

Computational Physics

Fast automatically differentiable matrix functions and applications in molecular simulations

Tina Torabi^{a,*}, Timon S. Gutleb^{b,1}, Christoph Ortner^a^a Department of Mathematics, University of British Columbia, Vancouver, V6T1Z2, BC, Canada^b School of Computer Science, University of Leeds, Leeds, LS2 9JT, UK

ARTICLE INFO

The review of this paper was arranged by Prof. Blum Volker

Dataset link: [EntropyGrad.jl \(Original data\)](#),
[Animations of migration paths for defects in silicon crystals at 100–K \(Original data\)](#)

ABSTRACT

We describe efficient differentiation methods for computing Jacobians and gradients of a large class of matrix functions including the matrix logarithm $\log(A)$ and p -th roots $A^{\frac{1}{p}}$. We exploit contour integrals and conformal maps as described by Hale et al. (2008) [34] for evaluation and differentiation and analyze the computational complexity as well as numerical accuracy compared to high accuracy finite difference methods. As a demonstrator application we compute properties of structural defects in silicon crystals at positive temperatures, requiring efficient and accurate gradients of matrix trace-logarithms.

1. Introduction

Matrix functions have long been integral to mathematical and scientific disciplines and thus have a wide range of applications [49,37,28] in fields such as network analysis [6], control systems [23], matrix function neural networks [4], and solid state physics [14,9,67]. For instance, in materials science, matrix functions are indispensable for calculating vibrational entropy, which provides insights into the impact of thermal vibrations on material properties [67,9].

Despite their broad applicability, the computation of derivatives of matrix functions—in particular for matrix families X_θ parameterized by a vector θ —remains a significant computational challenge. This challenge is particularly acute in high-dimensional systems, where conventional methods often become computationally prohibitive. For instance, in the study of thermally activated processes [12] precise calculations of free energy derivatives are required. Conventional approaches often rely on potential energy surface (PES) approximations, which, while computationally feasible, can introduce significant errors when entropy plays a major role. To accurately capture the full behavior of such systems, it is necessary to consider the free energy surface (FES) instead, which requires the computation of more complex matrix function derivatives.

In this paper we introduce a method that leverages analytic functional calculus in conjunction with both forward and reverse mode automatic differentiation (AD). Our objective is to develop a robust framework that significantly enhances the efficiency and scalability of com-

puting derivatives of matrix functions. Our approach streamlines computational implementation and expands the potential applications of matrix function derivatives to more complex higher-dimensional problems.

2. Matrix functions

2.1. Analytic matrix functions

We use the term *matrix function* to refer to a matrix-valued extension of a scalar function f such that $f(X)$ with $X \in \mathbb{C}^{n \times n}$ is a computable matrix retaining the dimensions of X . If $f(x) = a_0 + a_1x + \dots + a_Kx^K$ is a polynomial, this extension is given by

$$f(X) = \sum_{k=0}^K a_k X^k,$$

with $X^0 = I_n$, the $n \times n$ identity matrix. For general continuous f , $f(X)$ is defined via density of polynomials. Several approaches for evaluating matrix functions have been considered in the literature [37], e.g. via the eigendecomposition (or more generally the Jordan canonical form), or via a specific polynomial interpolation scheme. In this work we are primarily concerned with evaluating analytic matrix functions.

Suppose that $f : \mathbb{C} \rightarrow \mathbb{C}$ is analytic on a closed set $\overline{D} \subset \mathbb{C}$, where D is an open set containing the spectrum of a matrix $X \in \mathbb{C}^{n \times n}$; i.e.,

* Corresponding author.

E-mail addresses: torabit@student.ubc.ca (T. Torabi), T.S.Gutleb@leeds.ac.uk (T.S. Gutleb), ortner@math.ubc.ca (C. Ortner).¹ These authors contributed equally to this work.




Fig. 1. Figure (a) depicts the spectrum of X along with the appropriate contour C . Figure (b) demonstrates mapping the entire doubly connected domain of analyticity Ξ , onto an annulus region X through a conformal map. In this figure, the interval of singularities corresponds to the outer boundary of the annulus, while the interval including the spectrum aligns with the dashed inner boundary circle. We then apply the trapezoid rule over a circle in the annulus.

$\sigma(X) \subset D$. Assume further that the boundary, $C := \partial D$, is comprised of a finite number of closed rectifiable Jordan curves, encircling $\sigma(X)$ once in the counterclockwise direction. Then $f(X)$ can be expressed as a Riemann contour integral over C as follows:

$$f(X) = \frac{1}{2\pi i} \oint_C f(z)(zI - X)^{-1} dz. \quad (1)$$

2.2. Fast computation of matrix functions $f(X)$

In physics and engineering applications, performant numerical evaluation of matrix functions is often essential. Hale et al. [34] demonstrated that the Cauchy integral definition in (1) is particularly useful for numerical computations. They described an efficient approach utilizing (1) for matrices A whose eigenvalues are located on or near the positive real axis $(0, \infty)$, and for functions $f(z)$ such as z^α or $\log z$, which are analytic except for the presence of singularities or a branch cut near the negative real axis $(-\infty, 0]$. This section reviews their method which we use in subsequent sections to also compute derivatives of matrix functions.

Consider a matrix X with real entries and eigenvalues in $(0, \infty)$ as illustrated in Fig. 1a. While X is symmetric in many important cases, facilitating the application of trapezoidal sums, symmetry is not essential for the methods presented and will thus not be assumed. We will refer to the minimum and maximum eigenvalues of X as m and M , respectively. In cases where X is either symmetric or normal, the 2-norm condition number is given by the ratio M/m . We will generally assume that m and M are known though we note that they would typically only be estimated in practice. We further assume that the spectrum $\sigma(X)$ completely fills the interval $[m, M]$, without significant gaps and that the function f is analytic in the slit complex plane $\mathbb{C} \setminus (-\infty, 0]$. At first glance, one might think of surrounding $[m, M]$ by an appropriate contour, such as depicted in Fig. 1a, followed by applying the trapezoid rule to approximate the integral. However, for ill-conditioned matrices this is inefficient as it necessitates $O(M/m)$ linear solves to achieve typical accuracy requirements. The techniques suggested in [34] leverage variable transformations and conformal mappings to optimize the choice of contour points, effectively reducing the computational complexity to $O(\log(M/m))$.

The conformal mapping employed involves multiple transformations designed to map the region of analyticity of f and $(zI - X)^{-1}$, characterized by the doubly connected set $\Xi = \mathbb{C} \setminus ((-\infty, 0] \cup [\frac{m}{2}, \frac{M}{2}])$, onto an annulus $A = \{z \in \mathbb{C} : r < |z| < R\}$, where r and R represent the inner and outer radii of the annulus respectively. We show a schematic of this transformation in Fig. 1b. The question of how to map Ξ to the annulus A and vice versa was answered by Hale et al. [34] in three steps: Starting in the s -plane, we first map the annulus to a rectangle with vertices $\pm K$ and $\pm K + iK'$, using a logarithmic transformation

$$t(s) = \frac{2Ki}{\pi} \log\left(-\frac{is}{r}\right), \quad (2)$$

where K, K' denote the complete elliptic integrals. For more details on Jacobi elliptic functions and integrals, see Appendix A.1. Next, the rectangle is mapped to the upper half-plane in the u -plane by the Jacobian elliptic function

$$u(t) = \text{sn}(t|k^2), \quad k = \frac{(M/m)^{1/2} - 1}{(M/m)^{1/2} + 1}. \quad (3)$$

A Möbius transformation is then applied to map the upper half-plane to the z -plane

$$z(u) = (M/m)^{1/2} \left(\frac{k^{-1} + u}{k^{-1} - u} \right). \quad (4)$$

This final transformation is designed to distribute the eigenvalues of X evenly along the real axis, thus facilitating the application of the trapezoidal rule [34]. The steps of this conformal map are shown in Fig. 2. The integral in (1) can be reformulated as follows:

$$f(X) = -\frac{X}{2\pi i} \int_{-K+iK'/2}^{3K+iK'/2} f(z(t))(z(t) - X)^{-1} z^{-1} \frac{dz}{du} \frac{du}{dt} dt, \quad (5)$$

where the interval from $-K + iK'/2$ to $K + iK'/2$ reflects the segment of Γ in the upper half-plane, extended to $3K + iK'/2$ to include the lower half-plane contribution, with z transformations provided by:

$$\frac{dz}{du} = \frac{2k^{-1}\sqrt{mM}}{(k^{-1} - u)^2}, \quad \frac{du}{dt} = \text{sn}(t) = \sqrt{1 - k^2 u^2} = \text{cn}(t) \text{dn}(t). \quad (6)$$

Here, cn and dn are standard Jacobi elliptic functions [20, (22.2.4–22.2.6)]. Therefore, the expression for $f(X)$ becomes:

$$f(X) = -\frac{X\sqrt{mM}}{\pi k} \int_{-K+iK'/2}^{3K+iK'/2} \frac{f(z(t))(z(t) - X)^{-1} z^{-1} \text{cn}(t) \text{dn}(t)}{(k^{-1} - u)^2} dt. \quad (7)$$

Given that X is real, the integrand is real-symmetric, implying that $f(X)$ is effectively double the real part of the integral evaluated over the first half of the contour. Simplifying, we have:

$$f(X) = -\frac{2X\sqrt{mM}}{\pi k} \text{Im} \int_{-K+iK'/2}^{K+iK'/2} \frac{f(z(t))(z(t) - X)^{-1} z^{-1} \text{cn}(t) \text{dn}(t)}{(k^{-1} - u)^2} dt. \quad (8)$$

Applying the trapezoidal rule, with $t_j = -K + \frac{iK'}{2} + 2\frac{(j-1/2)K}{N}$ for $j = 1, \dots, N$, representing N equidistant points in $-K + iK'/2, K + iK'/2$, results in

$$f_N(X) = -\frac{4KX\sqrt{mM}}{\pi Nk} \text{Im} \left(\sum_{j=1}^N f(z(t_j)) \frac{(z(t_j)I - X)^{-1} \text{cn}(t_j) \text{dn}(t_j)}{z(t_j)(k^{-1} - u(t_j))^2} \right). \quad (9)$$

The convergence rate of the numerical approximation of $f(X)$ using an N -point quadrature formula to the true value of $f(X)$ is given in [34, Theorem 1]. Specifically, the error between the actual function $f(X)$ and the numerical approximation $f_N(X)$ is bounded by:

$$\|f(X) - f_N(X)\| = O\left(e^{-\pi^2 N / \log(M/m+3)}\right). \quad (10)$$

It is noteworthy that when f has a singularity at $z = 0$ but just a branch cut on $(-\infty, 0)$, the above mentioned method is not as efficient as in the case when f has singularities on $(-\infty, 0)$. To reach fast convergence, Hale et al. [34] proposed a change of variable $w = \sqrt{z}$, $dz = 2w dw$, under which (1) becomes

$$f(X) = \frac{X}{\pi i} \int_{C_w} w^{-1} f(w^2)(w^2 - X)^{-1} dw. \quad (11)$$




Fig. 2. Steps of the conformal map shown in detail. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

In the revised approach, the branch cut of the function $f(z)$ along the negative real axis in the z -plane is unfolded to the imaginary axis in the w -plane, allowing $f(z)$ to be analytically continued as $f(w^2)$ across the entire slit w -plane, i.e., $\mathbb{C} \setminus (-\infty, 0]$. The method involves a contour integration within a modified region, specifically enclosing $[m^{1/2}, M^{1/2}]$ in the w -plane. This adaptation improves the domain from $[m, M]$ to $[m^{1/2}, M^{1/2}]$, enabling a more effective application of the previous computational techniques.

