a training data set $\{\mathbf{X}, \mathbf{y}\}$, where $\mathbf{X} \in \mathbb{R}^{N \times D}$ denotes a matrix of N input vectors $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^D$ and \mathbf{y} is a vector of the corresponding N noisy observations. By choosing a Gaussian process to model function f , different prior assumptions can be made about the properties of the function, and after observing $\{\mathbf{X}, \mathbf{y}\}$, the posterior predictive probabilities for the function values at a set of new points can be calculated analytically as a multivariate Gaussian distribution. Here, the mean function is taken to be $m(\mathbf{x}) = 0$ and the covariance function is assumed to have the form

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

where η^2 and $\boldsymbol{\rho} = \{\rho_1, \dots, \rho_D\}$ are the hyperparameters of the GP model. The squared exponential covariance function is infinitely differentiable and thus favours smooth functions. The length scales $\boldsymbol{\rho}$ define how fast the function f can change, and η^2 controls the magnitude of the overall variation. The additional constant term c^2 has a similar effect as integration over an unknown constant mean function with a Gaussian prior distribution of variance of c^2 . The posterior predictive distribution for a function value of the function at a new point \mathbf{x}^* , denoted

as f^* , is described by a Gaussian distribution with mean

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

and variance

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

where \mathbf{I} is the identity matrix and the notation $K(\mathbf{X}, \mathbf{X}')$ represents a covariance matrix with entries $K_{ij} = k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$. The hyperparameter values $\boldsymbol{\theta} = \{\eta^2, \boldsymbol{\rho}\}$ are optimised by defining a prior probability distribution $p(\boldsymbol{\theta})$ and maximising the marginal posterior probability $p(\boldsymbol{\theta}|\mathbf{y}, \mathbf{X}) = p(\boldsymbol{\theta})p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta})$ after observing \mathbf{y} .

Since differentiation is a linear operation, the derivative of a Gaussian process is also a Gaussian process (see, e.g., [24, 25]), and this makes it possible to use observations of the derivative of the function and also to predict derivatives of the function f . The partial derivative observations can simply be included in the observation vector \mathbf{y} and the covariance matrix correspondingly extended with the covariances between the observations and the partial derivatives and the covariances between the partial derivatives themselves. In the case of the squared exponential covariance function, these entries are obtained by

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

and

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

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 useful also when predicting values of 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

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

and variance

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

In the present application, the vector \mathbf{x} represents coordinates of the atoms and the function f the energy of the system. The observations \mathbf{y} are the true values of the energy as well as the partial derivatives of the energy with respect to the coordinates of the atoms at the various sets of coordinates $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}$. With this input, the Gaussian process model is used to predict the most likely value of energy f^* and its derivatives $\frac{\partial f^*}{\partial x_d^*}$ at a new set of atom coordinates \mathbf{x}^* representing in this case an image in the discrete path representation between the initial and final state minima. Since the training data is assumed to be noiseless and include also derivative observations, the equations for the mean predictions can be presented as

$$E[f^*|\mathbf{x}^*, \mathbf{y}_{ext}, \mathbf{X}, \boldsymbol{\theta}] = K_{ext}^* K_{ext}^{-1} \mathbf{y}_{ext} \quad (1)$$

and

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

where

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

$$K_{ext}^* = \left[K(\mathbf{x}^*, \mathbf{X}) \quad \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_1} \quad \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_2} \quad \dots \quad \frac{\partial K(\mathbf{x}^*, \mathbf{X})}{\partial x_D} \right],$$

and

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

2.3. Algorithm for GP-aided MEP search

Input: the coordinates, energy and its gradient at the two minima on the energy surface, the number of images representing the path (N_p), convergence limit (CL), step coefficient (k_{step}).

Output: minimum energy path represented by N_p images.

1. Place the initial N_p images equally spaced along a straight line between the two minima.
2. Repeat until convergence (outer iteration loop):
 - A. Evaluate the true energy and its gradient at the $N_p - 2$ intermediate images of the path, and add them to the training data.
 - B. Calculate the negative energy gradient (e.g., force) component perpendicular to the path (ngc) for each intermediate image, and denote the mean of their norms as M_{ngc} .
 - C. If $M_{ngc} < CL$, the path has converged to the true MEP.
 - D. Optimise the hyperparameters of the GP model based on the training data, and calculate the matrix inversion in equation 1.
 - E. Define CL_{relax} as $\frac{1}{10}$ of the smallest M_{ngc} so far, and repeat (relaxation phase):
 - I. Move the intermediate images according to any stable path optimisation algorithm.
 - II. Update the GP posterior mean energy and gradient at the new intermediate images using equations 1 and 2.
 - III. Calculate ngc for each image using the GP posterior mean gradient, and denote the mean of their norms as M_{ngc}^{GP} .
 - IV. If $M_{ngc}^{GP} < CL_{relax}$, or if M_{ngc}^{GP} is increasing, exit the relaxation phase (E).

