

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Data Compression for Optimization of a Molecular Dynamics System: Preserving Basins of Attraction

Michael Retzlaff¹, Todd Munson², and Zichao (Wendy) Di²

¹ Department of Mathematics and Statistics, University of Maryland Baltimore County, Baltimore, MD, USA

² Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL, USA

Abstract. Understanding the evolution of atomistic systems is essential in various fields such as materials science, biology, and chemistry. The gold standard for these calculations is molecular dynamics, which simulates the dynamical interaction between pairs of molecules. The main challenge of such simulation is the numerical complexity, given a vast number of atoms over a long time scale. Furthermore, such systems often contain exponentially many optimal states, and the simulation tends to get trapped in local configurations. Recent developments leverage the existing temporal evolution of the system to improve the stability and scalability of the method; however, they suffer from large data storage requirements. To efficiently compress the data while retaining the basins of attraction, we have developed a framework to determine the acceptable level of compression for an optimization method by application of a Kantorovich-type theorem, using binary digit rounding as our compression technique. Choosing the Lennard-Jones potential function as a model problem, we present a method for determining the local Lipschitz constant of the Hessian with low computational cost, thus allowing the use of our technique in real-time computation.

Keywords: lossy compression · basins of attraction · nonlinear optimization · Lennard-Jones potential.

1 Introduction

Simulating atomistic evolution is essential for predicting materials properties for use in materials science, chemistry, and biology. Molecular dynamics (MD), the gold standard for atomistic simulations, simulates the interactions of atoms and molecules for a fixed period of time [1, 14]. The limited temporal scale is inherent from the system’s sequential nature but poses great challenges for observing various transitions on MD time scales. Recent advances leverage statistics and massively parallel computers to accelerate the simulations, such as the accelerated molecular dynamics method [18, 15] and the parallel replica dynamics method (ParRep) [19]. In particular, the parallel trajectory splicing method (ParSplice)

[12], as a generalization of ParRep, has drawn much attention in recent years because of its superior computational performance by utilizing large-scale high-performance computers to parallelize the state-to-state simulations in the time domain. One major burden, however, is the need to store all the local minimizers computed in the long MD simulation, since they are used as the initial guess for the next state simulation. Since the number of local minimizers is exponential in the number of atoms, storage space must be conserved, even though each individual minimizer is a small set of numbers (e.g., 3D coordinates of the atom positions). Lossy compression of these local minimizers can potentially result in significant savings. By compressing the values, we can cache more values locally on the compute node and send/receive less information from the global database. However, a critical step is to identify the basins of local convergence for a particular numerical method. As long as the compressed data stays in the same basin, we can recover the local minimizer by applying the numerical method to the stored iterate.

In this work, for an atomic system and Newton’s method as the underlying solver, we employ a Kantorovich-type theorem for studying the basins of attraction of potential energy minima. We propose a mathematical framework to estimate bounds for the compression level of an optimizer when using binary digit rounding as the compression routine. One critical step is to calculate the required information, namely, the local Lipschitz-type constant of the function’s derivative. As the system gets large, an analytic expression for this information is likely impossible. Therefore, a robust method to approximate this information is required. By exploring the topology of the potential energy landscape, we propose an efficient way to estimate the required constant. We prototype our method on the simple, yet common, Lennard-Jones potential function and validate our approach on various numerical examples.

2 Minimization of Lennard-Jones Potential

Given a configuration of n atoms in a cluster $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, we consider the Lennard-Jones (LJ) potential function [8], a simplified model that simulates the potential energy of the cluster based on inter-atom distances:

$$\hat{V}_n(\mathbf{X}) = 4\varepsilon \sum_{i < j} \left(\left(\frac{\sigma}{\hat{d}(\mathbf{x}_i, \mathbf{x}_j)} \right)^{12} - \left(\frac{\sigma}{\hat{d}(\mathbf{x}_i, \mathbf{x}_j)} \right)^6 \right), \quad (2.1)$$

where $\hat{d}(\mathbf{x}_i, \mathbf{x}_j)$ is the Euclidean distance between atoms $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^3$. The physical constants ε and σ are the depth of the potential well and inter-atom reaction limit, respectively, both depending on the type of atom. Because of its importance in computational chemistry, finding optimal configurations that locally minimize the potential energy (2.1) remains an active research area. For example, Maranas et al. [10] proposed an exotic optimization algorithm to find many stationary points, and Asenjo et al. [3] studied the mapping of the basins of attraction for various optimization algorithms.




Fig. 1: Left: landscape of function V_2 given x_2 as its only free variable. Right: landscape of function V_3 with reduced dimension by setting $x_2 = 2^{1/6}$ (its equilibrium distance) with DOF x_3, y_3 .

In this work, we follow the longstanding convention in [7] and consider the reduced-unit optimization problem³

$$\min_{\mathbf{X}} V_n(\mathbf{X}) = \sum_{i < j} \left(\frac{1}{d(\mathbf{x}_i, \mathbf{x}_j)^{12}} - \frac{1}{d(\mathbf{x}_i, \mathbf{x}_j)^6} \right), \quad (2.2)$$

whereby we set $d(\mathbf{x}_i, \mathbf{x}_j) = \frac{\hat{d}(\mathbf{x}_i, \mathbf{x}_j)}{\sigma}$ and $V_n = \frac{\hat{V}_n}{4\varepsilon}$ to eliminate the parameters in (2.1). The coordinates of the atoms $\mathbf{x}_i = (x_i, y_i, z_i)$, $i = 1, \dots, n$ are unknown variables; and we fix atom 1 at $(0, 0, 0)$, atom 2 at $(x_2, 0, 0)$, and atom 3 at $(x_3, y_3, 0)$. In addition to eliminating the need to postprocess a collection of distances in order to visualize the system, using coordinates in this manner offers some mathematical advantages. First, for $n \geq 3$ the function V_n has $3n - 6$ degrees of freedom (DOF), while the pairwise distance formulation has $O(n^2)$ DOF. Second, using coordinates eliminates the need for both configuration feasibility and non-negative constraints on the variables, resulting in an unconstrained optimization problem. Third, fixing the coordinates of atoms 1–3 in the described manner ensures distinguishable minimizers with reduced quantity. However, the reduced LJ potential function V_n still has $O(\exp(n^2))$ unique minimizers [7], and even enumerating the quantity of local minima of a cluster has been shown to be an NP-hard problem [21]. Furthermore, each unique minimizer has at least $O(n^3)$ equivalent permutations (for example, by interchange of atom numbering and rotation).