The expression used for this refined method is

$$f_N(X) = -\frac{8KX\sqrt[4]{mM}}{\pi Nk} \operatorname{Im} \left(\sum_{j=1}^N f(w(t_j)^2) \frac{(w(t_j)^2 I - X)^{-1} \operatorname{cn}(t_j) \operatorname{dn}(t_j)}{w(t_j)(k^{-1} - u(t_j))^2} \right), \quad (12)$$

where the t_j are specified by the modified equations, reflecting changes in the contour parameters:

$$w = \sqrt[4]{mM} \frac{k^{-1} + \operatorname{sn}(t)}{k^{-1} - \operatorname{sn}(t)}, \quad k = \frac{\sqrt[4]{M/m} - 1}{\sqrt[4]{M/m} + 1}.$$

These adjustments allow for precise integration over the new contour, efficiently capturing the integral's value while minimizing computational overhead.

2.3. Derivatives of matrix functions

The framework for evaluating matrix functions $f(X)$ reviewed in Section 2.2 can be efficiently extended to families of matrices denoted as $X(u)$, where $u \in \Omega \subset \mathbb{R}^m$ represents a multi-dimensional parameter, as well as derivatives of $f(X(u))$ with respect to the parameter $u \in \Omega$. Here, $\Omega \subset \mathbb{R}^m$ is an open parameter domain and $m \in \mathbb{N}$ the number of parameters.

Lemma 2.1. Consider a matrix $X(u)$, where each entry $x_{ij} : \Omega \rightarrow \mathbb{C}$ is continuously differentiable. Assume that f is analytic in a domain \bar{D} , where D is open and includes the spectrum of $X(u)$ for all $u \in \Omega$. Then:

The derivatives of the matrix function $f(X)$ with respect to the component u_k of u are given by:

$$\frac{\partial f(X)}{\partial u_k} = \frac{1}{2\pi i} \int_C f(z)(zI - X)^{-1} \left(\frac{\partial X}{\partial u_k} \right) (zI - X)^{-1} dz,$$

where C is a positively oriented contour enclosing the spectrum of X , and zI represents the identity matrix scaled by z .

The proof based on (1) is elementary and hence omitted.

Computing the derivatives $\frac{\partial f(X)}{\partial u_m}$ with respect to each component u_m of a multi-dimensional parameter vector u can lead to substantial computational complexity since it involves evaluating $\frac{\partial X}{\partial u_m}$, the Jacobian of X with respect to u . The resulting Jacobian $\frac{\partial X}{\partial u}$ is a tensor of dimensions $n \times n \times m$, representing a significant computational bottleneck. In the section which follows, we discuss the computational complexity of this approach in more detail and explain how reverse mode differentiation avoids this in many important scenarios.

Remark 2.1. (Effect of a clustered or ill-conditioned spectrum) The only spectrum-dependent factor in the error (10) is the logarithm of the 2-norm condition number M/m . Hence a poorly separated or ill-conditioned spectrum only slows down the convergence but it does not really effect stability or differentiability. In practice one can simply increase the number N of quadrature points until the desired tolerance is reached. When M/m is extremely large, a standard approach is to apply a spectrum shifting or scaling preconditioner before the contour evaluation [22,48].

Remark 2.2. (Generality of the contour-integral scheme) It is noteworthy that the contour integral framework above relies only on two key factors: (i) an analytic continuation of $f(z)$ in a slit or a domain that avoids the spectrum of X ; and (ii) the ability to evaluate the resolvent $(zI - X)^{-1}$ at the quadrature nodes. Therefore, any primary matrix function whose scalar version is analytic, e.g. the matrix exponential, matrix sign, powers, can be treated within exactly the same framework by choosing a suitable branch cut and, if necessary, applying the square-root substitutions done above or similar substitutions to regularize any singularities. The computational complexity associated with the method thus depends on the linear solver rather than on the particular analytic function, and the error bound (10) continues to apply in all these cases.

Remark 2.3. (Computational complexity) In practical settings with sparse or structured X , the Cauchy integral framework allows us to replace an $\mathcal{O}(n^3)$ dense eigendecomposition by ℓ sparse solves whose cost scales with the sparsity pattern, so the contour framework is both cheaper and easier to differentiate. The preceding Section 3 gives the detailed quantitative comparison.

3. Computational complexity

3.1. Objective and formulation

Let $g : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^p$, $f : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$, and $X : \mathbb{R}^m \rightarrow \mathbb{R}^{n \times n}$ be differentiable functions. Define the composite function $h : \mathbb{R}^m \rightarrow \mathbb{R}^p$ by $h(u) = g(f(X(u)))$, where the parameter vector $u \in \mathbb{R}^m$ is mapped to a




Fig. 3. Computational graphs for the forward and reverse modes of differentiation for the composite function $h(u) = g(f(X(u)))$.

p -dimensional output. As illustrated in Fig. 3, our goal is to evaluate the sensitivity of the output h with respect to changes in the input u , which is captured by the Jacobian $J_h(u) = \frac{\partial h}{\partial u}$. This Jacobian quantifies how infinitesimal changes in the input vector u affect the output $h(u)$, and can be computed by systematically applying the chain rule to differentiate through the layers of the composite function. Applying the chain rule and elementary matrix multiplication, we can express the Jacobian as:

$$J_h(u) = J_{g \circ f \circ X}(u) = J_g(f(X(u))) \cdot J_f(X(u)) \cdot J_X(u) \quad (13)$$

This chain rule can be applied in two distinct ways depending on how we propagate the derivatives: either by starting from the inputs and working towards the outputs (*forward mode differentiation*) shown in Fig. 3a, or by starting from the outputs and working back towards the inputs (*reverse mode differentiation*) shown in Fig. 3b. Although the chain rule is the same in both cases, the direction in which it is applied leads to different computational complexities. To analyze the computational complexity of evaluating the Jacobian $J_h(u)$ in reverse mode, we must first understand the costs associated with the forward evaluations of each function in the composition $h(u) = g(f(X(u)))$. Specifically, we denote the cost of evaluating each function without differentiation as $\text{cost}(X)$, $\text{cost}(f)$, and $\text{cost}(g)$, respectively.

Our aim is to compute these derivatives via computer programs. In order to evaluate numerical derivatives rather than symbolic expressions, automatic differentiation (AD) [38,31,30,35,7] computes derivatives by accumulating intermediate values during code execution. Standard code may be easily modified to incorporate AD with little effort. The target function h must first be constructed as a series of basic operations in a computer program before applying AD. After that, AD can be used in either forward or reverse mode.

3.2. Review: forward mode differentiation via dual numbers

To begin, we will briefly go over dual numbers and their applications in forward mode AD. Dual numbers [33,53] extend the real numbers by introducing an infinitesimal unit ϵ with the property

$$\epsilon^2 = 0, \quad \epsilon \neq 0.$$

A dual number can be written in the form:

$$x = a + b\epsilon,$$

where $a, b \in \mathbb{R}$, a is called the real part, and b is called the dual part. The Taylor series expansion of a function f evaluated at a dual number $a + b\epsilon$ around a is given by:

$$f(a + b\epsilon) = f(a) + f'(a)b\epsilon + \frac{f''(a)}{2}(b\epsilon)^2 + \frac{f'''(a)}{6}(b\epsilon)^3 + \dots$$

Since the property $\epsilon^n = 0$ holds for $n \geq 2$, the Taylor series simplifies to:

$$f(a + b\epsilon) = f(a) + f'(a)b\epsilon.$$

This simplification highlights the utility of dual numbers in computing first-order derivatives. Thus, dual numbers offer a straightforward and reliable approach for computing derivatives. There exist highly performant practical implementations of dual numbers, which simply extend basic arithmetic operations and a short list of standard mathematical functions from standard floating point numbers to dual numbers. For example, the `ForwardDiff.jl` [57] package in Julia and the `cpduals` [62] library in C++ provide efficient and well-documented implementations of dual number arithmetic.

Suppose we want to compute the directional derivative of a function h with respect to m' parameters of the input vector $x \in \mathbb{R}^m$. To achieve this, *forward mode AD* [5,72,32] utilizes m' dual numbers, where each dual number, ϵ_i , represents an independent infinitesimal perturbation associated with the direction $v_i \in \mathbb{R}^m$. Specifically, it evaluates

$$h\left(x + \sum_{i=1}^{m'} \epsilon_i v_i\right) = h(x) + \sum_{i=1}^{m'} \epsilon_i \nabla h(x) \cdot v_i. \quad (14)$$

In the dense case, this approach is conceptually equivalent to computing m' directional derivatives independently. However, in sparse cases efficient techniques can be used to reduce the computational cost. More generally, forward mode AD computes the action of the Jacobian $J_h(x)$ on a matrix V of size $m \times m'$, where the columns of V are the m' directional vectors:

$$V \mapsto J_h(x)V.$$

This is often referred to as the pushforward, adapted from differential geometry terminology. Forward mode AD computes these directional derivatives by following the process shown in Fig. 3a, where derivatives are propagated step by step:

Table 1

Computational steps and their associated costs during each step of forward mode AD.

Computation	Cost
$(v_1, v'_1) = (X(c), J_X(c)V)$	$m' \cdot \text{cost}(X)$
$(v_2, v'_2) = (f(v_1), J_f(v_1) \cdot v'_1)$	$m' \cdot \text{cost}(f)$
$(v_3, v'_3) = (g(v_2), J_g(v_2) \cdot v'_2)$	$m' \cdot \text{cost}(g)$

Computing the full Jacobian, corresponds to computing all m directional derivatives and essentially starting with V as the identity matrix I_m . The total cost of evaluating $J_h(c)$ is then $m' \cdot (\text{cost}(X) + \text{cost}(f) +$

$\text{cost}(g)$ and thus depends on the methods chosen to compute the functions X , f and g . We will discuss this further in the next subsection.