The GP calculations make use of the GPstuff toolbox [26]. For the hyperparameter optimisation which is carried out after each evaluation of the true energy and force, the computational effort scales as $\mathcal{O}((N(D+1))^3)$, where N is the number of observations and D is the number of degrees of freedom (here coordinate of movable atoms). Since the hyperparameters and observations stay the same during a search for the MEP on the approximate energy surface, the matrix inversion in equation 1 needs to be computed only once for each such optimization of the path. Thus, the complexity of one inner iteration on the GP posterior energy surface is $\mathcal{O}(N(D+1))$.

The length of any one displacement of an image is restricted to be less than half of the initial interval between the images in order to prevent the path from forming loops. Convergence of the path to the MEP is determined from the norm of the force component perpendicular to the path at each of the intermediate images. The path is considered to be converged to the MEP, when the mean of the true values of these norms is less than 0.001 eV/Å. During the relaxations, norms based on the current GP model are monitored and the mean of these used as a convergence criterion. Since it is not necessary to find a path that is accurately converged on the MEP of the inaccurate, approximate energy surface, the convergence limit for each relaxation phase is defined as $\frac{1}{10}$ of the smallest true mean of norms evaluated so far. Higher convergence limits at early relaxation steps speed up the algorithm and they also make it more stable by preventing the path from escaping too far from the true observation points. For the same reason, the relaxation is stopped before convergence if the convergence criterion starts to increase.

3. Applications

The method described above has been applied to two test problems: A simple two-dimensional problem where the energy surface can be visualised, and a more realistic problem involving the rearrangements of atoms in a heptamer island on a crystal surface.

3.1. Two-dimensional test problem

The two-dimensional problem is formulated by coupling a degree of freedom representing the simultaneous formation and breaking of chemical bonds with a degree of freedom representing a harmonic oscillator solvent environment. The model along with the detailed equations is described in the appendix A.2 of reference [11]. Here, one additional repulsive Gaussian was added to shift the saddle point away from the straight line interpolation between the two minima. A contour graph of the energy surface is shown in Fig. 1.