3.3. Review: reverse mode differentiation

In contrast to forward mode, *reverse mode* differentiation [31,18] computes derivatives by propagating information backward from the output(s) of a function. This mode computes *adjoints* (or sensitivities) with respect to the outputs, rather than computing directional derivatives with respect to the inputs and is particularly advantageous for functions where the output dimension is smaller than the input dimension. Given $Y \in \mathbb{R}^{p \times p}$, in the cotangent space, reverse mode AD computes the action of the transposed Jacobian on Y :

$$J_h(u)^T Y, \quad (15)$$

and computes the sensitivities with respect to p outputs, which essentially construct the rows of $J_h(u)$. Once the forward evaluation is complete as shown in Fig. 3b, the adjoints are propagated in reverse as follows:

Table 2

Computational steps and their associated costs during each sweep of reverse mode AD.

Computation	Cost
$\nabla_f = J_g(f(X(u)))^T \cdot Y$	$p \cdot \text{cost}(g)$
$\nabla_X = J_f(X(u))^T \cdot \nabla_f = \delta_X[f^T \cdot \nabla_f]$	$p \cdot \text{cost}(f)$
$\nabla_u = J_X(u)^T \cdot \nabla_X = \delta_u[X^T \cdot \nabla_X]$	$p \cdot \text{cost}(X)$

In reverse mode, computing the full Jacobian involves evaluating the sensitivities with respect to all p outputs. This is achieved by initializing Y as the identity matrix I_p . Consequently, the total cost of computing $J_h(c)$ in reverse mode is given by: $p \cdot (\text{cost}(X) + \text{cost}(f) + \text{cost}(g))$, where the overall cost naturally again depends on the methods used to compute the functions X , f and g . The choice of the most efficient mode of AD (forward or reverse) depends on comparing the number of outputs p and the number of inputs m . Reverse mode is more efficient when $p \ll m$, while forward mode is better suited for cases where $p \gg m$.

Additionally, evaluating the Jacobian of matrix functions $f(X(u))$ with respect to the parameter vector u , denoted as $J_f(u) = J_f(X(u))J_X(u)$, is of significant importance. In forward mode, this computation aligns with the first two rows of Table 1, and the total cost of evaluating the full Jacobian scales as $m \cdot (\text{cost}(f) + \text{cost}(X))$.

In reverse mode, the computation corresponds to the last two rows of Table 2, with the total cost scaling as $n^2 \cdot (\text{cost}(f) + \text{cost}(X))$. Clearly, unless $n^2 < m$, reverse mode Jacobian evaluation is significantly more computationally expensive than forward mode. The corresponding complexities for evaluating $J_f(u)$ are summarized in Table 4.

It is noteworthy that advanced techniques such as sparsity detection and matrix coloring can be employed to further reduce the computational cost of evaluating the Jacobian, by efficiently identifying and exploiting the structure within the matrices involved in the differentiation process. For a detailed description of this we refer to Appendix A.2. A numerical demonstration of this approach will be provided in section 4.3.

3.4. Simplified cost analysis

To estimate the cost of evaluating $f(X)$, we revisit the methods described in Section 2. The computational cost depends on the size of the matrix X , its structural properties (e.g., sparsity or bandedness), and the number of quadrature points ℓ used in the evaluation. If we take $g = \text{Trace}$, we can rewrite $(zI - X) = LU$, and we would then need to evaluate $\text{Trace}(X(zI - X)^{-1}) = \text{Trace}(XU^{-1}L^{-1})$. Using the Cauchy integral definition (12), the cost for different types of matrices is detailed below.

Dense matrices. For a dense $n \times n$ matrix X , the cost of evaluating $f(X)$ scales as $O(\ell n^3)$, since computing the inverse of a dense matrix and performing other required operations both scale cubically with n . However, in many applications, matrices often exhibit additional structures, such as sparsity or bandedness, which significantly reduce the computational cost. This reduction is achieved because sparse linear solvers can exploit these structural properties to minimize the number of operations required.

Sparse matrices. For a sparse matrix X , the cost of evaluating $f(X)$ arises from computing

$$f(X) = \sum_i e_i^T U^{-1} L^{-1} e_i = \sum_i (U^{-T} e_i)^T (L^{-1} e_i),$$

where U and L are the LU factors of X , and e_i represents the i -th standard basis vector. Each term in the summation involves two back-substitutions: one to compute $L^{-1} e_i$ and another for $U^{-T} e_i$. The cost of each back-substitution depends on the number of nonzeros (nnz) in the LU factors.

If X is a banded matrix with a bandwidth b , the LU factorization incurs minimal fill in [17]. The cost of LU factorization is $O(nb^2)$. Assuming b is constant, since the number of nonzeros in L and U factors scale as $O(nb)$, the total cost of evaluating $f(X)$ scales as $O(\ell bn^2)$.

For 2D sparse matrices, which often arise from discretizations of 2D grids or the Hessians of (quasi-)two-dimensional structures, LU factorization introduces moderate fill-in. The cost of factorization scales as $O(n^{3/2})$, while the number of nonzeros in the L and U factors is $O(n \log n)$. Consequently, the total cost of solving linear systems and evaluating $f(X)$ scales as $O(\ell n^2 \log n)$.

For 3D sparse matrices, LU factorization incurs more significant fill-in. The factorization cost scales as $O(n^2)$, and the number of nonzeros in the L and U factors is $O(n^{4/3})$. As a result, the total cost of evaluating $f(X)$ scales as $O(\ell n^{7/3})$.

On the other hand, if eigendecomposition is used to evaluate $f(X)$, the computational cost is $O(n^3)$, as the method does not leverage the sparsity of X . Thus, the contour integration approach with its ability to exploit sparsity has clear advantages for large matrices, though for three-dimensional systems that advantage is less pronounced than in lower dimension.

We summarize the computational complexities for each case in Table 3.

Improved complexity via selected inversion. The methods we use in our numerical tests have the complexity described above. In principle one can improve the computational cost further through a selected inversion algorithm: To compute $\text{Trace}(X(zI - X)^{-1})$ for evaluating f using (12), efficient algorithms such as `Pselinv` [40], which are specifically designed to compute selected elements of a matrix inverse without requiring the full inverse. By leveraging the sparsity and structure of the matrix, `Pselinv` enables significant computational savings compared to direct inversion. The trace can be expressed as:

$$\text{Trace}(X(zI - X)^{-1}) = \sum_{i,j} X_{ij} ((zI - X)^{-1})_{ji}, \quad (16)$$

where only the entries of $(zI - X)^{-1}$ that correspond to the sparsity pattern of X are required. Using `Pselinv`, these entries are computed efficiently, avoiding the computation of the full inverse. The computational cost of `Pselinv` is comparable to the cost of LU factorization, as it directly builds upon the LU decomposition of the matrix $zI - X$. For sparse matrices, the cost scales with the number of nonzero elements and the fill-in introduced during factorization. For banded matrices with constant bandwidth, the cost is $O(n)$, for 2D sparse matrices, the cost is $O(n^{3/2})$, and for 3D sparse matrices, the cost is $O(n^2)$. `Pselinv` is both memory-efficient and highly parallelizable, making it ideal for large-scale problems where sparsity in the matrix can be exploited.

Table 3Complexity comparison for evaluating $J_h(u)$ in the dense and various sparse cases.

Mode	Sparse 1D	Sparse 2D	Sparse 3D	Dense
Forward via contour	$O(\ell bmn^2)$	$O(\ell mn^2 \log n)$	$O(\ell mn^{7/3})$	$O(\ell mn^3)$
Reverse via contour	$O(\ell bpn^2)$	$O(\ell pn^2 \log n)$	$O(\ell pn^{7/3})$	$O(\ell pn^3)$
Forward via eigendecomposition	$O(mn^3)$	$O(mn^3)$	$O(mn^3)$	$O(mn^3)$
Reverse via eigendecomposition	$O(pn^3)$	$O(pn^3)$	$O(pn^3)$	$O(pn^3)$

Table 4Comparison of complexity for evaluating $J_f(u)$ in the dense, sparse 1D, sparse 2D, and sparse 3D cases. When $p \ll n^2$, the cost of J_f is significantly higher than the cost of $J_h(u)$ and evaluating J_f is therefore almost never advantageous.

Mode	Sparse 1D	Sparse 2D	Sparse 3D	Dense
Forward via contour	$O(\ell bmn^2)$	$O(\ell mn^2 \log n)$	$O(\ell mn^{7/3})$	$O(\ell mn^3)$
Reverse via contour	$O(\ell n^4)$	$O(\ell n^4 \log n)$	$O(\ell n^{13/3})$	$O(\ell n^5)$
Forward via eigendecomposition	$O(mn^3)$	$O(mn^3)$	$O(mn^3)$	$O(mn^3)$
Reverse via eigendecomposition	$O(n^5)$	$O(n^5)$	$O(n^5)$	$O(n^5)$

3.5. Accuracy

We now turn to the question of accuracy: In general, derivatives computed using automatic differentiation have effectively arbitrarily good precision since the chain rule and explicitly known derivative rules are used. This does not directly apply to our case, however, since the function whose derivative we compute is merely approximated using a contour integral with a finite number of quadrature points. In circumstances where a function evaluation is subject to approximation errors, it is often natural to use approximate derivative schemes such as finite differences schemes of the desired accuracy. It is thus a natural question to ask whether the approximation accuracy of the contour integral methods is sufficient to where use of automatic differentiation is sensible. We will provide numerical experiments to showcase the accuracy in Sections 4 and 5.

4. Numerical experiments

We present two toy problems to illustrate the complexity and accuracy of the discussed method. Following these examples, we discuss our methodology applied to problems arising from physics, showcasing the method's performance in relevant scenarios.

4.1. Implementation details

We used the `ComplexElliptic.jl` [63] package to compute the conformal maps required to implement the contour integral matrix function as described in [34]. The forward mode AD examples in the toy problems were then computed using `ForwardDiff.jl` [57] while the reverse mode AD examples used `Zygote.jl` [39] and `ChainRules.jl` [73]. Companion code which reproduces the presented results is available [66].

4.2. Scalar derivative

We consider the following toy problem: the composition of two functions $g(f(X))$, where the functions g and f are defined as

$$g(A) = \sum_{i,j} A_{ij}^3, \quad A = f(X) = X^{1/3}.$$

The matrix X is constructed as follows to have connectivity corresponding to a one-dimensional problem setting and to ensure that it is symmetric and positive semi-definite (SPD): Given a vector $u \in \mathbb{R}^n$ (i.e., $m = n$), we define

$$X(u) = B(u)^\top B(u), \quad \text{where}$$

$$B(u) = \begin{cases} \sin(u_i + u_j) + \cos(u_i u_j) + \exp(u_i - u_j), & \text{if } |i - j| \leq 3, \\ 0, & \text{otherwise.} \end{cases}$$

Fig. 4a illustrates the evaluation time required to compute the gradient of g with respect to u using a contour integral quadrature with $\ell = 35$ quadrature points. The evaluation times are shown for both reverse mode (blue curve) and forward mode (orange curve) automatic differentiation.

Remark 4.1. As expected the reverse mode is significantly more efficient as the matrix size increases for scalar derivatives. The observed deviation from the previously mentioned complexity orders is due to the sparsity structure present in the intended applications, which were thus also included in the toy examples. Matrices involved in the differentiation process, here X and B , usually exhibit additional structure such as sparsity or even bandedness that reduce the number of operations required [21]. In particular if as in this example $m = n$, then bandedness with bandwidth s (or equivalent sparsity structure) of the $n^2 \times m$ matrix $J_X(u)$ in Equation (13) reduces the expected asymptotic complexity in reverse mode to $O(\ell sn^2)$ in line with our observations.

Fig. 4b shows the error between the second-order centered finite difference method and reverse AD using the contour integral and conformal map approach, plotted against step size for various quadrature points $\ell = 10, 15, 20, 25$. The errors decrease as ℓ increases, with higher ℓ values (e.g., $\ell \geq 25$) yielding lower errors across all step sizes. The dashed line representing $\propto x^2$ indicates the expected quadratic convergence for the finite difference method. The figure also highlights the importance of choosing an appropriate ℓ to balance accuracy and computational cost.

4.3. Jacobian

Consider the function

$$g(u) := f(X(u)) = \sqrt{\nabla^2 E(u)},$$

where $X(u) = \nabla^2 E(u)$, $f(X) = \sqrt{X}$ and $E(u), u \in \mathbb{R}^n$, is given by

$$E(C; u) = \sum_i \left(\sum_{j \in \mathcal{N}_i} C_{ij} |u_i - u_j|^2 + \frac{1}{2} |u_i - u_j|^3 + |u_i - u_j|^4 \right). \quad (17)$$




Fig. 4. Figure (a) shows obtained evaluation time for computing the gradient of function g with respect to m -vector u using a contour integral quadrature with $\ell = 35$ quadrature points. Figure (b) shows convergence of a second order centered finite difference computation of the gradient compared to a gradient computed by reverse AD through the contour integral with conformal map using ℓ quadrature points.




Fig. 5. Figure (a) shows obtained evaluation time for computing the Jacobian of matrix X with respect to m -vector u using a contour integral quadrature with $\ell = 25$ quadrature points. Figure (b) shows convergence of a second order centered finite difference computation of the Jacobian of f compared to one computed by forward mode AD through the contour integral with conformal map using ℓ quadrature points.

Here, i indexes discrete points, \mathcal{N}_i represents the set of nearest neighbors of the i -th element, C_{ij} are coefficients characterizing the interactions between i and j . We can think of $E(u)$ as a toy model for an atomistic potential energy, making this example related to the more realistic cases we consider in the next section.

Here, we present the computational complexity of Jacobian evaluation using forward differentiation, comparing the approaches with and without coloring. The details of these methods are provided in Appendix A.2. As shown in Fig. 5a, leveraging the sparsity of X through coloring reduces the Jacobian evaluation cost to $O(m^3)$, whereas without coloring, the complexity remains at $O(m^4)$, as expected.

Fig. 5b shows the error between the second-order finite difference method and forward mode AD using the contour integral approach across step sizes for various quadrature points ℓ , demonstrating that appropriate choice of quadrature points yields highly accurate results.

5. Application to defects in crystalline silicon

To demonstrate practical applications of our approach, we explore thermally activated processes in materials science, specifically focusing on defect migration, where matrix function derivatives play a crucial role. Understanding these processes is essential for predicting and optimizing the mechanical and electronic properties of materials [71,69,54].

We consider the free energy

$$F(u) = \mathcal{E}(u) - TS(u), \quad (18)$$

where $\mathcal{E}(u)$ denotes the potential energy, $S(u)$ represents the entropy, and T is the temperature. The vibrational entropy [67,25,43] in solids, $S(u)$, is defined as:

$$S(u) := -\frac{1}{2} \text{Trace} \log^+(\mathbf{F}H(u)\mathbf{F}), \quad (19)$$

where u is the lattice displacement vector, and $H(u) = \nabla^2 \mathcal{E}(u)$ is the Hessian of the potential energy $\mathcal{E}(u)$. To create a spectral gap, we define a self-adjoint operator \mathbf{F} , which serves as a preconditioner and acts as $H_{\text{hom}}^{-1/2}$, where H_{hom} is the Hessian of the homogeneous lattice [67,9]. The \log^+ function is defined as follows: let \mathbf{T} be a bounded, self-adjoint operator on a Hilbert space with spectrum $\sigma(\mathbf{T}) \subset (-\infty, 0] \cup [m, M]$, where $0 < m \leq M$. Then, we can define a contour C that encircles the interval $[m, M]$ while remaining in the right half-plane. We can then define:

$$\log^+(\mathbf{T}) = \frac{1}{2\pi i} \oint_C \log(z) \cdot \text{Trace}(zI - \mathbf{T})^{-1} dz. \quad (20)$$

Defect migration is typically analyzed using the potential energy surface (PES), represented by $\mathcal{E}(\bar{u})$ in (18), which governs the structure, dynamics, and thermodynamics of the system [44]. Migration behavior is often studied by examining stationary points on the PES, such as local minima and transition states, which map out migration pathways, often visualized as steepest-descent paths.

The Nudged Elastic Band (NEB) method [36] is widely used to determine minimum energy paths (MEPs) between known initial and final states on the PES. NEB methods discretize the reaction pathway into a series of intermediate configurations, or ‘images’, connected by elastic springs to ensure even distribution along the path. These images are optimized to remain constrained to the PES while the spring forces maintain equidistant spacing.

While MEPs on the PES are frequently used due to their simplicity, they are only approximations [76] and can become increasingly inaccurate at elevated temperatures where entropic effects play an important role. For instance, it has been demonstrated in [27,29,60,15] that the MEPs on the free energy surface (FES), $\mathcal{F}(u)$, provides a more accurate and comprehensive depiction of defect migration under finite temperature conditions. By capturing more of the thermodynamic behavior of the system, the FES allows for the analysis of migration mechanisms, rate constants, and material properties that the PES alone cannot adequately describe. Strong entropic effects have also been observed in a variety of processes, such as the nucleation of dislocation loops [58,2], the transformation of vacancy clusters into stacking-fault tetrahedra [70], the growth of nano-voids under tensile stress [52], and dislocation emission from crack tips [71]. In these cases, the role of entropy, particularly at high temperatures, becomes dominant in determining the system’s kinetics and overall behavior. Such entropic contributions, while increasingly recognized as fundamental, are impossible to capture using the PES alone, as it only accounts for potential energy minima and transition states, neglecting the broader thermodynamic landscape.

Applying optimization methods such as NEB methods to the free energy surface requires gradients of the free energy functional $\mathcal{F}(u)$ as expressed in Equation (18). These gradients which drive the optimization process necessitate computing derivatives that account for both energetic and entropic contributions. Specifically, the gradient of $\mathcal{F}(u)$ with respect to the displacement field u is

$$\frac{\partial \mathcal{F}(u)}{\partial u} = \frac{\partial \mathcal{E}(u)}{\partial u} - T \frac{\partial S(u)}{\partial u}, \quad (21)$$

where $\frac{\partial \mathcal{E}(u)}{\partial u}$ is the derivative of the potential energy, and $\frac{\partial S(u)}{\partial u}$ is the derivative of the entropy. The latter involves matrix derivatives that capture changes in entropy due to structural perturbations, which can be computationally expensive.

To address these challenges, various methods have been proposed for exploring the free energy surface in collective variable spaces. Techniques such as metadynamics [42,13], adaptive biasing force methods [19], and umbrella sampling [68] effectively sample the FES in low-dimensional spaces. However, these methods are typically less effective at locating saddle points in high-dimensional spaces, which are critical for accurately determining transition states and reaction pathways in complex systems. These limitations highlight the need for more efficient

approaches to calculate free energy derivatives in high-dimensional contexts [60].

This motivates the development of computationally efficient and highly accurate methods for matrix differentiation to advance our understanding of defect migration and similar thermally activated processes. Next, we thus illustrate the complexity and accuracy of our method for evaluating entropy derivatives and then showcase its application in some practical settings.

5.1. Implementation details

In the numerical experiments that follow we explore minimum free energy pathways for various defects in Silicon using the differentiable Stillinger-Weber interatomic potential implementation in `EmpiricalPotentials.jl` [50]. Forward mode AD is supported in `EmpiricalPotentials.jl` via the `ForwardDiff.jl` [57] package but manual adjoints had to be added for the Stillinger-Weber potential to allow for `Zygote.jl` [39] and `ChainRules.jl` [73] support. The search for minimum free energy pathways is performed using the `SaddleSearch.jl` [51] package which supports string and NEB based searches. The Stillinger-Weber potential in our examples is employed purely for the sake of convenience; any interatomic potential supporting the `ForwardDiff.jl` and `ChainRules.jl` interfaces can be used instead.

The source code used to generate the results is available in the companion code repository [66]. Animations of the migration processes are made available via FigShare, see [64], providing a visual representation of these dynamic effects. Throughout our experiments all atoms in the lattice are allowed to move during defect migration.

5.2. Complexity and convergence study

To demonstrate the computational complexity and convergence behavior of evaluating the derivative of entropy $\frac{\partial S(u)}{\partial u}$ using our implementation, we conducted tests on a system of silicon comprising 64 atoms. Fig. 6a illustrates the CPU time complexities for reverse mode and forward mode automatic differentiation (AD) based on the proposed contour integral formulation. The observed complexity for forward and reverse mode AD, for a fixed number of quadrature points, ℓ scale as $O(m^4)$ and $O(m^2)$ respectively. Fig. 6b showcases the error between a simple second order centered finite difference approach and the value given by reverse AD through the contour integral approach. The FD method effectively serves as a proxy for the true gradient up to approximately 10^{-10} , beyond which numerical precision errors limit its reliability. These observations allow one to assess the contour integral method’s accuracy as a function of ℓ and emphasize the importance of selecting an appropriate number of quadrature points. These results are fully consistent with our numerical study on toy problems in the previous section.

5.3. Vacancy migration

Vacancies in silicon remain an extensively investigated point defect in semiconductor physics, due to their role in facilitating impurity diffusion. Here we consider vacancy migration, schematically illustrated in Fig. 7, which alleviates internal stresses and influences the material’s mechanical strength and ductility [16].