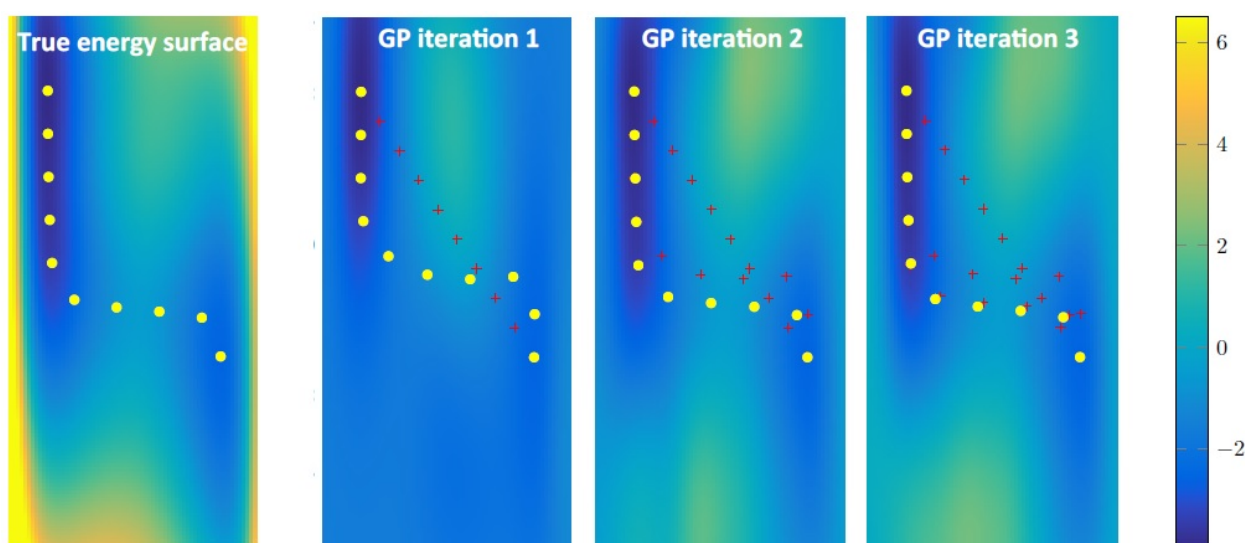


FIG. 1. The true and Gaussian process approximated energy surface and minimum energy path for a two-dimensional test problem. Far left: The true energy surface and points on the minimum energy path (yellow dots). Far right and intermediate figures: The approximate energy surface generated by the Gaussian process regression after one, two and three iterations, points ('images') on the estimated minimum energy path and points where the true energy and force have been calculated (red + signs) at each stage of the calculation the period is missing.

This example shows how the GP model of the energy surface is gradually built up and refined as more observations, i.e. calculations of the true energy and partial derivatives of the energy, are made. Here, $N_p = 10$ images are used to represent the path and the calculation is started by placing the images along a straight line between the two minima on the energy surface. The first observations are made at those points (see red + signs on the figure second from the left). Based on the energy and partial derivatives of the energy at those points, the GP model already shows some of the most important features of the energy surface close to the linear interpolation, but completely misses the increase in energy in the lower half of the figure. The relaxation of the images on this rough estimate of the energy surface does not, however, bring the images too far from the initial placement because of the condition that images cannot be moved in a single iteration by more than half the initial distance between the images. In the second GP iteration, observations are made at the position of the images at the end of the first GP iteration. When those data points are fed into the GP model, the energy surface is already showing the essential features around the MEP, but of course misses the steep increase in the energy far from the MEP. The relaxation of the images during the second GP iteration brings them quite close to the MEP. The addition of observations at those points at the beginning of the third GP iteration refines the model energy surface further. While a total of six GP iteration are required to bring the images onto the MEP to within the tight tolerance of 0.001 eV/\AA in the mean magnitude of the force component perpendicular to the path, no visible changes occur in the contour graph or the location of the images, so the results are not displayed in the figure.

3.2. Heptamer island on a crystal surface

A more realistic test problem which has been used in several studies of MEP and saddle point searches involves an island of 7 atoms on the (111) surface of a face centered cubic (FCC) crystal (see, for example, references [27,28]). Roughly, this represents a metallic system, but the interaction between the atoms is described here with a simple Morse potential to make it easier to implement the benchmark calculation. The initial, saddle point and final configurations of the atoms for three possible rearrangements of the atoms is shown in Fig. 2. Several other transitions are possible (see reference [27]), but these three are chosen as examples.




FIG. 2. On-top view of the surface and the seven atom island used to test the efficiency of the Gaussian process regression method. The initial state is shown to the left. The saddle point configurations and the final state configurations of three example transitions are also shown. Transition 1 corresponds to a pair of edge atoms sliding to adjacent FCC sites. In transition 2, an atom half way dissociates from the island. In transition 3, 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. At the same time the other island atoms as well as some of the underlying atoms also move but in the end return to nearly the same position as they had initially.

The three examples chosen here represent three types of transitions that can occur in the shape of the island. In one case, a pair of edge atoms slides to adjacent FCC sites, in another an atom half way dissociates from the island, and in the third case 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 energy along the MEP for the transition 3 is shown in Fig. 3 as well as the energy of the $N_p = 7$ images at the end of GP iterations 1 to 7. After the first and second GP iteration, the estimates of the MEP is quite inaccurate and the energy rises along those paths by more than 3 eV, but already after the third GP iteration, the estimated energy barrier is not too far from the accurate value. After the fifth GP iteration, the shape of the energy curve is quite well reproduced, and after seven iterations the energy along the MEP of the approximate energy surface is nearly indistinguishable from the energy along the true MEP.




FIG. 3. Energy along paths for transition 3 shown in Fig. 2. The energy of images on the true MEP are shown in blue, but the energy of images on MEPs of approximate models of the energy surface obtained after 1 to 7 Gaussian process iterations are shown in red. After the first two Gaussian process iterations, the energy barrier for this transition is greatly overestimated, but already after three iterations the estimated energy barrier is quite close to the true value, and after 7 iterations an accurate estimate is obtained from the model energy surface.