We begin by generating a diamond cubic bulk silicon system containing 64 atoms per supercell. To create the initial configuration, one atom is removed, introducing a vacancy. The system is then relaxed by minimizing its energy to reach equilibrium. The free-energy profiles for vacancy migration at various temperatures is illustrated in Fig. 7b. The migration path connects the minimized initial and final configurations on the PES. Free-energy profiles were calculated at three different temperatures: 0 K, 100 K, and 200 K. The 0 K pathway corresponds to the minimum energy path (MEP) on the potential energy surface (PES). The




Fig. 6. Figure (a) shows obtained evaluation time for computing the gradient of vibrational entropy in (19) with respect to a displacement vector u of length m . Both methods used contour integral quadrature with $\ell = 20$ quadrature points. Figure (b) shows convergence of a second order centered finite difference computation of the gradient in (19) for a bulk Silicon example to a gradient computed by reverse AD through the contour integral with conformal map using ℓ quadrature points.




Fig. 7. (a) Illustration of vacancy migration in silicon. The orange atom in (1) migrates to occupy the adjacent vacancy site, resulting in configuration (2). (b) Free-energy profiles for vacancy migration in silicon at various temperatures, illustrating the influence of temperature on the migration barrier and the entropic contributions that facilitate the migration process.

vacancy migration energy at 0 K is calculated to be 0.515 eV, which is in agreement with the literature value of 0.52 eV reported by [59]. Our computed migration path at 0 K also aligns closely with the one reported in [55].

As the temperature increases, the energy barrier shows a significant variation, reflecting the dynamic nature of atomic movements and their influence on the vacancy migration pathway.

5.3.1. Interstitial

Self-interstitials in crystalline Si introduce much stronger local distortions in the crystal structure than vacancies. Interstitials typically occupy high-symmetry sites such as the tetrahedral or hexagonal interstitial sites. The migration of an interstitial atom from one site to an adjacent site, showcased in Fig. 8, is a process that influences material properties such as diffusivity and mechanical strength.

We generated a bulk silicon system consisting of 64 atoms per super-cell and generated an initial state with a tetrahedral interstitial defect by inserting an atom at $(\frac{a}{2}, \frac{a}{2}, \frac{a}{2})$, where a is the lattice constant. The final state has an interstitial inserted at an adjacent tetrahedral site. To computed the migration path between these two states by applying the

NEB method to the FES, to obtain free energy profiles at temperatures of 0 K, 200 K, and 600 K. Fig. 8 illustrates that the free energy barrier decreases slightly with increasing temperature. Thermal effects reduce the energy barrier, making interstitial migration slightly more likely at higher temperatures.

5.3.2. Variational TST

Understanding the rates of various thermally activated processes, ranging from defect formation [2] and migration [47] to creep [56] and catalysis [10], is crucial for the advancement and optimization of material properties and chemical processes. These rates are often calculated using harmonic transition state theory (HTST), where the potential energy of the lattice is approximated by a second-order Taylor expansion, effectively assuming a quadratic potential near equilibrium positions. HTST can yield inaccurate predictions when the potential energy surface significantly deviates from this quadratic assumption or when thermal effects substantially influence the system's dynamics [61]. As such, a more precise approach involves including non-harmonic corrections to resolve additional complexity found in the system's potential landscape.




Fig. 8. (a) Illustration of interstitial migration in silicon, with the red atom indicating the interstitial defect migrating from position 1 to 2. (b) Evolution of the free-energy profiles for the migration of an interstitial defect in Silicon.

Once the minimum energy path (MEP) for a reaction is determined, the transition rate can be defined within the framework of harmonic transition state theory (HTST) as:

$$\mathcal{K}^{\text{HTST}} = \exp(-\beta [\mathcal{F}(\xi_{\text{saddle}}) - \mathcal{F}(\xi_{\text{min}})]), \quad (22)$$

where ξ_{min} denotes the positions corresponding to the minimum energy configuration, and ξ_{saddle} represents the positions of the saddle configuration. The MEP can be parameterized by a reaction coordinate ξ , a scalar variable that describes the progress of the system along the reaction path. If $u(\xi)$ represents the reaction path, then ξ serves as a continuous parameter ranging from initial minimum energy configuration ($\xi = 0$) to the final minimum configuration ($\xi = 1$). The reaction coordinate provides a natural way to track changes in energy and other system properties along the MEP. The saddle point is identified as the maximum energy point along the reaction coordinate. Here, \mathcal{F} denotes the free energy, and $\beta = 1/(k_B T)$, where k_B is the Boltzmann constant and T is the temperature.

To achieve a more accurate description of defect migration and formation in silicon, one can incorporate higher-order terms in the transition state formulation. These additional terms introduce corrections to both the energy and entropy contributions, impacting the calculated transition rates and leading to a more precise and reliable model of defect dynamics. Variational transition state theory (VTST) [2,3] provides a simple correction of this kind. Below, we briefly review the idea.

One can equivalently express the energy \mathcal{E} and entropy S as a function of the reaction coordinate ξ and expand both along the reaction coordinate on the MEP around the saddle point $\xi = \xi_{\text{saddle}}$:

$$\mathcal{E}(\xi_{\text{saddle}} + \xi) \approx \mathcal{E}(\xi_{\text{saddle}}) + \frac{1}{2} \mathcal{E}_{\xi\xi}(\xi_{\text{saddle}}) \xi^2, \quad (23)$$

$$S(\xi_{\text{saddle}} + \xi) \approx S(\xi_{\text{saddle}}) + S_{\xi}(\xi_{\text{saddle}}) \xi, \quad (24)$$

where $\mathcal{E}_{\xi\xi} = \frac{\partial^2 \mathcal{E}}{\partial \xi^2}$ and $S_{\xi} = \frac{\partial S}{\partial \xi}$ are scalar functions. We can formally approximate

$$\mathcal{F}(\xi) \approx \mathcal{E}(\xi_{\text{saddle}}) + \frac{1}{2} \mathcal{E}_{\xi\xi}(\xi_{\text{saddle}}) \xi^2 - T S(\xi_{\text{saddle}}) - T S_{\xi}(\xi_{\text{saddle}}) \xi. \quad (25)$$

Since we have a maximum at the saddle point along the reaction coordinates on the MEP, we take the derivative of the free energy with respect to ξ , the displacement along the reaction coordinates. This leads to the following condition for the stationary point:

$$\frac{\partial \mathcal{F}}{\partial \xi} = \mathcal{E}_{\xi\xi}(\xi_{\text{saddle}}) \xi - T S_{\xi}(\xi_{\text{saddle}}) = 0, \quad (26)$$

with solution

$$\xi = T \frac{S_{\xi}(\xi_{\text{saddle}})}{\mathcal{E}_{\xi\xi}(\xi_{\text{saddle}})}. \quad (27)$$




Fig. 9. Estimated migration rates obtained from the variational free energy barriers for vacancy and interstitial migration.

Substituting this back into the free energy expression, we get:

$$\mathcal{F}(T) \approx \mathcal{E}(\xi_{\text{saddle}}) - T S(\xi_{\text{saddle}}) - \frac{1}{2} T^2 \frac{S_{\xi}^2(\xi_{\text{saddle}})}{\mathcal{E}_{\xi\xi}(\xi_{\text{saddle}})}. \quad (28)$$

Thus,

$$\mathcal{F}_{\text{VTST}}(T) \approx \mathcal{F}_{\text{HTST}}(T) - \frac{1}{2} T^2 \frac{S_{\xi}^2(\xi_{\text{saddle}})}{\mathcal{E}_{\xi\xi}(\xi_{\text{saddle}})}. \quad (29)$$

Therefore, the temperature-dependent variational transition rate is:

$$k_{\text{VTST}}(T) = \exp(-\beta \mathcal{F}_{\text{HTST}}) \exp\left(\beta T^2 \frac{S_{\xi}^2(\xi_{\text{saddle}})}{2 \mathcal{E}_{\xi\xi}(\xi_{\text{saddle}})}\right) = k_{\text{HTST}} \exp(\beta T^2 T_e) \quad (30)$$

where $T_e = \frac{S_{\xi}^2(\xi_{\text{saddle}})}{2 \mathcal{E}_{\xi\xi}(\xi_{\text{saddle}})}$, and is referred to as the “effective temperature”.

As shown in previous sections, S_{ξ} can be easily obtained using our differentiation framework without the need to rely on approximation or sampling techniques which are often used for this evaluation [2]. To compare the formation rates estimated using harmonic and variational TST, we examine the migration of vacancy and interstitial defects based on the free energy profiles discussed in the previous subsections.

Fig. 9 presents the migration rates obtained using harmonic transition state theory (HTST) and variational transition state theory (VTST).

For vacancy migration, as the temperature increases, the difference between the predictions of VTST and HTST becomes more apparent. This indicates that harmonic TST overestimates the migration rate and is, therefore, less reliable for high-temperature processes. This is evident from the fact that $\mathcal{E}_{ww}(\xi_{\text{saddle}}) < 0$, leading to

$$\exp\left(-\beta \frac{T^2}{2} T_e\right) < 1.$$

In contrast, for interstitial migration, the temperature correction T_e is small, and the overestimation of the rate by HTST is negligible even at elevated temperatures.

For certain systems entropic contributions become increasingly significant at elevated temperatures as the vibrational modes of atoms exhibit greater anharmonicity. This leads to a wider distribution of energy states that harmonic TST can fail to capture. More advanced methods such as VTST can be used in such circumstances to obtain more reliable predictions of migration rates.

Remark 5.1. (Scalability) The numerical experiments of Sections 4–5 employ matrices up to hundreds of rows, while many problems of scientific interest may be of a much larger scale. The cost estimates in Section 3 show that the contour formulation becomes increasingly asymptotically favorable for much larger systems, provided that the linear solvers optimally exploit sparsity. When the linear systems become very large, the bottleneck is no longer the contour integral itself but the memory needed for the sparse LU factorizations. At that point it is natural to switch to a distributed sparse solver [8,24,75,45], which splits both the storage and the computations over several processors. Advanced techniques such as selected inversion [40,46] can further reduce the computational cost for very large systems.

6. Conclusion

In this paper we showed how to accurately and efficiently compute the derivatives of functions of matrix families X_θ and their derivatives with respect to the entries of θ_j of θ . The method is an extension of an established approach for matrix functions due to Hale, Higham and Trefethen [34] and thus begins with the Cauchy integral definition of matrix functions and computes appropriate conformal maps and quadrature in the complex plane. We explained in detail the resulting computational complexity for both reverse and forward mode differentiation.