The number of energy and force evaluations needed to converge the five intermediate images to the MEP in both a regular NEB calculation and in a GP aided calculation was found for varying number of degrees of freedom. The average for the three transitions depicted in Fig. 2 is shown in Fig. 4. The number of degrees of freedom varies from 21 (as only the island atoms are allowed to move while all the substrate atoms are kept immobile), to 42 (as seven of the closest substrate atoms are also allowed to move during the transition). The number of energy and force evaluations for the NEB method obtained here is similar to what has been reported earlier for this test problem, see references [27,28]. It is possible to use a more efficient minimisation scheme to relax the images in NEB calculations [22], but the difference is not large.

A large reduction in the number of energy and force evaluations is obtained by using the GP regression, as shown in Fig. 4. With the GP regression, the reduction is to less than a fifth as compared with the regular NEB calculation. In calculations involving *ab initio* or density functional theory evaluation of the energy and force, the computational effort is essentially proportional to this number of observations and the additional calculations involved in the GP regression is insignificant in comparison. This test problem, therefore, shows that the use of GP regression can significantly reduce the computational effort in, for example, calculations of surface processes.

4. Discussion

The results presented in this article indicate that GP regression is a powerful approach for significantly reducing the computational effort in calculations of MEPs for transitions. This is important since a great deal of computer time is used in such calculations, especially when *ab initio* or density functional theory calculations are used to evaluate the energy and atomic forces. The heptamer island test problem studied here indicates that the

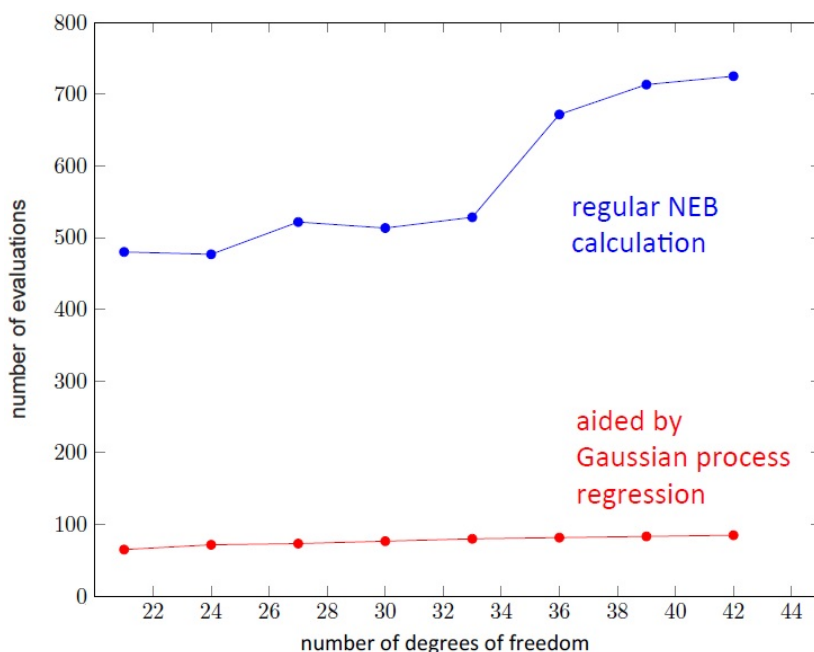


FIG. 4. The average number of energy and force evaluations needed to converge five intermediate images on the minimum energy paths of the three heptamer island transitions shown in Fig. 2 as a function of the number of degrees of freedom included in the calculations. The convergence tolerance is 0.001 eV/\AA for the magnitude of the perpendicular component of the force on any one of the images. For the smallest number, 21, only the seven island atoms are allowed to move and all substrate atoms are immobile. For a larger number of degrees of freedom, some of the substrate atoms are also allowed to move during the transition. In the regular NEB calculations (blue dots), the minimization method for relaxing the images to the MEP is based on velocity Verlet algorithm, as described in reference [11]. In the Gaussian process regression calculations (red dots), the number of true energy and function evaluations is less than a fifth of what is needed in the regular NEB calculation. This illustrates well the large reduction in the computational effort that Gaussian process regression can provide in a typical surface process calculation.

computational effort can be reduced to less than a fifth. But, this study represents only an initial proof-of-principle demonstration of the GP regression in this context. There are several ways in which the implementation can be improved and made more efficient. One of the advantages of GP regression over, for example, neural networks is the availability of uncertainty estimates which can be used to make the observations more selective. In the present case, an observation (i.e., evaluation of the true energy and force) was made for all the images in each GP iteration. Alternatively, an observation may only be made for the image for which there is greatest uncertainty. This could target the calculations better and thereby reduce the total number of energy and force evaluations needed to converge to the true MEP.

While the whole path has to be converged well enough to provide an accurate estimate of the tangent, the part of the path that is most important for practical purposes is the region around the first order saddle point. In most cases, the MEP is needed mainly to find the highest energy point along the path, i.e. the first order saddle point on the energy surface that is required for evaluating the transition rate within harmonic transition state theory. The algorithm can be refined to take this into account by, for example, applying the climbing-image NEB [20] where one of the images is driven to the maximum energy along the path, and at the same time the tolerance for the convergence of other images can be increased.

In a typical case, the goal is to evaluate the transition rate using harmonic transition state theory. There, the second derivative matrix, the Hessian matrix, and the frequency of vibrational modes needs to be evaluated at the end points as well as at the (highest) first order saddle point. While the saddle point is not known until the MEP calculation has been carried out, the minima are, and the second derivative matrices at those points might as well be calculated right from the start. This would provide additional information that could be fed into the GP regression so as to improve the accuracy of the approximate energy surface right from first GP iteration. It

remains an interesting challenge to extend the GP regression approach to include in some way such information on the second derivatives.

The test problems studied here are quite simple, and it will be important to test the method on more complex systems to fully assess its utility and to develop it further. One issue that can arise is that more than one MEP connects the two endpoint minima. Then, some kind of sampling of MEPs needs to be carried out [29]. Also, some energy surfaces have multiple local minima and highly curved MEPs, which can lead to convergence problems unless a large number of images is included in the calculation. The scaling of the GP regression approach to such more challenging problems needs to be tested. There will, however, clearly be a large set of important problems, such as calculations of catalytic processes, which often involve rather small molecules adsorbed on surfaces, where the complexity is quite similar to the heptamer island test problem studied here, and where the GP regression is clearly going to offer a significant reduction in computational effort.

At low enough temperature, quantum mechanical tunneling becomes the dominant transition mechanism and the task is then to find the minimum action path [5,30,31]. Calculations of tunneling paths requires exploring the energy surface over a wider region than a calculation of MEPs and here again the GP regression approach can lead to a significant reduction in computational effort, even more than for MEP calculations since each iteration necessarily involves more observations and thereby more input for the modeling of the energy surface.

The discussion has focused here on atomic rearrangements, but it will, furthermore, be interesting to apply the GP regression approach to magnetic transitions where the evaluation of the magnetic properties of the system is carried out using computationally intensive *ab initio* or density functional theory calculations. There, the relevant degrees of freedom 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 [7–9].

Acknowledgements

HJ would like to thank Prof. Andrew Peterson at Brown University for helpful discussions. This work was supported by the Academy of Finland (FiDiPro program grant no. 263294) and by the Icelandic Research Fund.