With the extension to derivatives in hand, we demonstrated the utility of the approach in natural applications in the context of molecular simulations and material modeling. We considered questions involving the entropy of Silicon systems which require repeated computation of matrix trace-logarithms and gradients thereof for saddle point searches. As the conformal map based method allows defining straightforward adjoints for the matrix logarithm, all that is required in principle for efficient reverse mode automatic differentiation of such matrix functions is a reverse mode compatible implementation of the involved interatomic potential (in the case of our Silicon example this role is played by the Stillinger-Weber potential).

Our aim was to demonstrate that employing modern differentiation algorithms not only results in numerically robust and accurate derivatives but also in improved computational performance, and that such an approach is entirely practical within a modern software stack.

CRedit authorship contribution statement

Tina Torabi: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Investigation, Formal analysis, Conceptualization. **Timon S. Gutleb:** Writing – review & editing, Writing – original draft, Software, Investigation, Formal analysis, Conceptualization. **Christoph Ortner:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

TT and CO were supported by NSERC Discovery grant RGPIN-2021-03489 and NFRF Exploration Grant GR022937. TSG was partially supported by a PIMS-Simons postdoctoral fellowship, jointly funded by the Pacific Institute for the Mathematical Sciences (PIMS) and the Simons Foundation.

Appendix A

A.1. Jacobi elliptic functions

We begin with some definitions and properties of Jacobi elliptic functions as e.g. detailed in [1,11].

A.1.1. Elliptic functions

A doubly periodic meromorphic function is called an *elliptic function*. Let a parameter m and complementary parameter p be given satisfying

$$m + p = 1. \quad (31)$$

If $m \in \mathbb{R}$ then in what follows one can assume without loss of generality that $0 \leq m \leq 1$.

A.1.2. Quarter-periods

The quarter-periods K and iK' are defined by the integrals

$$K(m) = K = \int_0^{\frac{\pi}{2}} \frac{d\theta}{\sqrt{1 - m \sin^2 \theta}}, \quad (32)$$

$$iK'(m) = iK' = \int_0^{\frac{\pi}{2}} \frac{d\theta}{\sqrt{1 - p \sin^2 \theta}}, \quad (33)$$

where $K, K' \in \mathbb{R}$. K is called the real quarter-period and iK' the imaginary quarter-period. These quantities satisfy the relationships

$$K(m) = K(p) = K'(1 - m). \quad (34)$$

Moreover, if any one of the numbers $m, p, K(m), K'(m), \frac{K'(m)}{K(m)}$ are given, all the rest are uniquely determined through the above relations, i.e. K and K' cannot be independently chosen.

A.1.3. Jacobi elliptic functions

The Jacobi elliptic functions are widely studied standard forms of elliptic functions represented by $\text{sn}(u, k)$, $\text{cn}(u, k)$, and $\text{dn}(u, k)$, where k is termed the elliptic modulus. They originate from the inverse of the elliptic integral of the first kind,

$$u = F(\varphi, k) = \int_0^\varphi \frac{dt}{\sqrt{1 - k^2 \sin^2 t}}, \quad (35)$$

where φ as $\text{am}(u, k)$ denotes the Jacobian amplitude. This leads to the relationships

$$\sin \varphi = \text{sn}(\text{am}(u, k)) = \text{sn}(u, k), \quad (36)$$

$$\cos \varphi = \text{cn}(\text{am}(u, k)) = \text{cn}(u, k), \quad (37)$$

$$\sqrt{1 - k^2 \sin^2 \varphi} = \sqrt{1 - k^2 \sin^2(\text{am}(u, k))} = \text{dn}(u, k). \quad (38)$$

These functions extend trigonometric functions to be doubly periodic, satisfying:

$$\text{sn}(u, 0) = \sin u, \quad (39)$$

$$\text{cn}(u, 0) = \cos u, \quad (40)$$

$$\text{dn}(u, 0) = 1. \quad (41)$$

A.1.4. Identities for Jacobi elliptic functions

The established identities for Jacobi elliptic functions are given by the following equations:

$$\text{sn}^2 u + \text{cn}^2 u = 1, \quad (42)$$

$$k^2 \text{sn}^2 u + \text{dn}^2 u = 1, \quad (43)$$

$$k^2 \text{cn}^2 u + k'^2 = \text{dn}^2 u, \quad (44)$$

$$\text{cn}^2 u + k'^2 \text{sn}^2 u = \text{dn}^2 u. \quad (45)$$

A.1.5. Specific values

Specific noteworthy values are listed below:

$$\text{cn}(0, k) = \text{cn}(0) = 1, \quad (46)$$

$$\text{cn}(K(k), k) = \text{cn}(K(k)) = 0, \quad (47)$$

$$\text{dn}(0, k) = \text{dn}(0) = 1, \quad (48)$$

$$\text{dn}(K(k), k) = \text{dn}(K(k)) = k' = \sqrt{1 - k^2}, \quad (49)$$

$$\text{sn}(0, k) = \text{sn}(0) = 0, \quad (50)$$

$$\text{sn}(K(k), k) = \text{sn}(K(k)) = 1, \quad (51)$$

where $K = K(k)$ signifies the complete elliptic integral of the first kind and $k' = \sqrt{1 - k^2}$ represents the complementary elliptic modulus [74]. Fig. 10 illustrates the Jacobi elliptic functions for $k = 0.7$.

A.1.6. Complex arguments

When dealing with complex arguments, the Jacobi elliptic functions can be extended as follows:

$$\text{sn}(u + iv) = \frac{\text{sn}(u, k)\text{dn}(v, k')}{1 - \text{dn}^2(u, k)\text{sn}^2(v, k')} + \frac{i\text{cn}(u, k)\text{dn}(u, k)\text{sn}(v, k')\text{cn}(v, k')}{1 - \text{dn}^2(u, k)\text{sn}^2(v, k')}, \quad (52)$$

$$\text{cn}(u + iv) = \frac{\text{cn}(u, k)\text{cn}(v, k')}{1 - \text{dn}^2(u, k)\text{sn}^2(v, k')} - \frac{i\text{sn}(u, k)\text{dn}(u, k)\text{sn}(v, k')\text{dn}(v, k')}{1 - \text{dn}^2(u, k)\text{sn}^2(v, k')}, \quad (53)$$

$$\text{dn}(u + iv) = \frac{\text{dn}(u, k)\text{cn}(v, k')\text{dn}(v, k')}{1 - \text{dn}^2(u, k)\text{sn}^2(v, k')} - \frac{ik^2\text{sn}(u, k)\text{cn}(u, k)\text{sn}(v, k')}{1 - \text{dn}^2(u, k)\text{sn}^2(v, k')}. \quad (54)$$

A.2. Strategies for utilizing matrix sparsity

Leveraging potential sparsity of the matrix X can significantly enhance the efficiency of computing its Jacobian. Consider the toy model presented in Section 4.3:

$$E(C; u) = \sum_i \left(\sum_{j \in \mathcal{N}_i} C_{ij} |u_i - u_j|^2 + \delta |u_i - u_j|^3 + |u_i - u_j|^4 \right). \quad (55)$$

Consider a system composed of 8 atoms. Our objective is to compute the Jacobian of the Hessian matrix, which corresponds to the matrix of second derivatives of the energy (17) with respect to the displacements in such a system. The Jacobian is shown in Fig. 11 (a). Looking at this sparse matrix, we can see that there is a lot of free space in the matrix. We aim to compress the sparse matrix into a denser format shown in Figs. 11 (b) and (c), allowing for the processing of non-zero entries




Fig. 10. Jacobi Elliptic Functions $\text{sn}(u, k)$, $\text{cn}(u, k)$, $\text{dn}(u, k)$ for $k = 0.7$.

with significantly fewer function calls than previously necessary. Although both matrices (d) and (e) represent the Jacobian in a denser format, the main task is to find the smallest dense matrix which in our case is (e). This technique uses a strategy from graph theory [26,41], aiming to combine columns with non-overlapping non-zero elements, referred to as *structurally orthogonal columns* into single groups, thus reducing the total number of groups needed. To achieve this, one can use a number of graph coloring algorithms including Contraction Coloring, Greedy distance- k coloring, and Backtracking Sequential Coloring. We will briefly go over the Greedy distance-1 coloring algorithm which results were also shown in Fig. 11 (f)-(g) [65].

We begin with a brief overview of essential graph theory terminology. A graph G is formally defined as an ordered pair $G = (V, E)$, where V is a finite, non-empty set of vertices, and E is a set comprising unordered pairs of distinct vertices, known as edges. Vertices u and v are adjacent if an edge (u, v) is included in E ; otherwise, they are described as non-adjacent. Within a graph, a path of length l (measured in edges) is a sequence of vertices v_1, v_2, \dots, v_{l+1} , where each consecutive pair (v_i, v_{i+1}) is adjacent, for all $1 \leq i \leq l$, with each vertex appearing uniquely in the sequence. Vertices u and v are considered distance- k neighbors if the shortest path between them has a length of k or less. The set of all distance- k neighbors of a vertex u , denoted $N_k(u)$, does not include u itself. Additionally, if two vertices are distance- k neighbors, they also qualify as distance- k' neighbors for any $k' > k$.

A graph is categorized as bipartite if its vertex set V can be partitioned into two disjoint subsets V_1 and V_2 , such that every edge connects a vertex from V_1 to one from V_2 . There are no edges between vertices within the same subset, which ensures that V_1 and V_2 comprehensively separate the vertices of the graph.

A distance- k vertex coloring of a graph $G = (V, E)$ is a labeling function $\varphi : V \rightarrow \{1, 2, \dots, p\}$ that assigns different colors to any pair of distance- k neighbors. The minimum number of colors needed to establish such a coloring for graph G is termed the k -chromatic number, denoted as $\chi_k(G)$. If only a specific subset $W \subset V$ of the vertices is colored, the coloring is referred to as partial. Specifically, a partial distance- k coloring on W is defined by a function $\varphi : W \rightarrow \{1, 2, \dots, p\}$ such that $\varphi(u) \neq \varphi(v)$ for any two vertices u and v within W that are distance- k neighbors. An example of such a distance-2 coloring is illustrated on the right side of Fig. 11.

The following lemma provides a graph-theoretical characterization of structural orthogonality for a nonsymmetric matrix.

Lemma A.1. [26, Lemma 3.4] Let A be an $m \times n$ matrix and $G_b(A) = (V_1, V_2, E)$ be its bipartite graph. Two columns (or rows) in A are structurally




Fig. 11. The process of compressing the Jacobian matrix of the Hessian of toy model (17) with 8 atoms in two partitions (a)–(e). Each partition is also represented as a distance-1 coloring in a column intersection graph (f)–(g).

orthogonal if and only if the corresponding vertices in $G_b(A)$ are at a distance greater than two from each other.

According to Lemma A.1, determining a structurally orthogonal partition of the columns of a matrix A corresponds to attaining a partial distance-2 coloring of its bipartite graph $G_b(A) = (V_1, V_2, E)$ specifically applied to V_2 . This coloring is termed partial as it does not extend to V_1 . Theorem 3.5 provides a formal statement of this relationship.

Theorem A.1. [26, Theorem 3.5] Let A be an $m \times n$ matrix and $G_b(A) = (V_1, V_2, E)$ represent its bipartite graph. A function φ constitutes a partial distance-2 coloring of $G_b(A)$ on V_2 if and only if φ generates a structurally orthogonal partition of the columns of A .

Now we can reformulate our main problem as follows: Given the bipartite graph $G_b(A) = (V_1, V_2, E)$ that represents the sparsity structure of an $m \times n$ matrix A , the objective is to find a partial distance-2 coloring of $G_b(A)$ on V_2 that utilizes the minimal number of colors. This coloring strategy aims to efficiently partition the columns of A into structurally orthogonal sets, corresponding to the vertices in V_2 , while minimizing the color count.

To solve this problem we can use the greedy distance-2 coloring algorithm, shown in 1. We initially need to represent our Jacobian using a graph. Given a matrix $A \in \mathbb{R}^{m \times n}$, we can represent the structure of A using a bipartite graph $G = (V, E)$. The construction process for the bipartite graph is outlined as follows:

1. Create two disjoint sets of vertices, U and W , where U corresponds to the rows of A and W corresponds to the columns of A . Thus, $U = \{u_1, u_2, \dots, u_m\}$ and $W = \{w_1, w_2, \dots, w_n\}$.
2. The set of vertices V is the union of U and W , i.e., $V = U \cup W$.
3. For each non-zero entry A_{ij} in the matrix A , add an edge (u_i, w_j) to the set of edges E . This implies there is an edge between vertex $u_i \in U$ and vertex $w_j \in W$ if and only if the entry A_{ij} is non-zero.

This bipartite graph representation, G , captures the interactions between the rows and columns of A based on its non-zero entries. It allows for a visual and analytical understanding of the matrix's structure, and can help us use the coloring algorithms that exploit sparsity patterns for improved computational efficiency. Below, we outline the detailed

steps and the corresponding algorithm that collectively streamline this process:

1. **Graph Representation of Sparsity:** First, the sparsity pattern of the Jacobian matrix J is modeled as a graph $G = (V, E)$, where the vertices V represent the matrix's rows and columns, and the edges E correspond to non-zero entries in J .
2. **Applying a Greedy Distance-2 Coloring Algorithm:** The coloring of graph G is executed through a methodical greedy algorithm that ensures no two vertices within two edges of each other share the same color, reflecting the distance-2 coloring strategy necessary for avoiding computational interference:
 - (a) Initially, all vertices are uncolored.
 - (b) The coloring process is iterative, with each uncolored vertex being assigned the least positive integer color that is not used by its distance-1 and distance-2 neighbors, ensuring that the coloring satisfies the distance-2 constraints.
3. **Constructing Perturbation Vectors:** For every unique color assigned, a corresponding perturbation vector d is created. This vector has elements set to 1 for indices colored with the current color and 0 elsewhere.
4. **Evaluating the Jacobian:** Using each perturbation vector d , the function $F(x + \epsilon d)$ is evaluated to ascertain the non-zero components of the Jacobian matrix that correspond to each color. This step is crucial for identifying which parts of the matrix can be independently calculated, thereby enhancing computational efficiency.

Algorithm 1, operationalizes the second step of this process. Once a group is determined, centered finite difference or automatic differentiation [63] can be used to calculate the directional derivatives along the compressed matrix directions. For more details on the graph coloring method for computing derivatives we refer to [26].

Data availability

I have shared the code at the attach file step.

[EntropyGrad.jl \(Original data\)](#) (Github)
[Animations of migration paths for defects in silicon crystals at 100 K \(Original data\)](#) (Figshare)

Algorithm 1 A greedy distance-2 coloring algorithm. [26, Algorithm 3.1].

```

1: procedure D2COLORINGALG( $G = (V, E)$ )
2:   Let  $v_1, v_2, \dots, v_{|V|}$  be a given ordering of  $V$ 
3:   Initialize ForbiddenColors with some value  $a \notin V$ 
4:   for  $i \leftarrow 1$  to  $|V|$  do
5:     for each colored vertex  $w \in N_1(v_i)$  do
6:       ForbiddenColors[ $color[w]$ ]  $\leftarrow v_i$ 
7:     for each colored vertex  $x \in N_1(w)$  do
8:       ForbiddenColors[ $color[x]$ ]  $\leftarrow v_i$ 
9:     end for
10:   end for
11:    $color[v_i] \leftarrow \min\{c > 0 : \text{ForbiddenColors}[c] \neq v_i\}$ 
12: end for
13: end procedure