References

- [1] Wigner E. The Transition State Method. *Trans. Faraday Soc.*, 1938, **34**, P. 29.
- [2] Kramers H.A. Brownian Motion in a Field of Force and the Diffusion Model of Chemical Reactions. *Physica*, 1940, **7**, P. 284.
- [3] Keck J.C. Variational theory of reaction rates. *J. Chem. Phys.*, 1967, **13**, P. 85.
- [4] Vineyard G.H. Frequency factors and isotope effects in solid state rate processes. *J. Phys. Chem. Solids*, 1957, **3**, P. 121.
- [5] Jónsson H. Simulation of Surface Processes. *Proceedings of the National Academy of Sciences*, 2011, **108**, P. 944.
- [6] Bessarab P.F., Uzdin V.M., Jónsson H. Harmonic transition state theory of thermal spin transitions. *Phys. Rev. B*, 2012, **85**, P. 184409.
- [7] Bessarab P.F., Uzdin V.M., Jónsson H. Potential Energy Surfaces and Rates of Spin Transitions. *Z. Phys. Chem.*, 2013, **227**, P. 1543.
- [8] Bessarab P.F., Uzdin V.M., Jónsson H. Calculations of magnetic states and minimum energy paths of transitions using a noncollinear extension of the Alexander-Anderson model and a magnetic force theorem. *Phys. Rev. B*, 2014, **89**, P. 214424.
- [9] Bessarab P.F., Skorodumov A., Uzdin V.M., Jónsson H. Navigation on the energy surface of the noncollinear Alexander-Anderson model. *Nanosystems: Physics, Chemistry, Mathematics*, 2014, **5**, P. 757.
- [10] Mills G., Jónsson H., Schenter G.K., Reversible work based transition state theory: Application to H₂ dissociative adsorption. *Surface Science*, 1995, **324**, P. 305.
- [11] 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”, edited by B.J. Berne, G. Ciccotti, D.F. Coker, pages 385-404. World Scientific, 1998.
- [12] Bessarab P.F., Uzdin V.M., Jónsson H. Method for finding mechanism and activation energy of magnetic transitions, applied to skyrmion and antivortex annihilation. *Comput. Phys. Commun.*, 2015, **196**, P. 335.
- [13] Peterson A.A. Acceleration of saddle-point searches with machine learning. *J. Chem. Phys.*, 2016, **145**, P. 074106.
- [14] O’Hagan A. Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society (Series B)*, 1978, **40**, P. 1.
- [15] MacKay D.J.C. Introduction to Gaussian processes. In “Neural Networks and Machine Learning”, Editor C.M. Bishop, pages 133-166. Springer Verlag, 1998.
- [16] Neal R.M. Regression and classification using Gaussian process priors (with discussion). In “Bayesian Statistics 6”, Editors J.M. Bernardo, J.O. Berger, A.P. Dawid, A.F.M. Smith, pages 475-501 (Oxford University Press, 1999).
- [17] Rasmussen C.E., Williams C.K.I. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [18] Smidstrup S., Pedersen A., Stokbro K., Jónsson H. Improved initial guess for minimum energy path calculations. *J. Chem. Phys.*, 2014, **140**, P. 214106.
- [19] Henkelman G., Jónsson H. Improved Tangent Estimate in the NEB Method for Finding Minimum Energy Paths and Saddle Points. *J. Chem. Phys.*, 2000, **113**, P. 9978.
- [20] Henkelman G., Uberuaga B.P., Jónsson H. A Climbing-Image NEB Method for Finding Saddle Points and Minimum Energy Paths. *J. Chem. Phys.*, 2000, **113**, P. 9901.
- [21] Melander M., Laasonen K., Jónsson H. Removing external degrees of freedom from transition state search methods using quaternions. *Journal of Chemical Theory and Computation*, 2015, **11**, P. 1055.
- [22] Sheppard D., Terrell R., Henkelman G. Optimization methods for finding minimum energy paths. *J. Chem. Phys.*, 2008, **128**, P. 134106.

- [23] Lampinen J., Vehtari A. Bayesian approach for neural networks – review and case studies. *Neural Networks*, 2001, **14**, P. 7.
- [24] Solak E., Murray-Smith R., Leithead W.E., Leith D.J., Rasmussen C.E. Derivative observations in Gaussian process models of dynamic systems. In “Advances in Neural Information Processing Systems 15”, pages 1033–1040. MIT Press, 2003.
- [25] Rasmussen C.E. Gaussian processes to speed up Hybrid Monte Carlo for expensive Bayesian integrals. In “Bayesian Statistics” 7, pages 651–659. Oxford University Press, 2003.
- [26] Vanhatalo J., Riihimäki J., Hartikainen J., Jylänki P., Tolvanen V., Vehtari A. GPstuff: Bayesian Modeling with Gaussian Processes. *Journal of Machine Learning Research*, 2013, **14**, P. 1175.
- [27] Henkelman G., Jóhannesson G.H., Jónsson H. Methods for finding saddle points and minimum energy paths. In “Progress in Theoretical Chemistry and Physics”, ed. S. D. Schwartz, Vol. 5, chapter 10, pages 269–300. Kluwer Academic, Dordrecht, 2000.
- [28] Chill S.T., Stevenson J., Ruhle V., Shang C., Xiao P., Farrell J., Wales D., Henkelman G. Benchmarks for characterization of minima, transition states and pathways in atomic, molecular, and condensed matter systems. *J. Chem. Theory Comput.*, 2014, **10**, P. 5476.
- [29] Maras E., Trushin O., Stukowski A., Ala-Nissila T., Jónsson H. Global transition path search for dislocation formation in Ge on Si(001). *Comput. Phys. Commun.*, 2016, **205**, P. 13.
- [30] Mills G., Schenter G.K., Makarov D., Jónsson H. Generalized Path Integral Based Quantum Transition State Theory. *Chemical Physics Letters*, 1997, **278**, P. 91.
- [31] Mills G., Schenter G.K., Makarov D. Jónsson H. RAW Quantum Transition State Theory. In “Classical and Quantum Dynamics in Condensed Phase Simulations”, editors B.J. Berne, G. Ciccotti and D.F. Coker, page 405. World Scientific, 1998.