```

References

- [1] M. Abramowitz, I.A. Stegun (Eds.), *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Dover Publications Inc., New York, 1965, p. 1046.
- [2] S. Bagchi, D. Perez, Anomalous entropy-driven kinetics of dislocation nucleation, <https://arxiv.org/abs/2402.00810>, 2024.
- [3] J.L. Bao, D.G. Truhlar, Variational transition state theory: theoretical framework and recent developments, *Chem. Soc. Rev.* 46 (2017) 7548–7596, <https://doi.org/10.1039/C7CS00602K>.
- [4] I. Batatia, et al., Equivariant matrix function neural networks, <https://arxiv.org/abs/2310.10434>, 2024.
- [5] A.G. Baydin, et al., Automatic differentiation in machine learning: a survey, *J. Mach. Learn. Res.* (ISSN 1532-4435) 18 (1) (Jan. 2017) 5595–5637.
- [6] M. Benzi, P. Boito, Matrix functions in network analysis, *GAMM-Mitt.* 43 (3) (2020) e202000012, <https://doi.org/10.1002/gamm.202000012>.
- [7] B. van den Berg, et al., Forward- or reverse-mode automatic differentiation: what's the difference?, *Sci. Comput. Program.* (ISSN 0167-6423) 231 (2024) 103010, <https://doi.org/10.1016/j.scico.2023.103010>.
- [8] L.S. Blackford, et al., *SciLAPACK Users' Guide*, Society for Industrial and Applied Mathematics, Philadelphia, PA, ISBN 0-89871-397-8, 1997 (paperback).
- [9] J. Braun, M.H. Duong, C. Ortner, Thermodynamic limit of the transition rate of a crystalline defect, *Arch. Ration. Mech. Anal.* 238 (3) (2020) 1413–1474.
- [10] R.J. Bunting, et al., Reactivity of single-atom alloy nanoparticles: modeling the dehydrogenation of propane, *J. Am. Chem. Soc.* 145 (27) (2023) 14894–14902, <https://doi.org/10.1021/jacs.3c04030>.
- [11] P.F. Byrd, M.D. Friedman, *Handbook of Elliptic Integrals for Engineers and Scientists*, 2nd ed., Grundlehren der mathematischen Wissenschaften, Springer-Verlag Berlin Heidelberg, Berlin, Heidelberg, ISBN 978-3-642-65140-3, 1971, pp. XVI, 360.
- [12] D. Caillard, J.-L. Martin, *Thermally Activated Mechanisms in Crystal Plasticity*, Pergamon Materials Series, vol. 8, Pergamon, Amsterdam, Boston, Mass, ISBN 978-0080427034, 2003.
- [13] E. Carter, et al., Constrained reaction coordinate dynamics for the simulation of rare events, *Chem. Phys. Lett.* (ISSN 0009-2614) 156 (5) (1989) 472–477, [https://doi.org/10.1016/S0009-2614\(89\)87314-2](https://doi.org/10.1016/S0009-2614(89)87314-2).
- [14] H. Chen, C. Ortner, QM/MM methods for crystalline defects. Part 1: locality of the tight binding model, *Multiscale Model. Simul.* 14 (1) (2016) 232–264, <https://doi.org/10.1137/15M1022628>.
- [15] M. Chen, T.-Q. Yu, M.E. Tuckerman, Locating landmarks on high-dimensional free energy surfaces, *Proc. Natl. Acad. Sci. USA* 112 (11) (2015) 3235–3240, <https://doi.org/10.1073/pnas.1418241112>.
- [16] P.G. Coleman, Activation energies for vacancy migration, clustering and annealing in silicon, *J. Phys. Conf. Ser.* 265 (1) (Jan. 2011) 012001, <https://doi.org/10.1088/1742-6596/265/1/012001>.
- [17] T.A. Davis, S. Rajamanickam, W.M. Sid-Lakhdar, A survey of direct methods for sparse linear systems, *Acta Numer.* 25 (2016) 383–566, <https://doi.org/10.1017/S0962492916000076>.
- [18] J.E. Dennis, R.B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Society for Industrial and Applied Mathematics, 1996.
- [19] B.M. Dickson, et al., Free energy calculations: an efficient adaptive biasing potential method, *J. Phys. Chem. B* 114 (17) (May 2010) 5823–5830.
- [20] NIST Digital Library of Mathematical Functions, Release 1.2.1 of 2024-06-15. in: F.W.J. Olver, A.B. Olde Daalhuis, D.W. Lozier, B.I. Schneider, R.F. Boisvert, C.W. Clark, B.R. Miller, B.V. Saunders, H.S. Cohl, M.A. McClain (Eds.), <https://dlmf.nist.gov/>.
- [21] I.S. Duff, A.M. Erisman, J.K. Reid, *Direct Methods for Sparse Matrices*, Oxford University Press, Jan. 2017, 9780198508380.
- [22] Y. Erlangga, C. Vuik, C. Oosterlee, On a class of preconditioners for solving the Helmholtz equation, *Appl. Numer. Math.* (ISSN 0168-9274) 50 (3) (2004) 409–425, <https://doi.org/10.1016/j.apnum.2004.01.009>.
- [23] H. Fatoorehchi, S. Djilali, Stability analysis of linear time-invariant dynamic systems using the matrix sign function and the Adomian decomposition method, *Int. J. Dyn. Control* 11 (July 2022), <https://doi.org/10.1007/s40435-022-00989-3>.
- [24] S. Filippone, M. Colajanni, PSBLAS: a library for parallel linear algebra computation on sparse matrices, *ACM Trans. Math. Softw.* (ISSN 0098-3500) 26 (4) (Dec. 2000) 527–550, <https://doi.org/10.1145/365723.365732>.
- [25] B. Fultz, Vibrational thermodynamics of materials, *Prog. Mater. Sci.* 55 (4) (2010) 247–352, <https://doi.org/10.1016/j.pmatsci.2009.05.002>.
- [26] A. Gebremedhin, F. Manne, A. Pothen, What color is your Jacobian? Graph coloring for computing derivatives, *SIAM Rev.* 47 (2005) 629–705, <https://doi.org/10.1137/S0036144504444711>.
- [27] A. Glensk, et al., Breakdown of the Arrhenius law in describing vacancy formation energies: the importance of local anharmonicity revealed by ab initio thermodynamics, *Phys. Rev. X* 4 (Feb. 2014) 011018, <https://doi.org/10.1103/PhysRevX.4.011018>.
- [28] G.H. Golub, C.F. Van Loan, *Matrix Computations*, Eng. Baltimore, 2013.
- [29] B. Grabowski, et al., Ab initio up to the melting point: anharmonicity and vacancies in aluminum, *Phys. Rev. B* 79 (Apr. 2009) 134106, <https://doi.org/10.1103/PhysRevB.79.134106>.
- [30] A. Griewank, K. Kulshreshtha, A. Walther, On the numerical stability of algorithmic differentiation, *Computing* 94 (2012) 125–149.
- [31] A. Griewank, A. Walther, *Evaluating Derivatives*, second edition, Society for Industrial and Applied Mathematics, 2008.
- [32] A. Griewank, On automatic differentiation, in: M. Iri, K. Tanabe (Eds.), *Mathematical Programming*, Kluwer Academic Publishers, Dordrecht, ISBN 079230490X, 1989, pp. 83–108.
- [33] Y.-L. Gu, J.Y.S. Luh, Dual-number transformation and its applications to robotics, *IEEE J. Robot. Autom.* 3 (1987) 615–623.
- [34] N. Hale, N.J. Higham, L.N. Trefethen, Computing A^a , $\log(A)$, and related matrix functions by contour integrals, *SIAM J. Numer. Anal.* 46 (5) (2008) 2505–2523, <https://doi.org/10.1137/070700607>.
- [35] R. Hecht-Nielsen, Theory of the backpropagation neural network, in: *International 1989 Joint Conference on Neural Networks*, vol. 1, 1989, pp. 593–605.
- [36] G. Henkelman, H. Jónsson, Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points, *J. Chem. Phys.* (ISSN 0021-9606) 113 (22) (Dec. 2000) 9978–9985, <https://doi.org/10.1063/1.1323224>.
- [37] N.J. Higham, *Functions of Matrices*, Society for Industrial and Applied Mathematics, 2008.
- [38] P. Hoffmann, A Hitchhiker's guide to automatic differentiation, *Numer. Algorithms* 72 (July 2016), <https://doi.org/10.1007/s11075-015-0067-6>.
- [39] M. Innes, Don't unroll adjoint: differentiating SSA-form programs, *CoRR*, arXiv:1810.07951, 2018, <http://arxiv.org/abs/1810.07951>.
- [40] M. Jacquelin, L. Lin, C. Yang, PSellInv—a distributed memory parallel algorithm for selected inversion: the symmetric case, *ACM Trans. Math. Softw.* (ISSN 0098-3500) 43 (Dec. 2016) 3, <https://doi.org/10.1145/2786977>.
- [41] M. Kubale, *Graph Colorings*, Contemporary Mathematics, vol. 352, American Mathematical Society, ISBN 978-0821834589, 2004.
- [42] A. Laio, M. Parrinello, Escaping free-energy minima, *Proc. Natl. Acad. Sci. USA* 99 (20) (Oct. 2002) 12562–12566.
- [43] C. Lapointe, et al., Machine learning surrogate models for strain-dependent vibrational properties and migration rates of point defects, *Phys. Rev. Mater.* 6 (Nov. 2022) 113803, <https://doi.org/10.1103/PhysRevMaterials.6.113803>.
- [44] E. Lewars, *Computational Chemistry: Introduction to the Theory and Applications of Molecular and Quantum Mechanics*, Springer, ISBN 978-90-481-3860-9, Jan. 2011, pp. 1–664.
- [45] X.S. Li, J.W. Demmel, SuperLU_DIST: a scalable distributed-memory sparse direct solver for unsymmetric linear systems, *ACM Trans. Math. Softw.* (ISSN 0098-3500) 29 (2) (June 2003) 110–140, <https://doi.org/10.1145/779359.779361>.
- [46] L. Lin, SellInv - an algorithm for selected inversion of a sparse symmetric matrix, in: Lawrence Berkeley National Laboratory, 2010.
- [47] S. List, H. Ryssel, Atomistic analysis of the vacancy diffusion mechanism, in: 1996 International Conference on Simulation of Semiconductor Processes and Devices, SISPAD '96 (IEEE Cat. No.96TH8095), 1996, pp. 27–28.
- [48] X. Liu, et al., Solving the three-dimensional high-frequency Helmholtz equation using contour integration and polynomial preconditioning, *SIAM J. Matrix Anal. Appl.* 41 (1) (2020) 58–82, <https://doi.org/10.1137/18M1228128>.
- [49] A.M. Mathai, *A Handbook of Generalized Special Functions for Statistical and Physical Sciences*, Eng. Oxford, 1993.
- [50] C. Ortner, et al., EmpiricalPotentials.jl, <https://github.com/JuliaMolSim/EmpiricalPotentials.jl>.
- [51] C. Ortner, et al., SaddleSearch.jl, <https://github.com/cortner/SaddleSearch.jl>.
- [52] D. Perez, et al., Entropic stabilization of nanoscale voids in materials under tension, *Phys. Rev. Lett.* 110 (May 2013) 206001, <https://doi.org/10.1103/PhysRevLett.110.206001>.
- [53] D. Piponi, Automatic differentiation, C++ templates, and photogrammetry, *J. Graph. Tools* 9 (Jan. 2004), <https://doi.org/10.1080/10867651.2004.10504901>.
- [54] S. Pogatscher, et al., Diffusion on demand to control precipitation aging: application to Al-Mg-Si alloys, *Phys. Rev. Lett.* 112 (22 June 2014) 225701, <https://doi.org/10.1103/PhysRevLett.112.225701>.

- [55] M. Posselt, F. Gao, H. Bracht, Correlation between self-diffusion in Si and the migration mechanisms of vacancies and self-interstitials: an atomistic study, *Phys. Rev. B* 78 (3 July 2008) 035208, <https://doi.org/10.1103/PhysRevB.78.035208>.
- [56] M. Rajaguru, S.M. Keralavarma, A discrete dislocation model of creep in single crystals, in: TMS 2016 145th Annual Meeting & Exhibition, Springer International Publishing, Cham, ISBN 978-3-319-48254-5, 2016, pp. 351–358.
- [57] J. Revels, M. Lubin, T. Papamarkou, Forward-mode automatic differentiation in Julia, arXiv:1607.07892 [cs.MS], 2016, <https://arxiv.org/abs/1607.07892>.
- [58] S. Ryu, K.W. Kang, W. Cai, Entropic effect on the rate of dislocation nucleation, *Proc. Natl. Acad. Sci. USA* 108 (2010) 5174–5178.
- [59] P. Śpiewak, K.J. Kurzydłowski, Formation and migration energies of the vacancy in Si calculated using the HSE06 range-separated hybrid functional, *Phys. Rev. B* 88 (19 Nov. 2013) 195204, <https://doi.org/10.1103/PhysRevB.88.195204>.
- [60] G. Shrivastav, E. Vanden-Eijnden, C.F. Abrams, Mapping saddles and minima on free energy surfaces using multiple climbing strings, *J. Chem. Phys.* (ISSN 0021-9606) 151 (12) (Sept. 2019) 124112, <https://doi.org/10.1063/1.5120372>.
- [61] T.D. Swinburne, M.-C. Marinica, Unsupervised calculation of free energy barriers in large crystalline systems, *Phys. Rev. Lett.* 120 (Mar. 2018) 135503, <https://doi.org/10.1103/PhysRevLett.120.135503>.
- [62] M. Tesch, cppduals: a nestable vectorized templated dual number library for C++11, *J. Open Source Softw.* 4 (43) (2019) 1487, <https://doi.org/10.21105/joss.01487>.
- [63] T. Torabi, tinatorabi/ComplexElliptic.jl: ComplexElliptic (v1.0), 2024.
- [64] T. Torabi, T.S. Gutleb, C. Ortner, Animations of migration paths for defects in silicon crystals at 100 K, in: FigShare, Sept. 2024.
- [65] T. Torabi, Vibrational entropy: spatial decomposition and surrogate models, MA thesis, University of British Columbia, 2024.
- [66] T. Torabi, T.S. Gutleb, EntropyGrad.jl, <https://doi.org/10.5281/zenodo.14510426>.
- [67] T. Torabi, Y. Wang, C. Ortner, Surrogate models for vibrational entropy based on a spatial decomposition, arXiv preprint arXiv:2402.12744, 2024.
- [68] G. Torrie, J. Valleau, Nonphysical sampling distributions in Monte Carlo free-energy estimation: umbrella sampling, *J. Comput. Phys.* (ISSN 0021-9991) 23 (2) (1977) 187–199, [https://doi.org/10.1016/0021-9991\(77\)90121-8](https://doi.org/10.1016/0021-9991(77)90121-8).
- [69] D. Turnbull, Theory of grain boundary migration rates, *JOM* (ISSN 1543-1851) 3 (8) (Aug. 1951) 661–665, <https://doi.org/10.1007/BF03397362>.
- [70] B.P. Uberuaga, et al., Direct transformation of vacancy voids to stacking fault tetrahedra, *Phys. Rev. Lett.* 99 (Sept. 2007) 135501, <https://doi.org/10.1103/PhysRevLett.99.135501>.
- [71] D. Warner, W. Curtin, Origins and implications of temperature-dependent activation energy barriers for dislocation nucleation in face-centered cubic metals, *Acta Mater.* (ISSN 1359-6454) 57 (14) (2009) 4267–4277, <https://doi.org/10.1016/j.actamat.2009.05.024>.
- [72] R. Wengert, A simple automatic derivative evaluation program, *Commun. ACM* 7 (8) (1964) 463–464.
- [73] F. White, JuliaDiff/ChainRules.jl: v1.61.0, 2024.
- [74] E.T. Whittaker, G.N. Watson, A Course of Modern Analysis, 4th ed., Cambridge Mathematical Library, Cambridge University Press, 1996.
- [75] J. Zhang, et al., The PetscSF scalable communication layer, *IEEE Trans. Parallel Distrib. Syst.* 33 (4) (2022) 842–853.
- [76] J.A. Zimmerman, H. Gao, F.F. Abraham, Generalized stacking fault energies for embedded atom FCC metals, *Model. Simul. Mater. Sci. Eng.* 8 (2) (Mar. 2000) 103, <https://doi.org/10.1088/0965-0393/8/2/302>.