Figure 1 illustrates the functions V_2 and V_3 , respectively. We can see that minimizing even such small systems can be challenging. First, when any pairwise atom distance approaches zero, there exists a singularity in V_n . Second, when any pairwise distance $d(\mathbf{x}_i, \mathbf{x}_j)$ is large, the norm of the corresponding gradient components are (numerically) very small. That is, there can be an infinite number of (nearly) critical points when $\|\mathbf{x}_i\| \rightarrow \infty$ for some i . As such, a careful

³ Simplifying the units in this manner does not alter the topology of the critical points.

choice of initial guess is necessary to ensure that the optimization algorithm converges to the desirable minimizer. Many researchers (e.g., [17, 23]) leverage geometric principles in their numerical methods to generate initial iterates.

3 Method and Small System Validation

In this section, we survey existing theorems that can guarantee the local convergence of problem (2.2). The application of these theorems allows us to determine a maximal amount that a local minimizer can be perturbed under a lossy compression technique, while preserving its basin of attraction. That is, the given optimization algorithm applied to the compressed data point converges to the same optimal point. To quantify the amount of compression in terms of the acceptable perturbation, we use binary digit rounding as a simple compression technique.

For the optimization algorithm applied to recover the minimizer, we use Newton's method to solve the first-order optimality conditions of (2.2), since convergence can be guaranteed under certain conditions when the iterates are *near* a root. Among the most well-known convergence results for Newton's method, the Kantorovich Theorem (KT) [9] continues to draw attention by many researchers (e.g., [2, 13, 24]). The crux of the KT result guarantees that a solution \mathbf{X}^* of $U_n(\mathbf{X}) = V'_n(\mathbf{X}) = 0$ exists in a neighborhood of a point $\mathbf{X}^{(0)}$ under some assumptions, and that Newton's method converges to \mathbf{X}^* when $\mathbf{X}^{(0)}$ is the initial iterate. Note that KT does not require knowledge of the optimal point, only information based on the initial iterate, a workflow opposite to our compression framework where \mathbf{X}^* is given. Instead, we exploit a variant of KT called the *ball of convergence about \mathbf{X}^** that assumes its prior knowledge.

First, we define $B(\mathbf{X}^*, r)$ as an open ball in a Banach space \mathcal{B} with center \mathbf{X}^* and radius $r > 0$, that is, $B(\mathbf{X}^*, r) = \{\mathbf{X} \in \mathcal{B} : \|\mathbf{X} - \mathbf{X}^*\| < r\}$. In our case, for any U_n , $n \geq 3$, we have $\mathcal{B} = \mathbb{R}^{3n-6}$, and we choose the Euclidean norm on \mathcal{B} , abbreviating $\|\mathbf{X}\| := \|\mathbf{X}\|_2$. We use the same convention for the compatible matrix operator norm, namely, $\|U'_n(\mathbf{X})\| := \|U'_n(\mathbf{X})\|_2 = \max_{\mathbf{Y} \in \mathcal{B}} \frac{\|U'_n(\mathbf{X})\mathbf{Y}\|_2}{\|\mathbf{Y}\|_2}$.

Theorem 1. (*Rheinboldt [16]*) For $U_n : N \subseteq \mathbb{R}^m \rightarrow \mathbb{R}^m$, suppose $U_n(\mathbf{X}^*) = 0$ and that $(U'_n(\mathbf{X}^*))^{-1}$ exists. Furthermore, for some $r_1 > 0$, there exists a constant $L_* > 0$ such that

$$\|(U'_n(\mathbf{X}^*))^{-1}(U'_n(\mathbf{X}) - U'_n(\mathbf{Y}))\| \leq L_* \|\mathbf{X} - \mathbf{Y}\| \quad (3.3)$$

for all $\mathbf{X}, \mathbf{Y} \in B(\mathbf{X}^*, r_1)$. Then, if $r_1 < \frac{2}{3L_*}$, Newton's method converges uniquely to \mathbf{X}^* for any $\mathbf{X}^{(0)} \in B(\mathbf{X}^*, r_1)$.

Alternatively, we have the following.

Theorem 2. (Wang [20]) For $U_n : N \subseteq \mathbb{R}^m \rightarrow \mathbb{R}^m$, suppose $U_n(\mathbf{X}^*) = 0$ and that $(U'_n(\mathbf{X}^*))^{-1}$ exists. Furthermore, for some $r_2 > 0$, there exists a constant $L_{**} > 0$ such that

$$\|(U'_n(\mathbf{X}^*))^{-1}U'_n(\mathbf{X}) - I\| \leq L_{**}\|\mathbf{X} - \mathbf{X}^*\| \quad (3.4)$$

for all $\mathbf{X} \in B(\mathbf{X}^*, r_2)$. Then, if $r_2 \leq \frac{2}{L_{**}}$, Newton's method converges uniquely to \mathbf{X}^* for any $\mathbf{X}^{(0)} \in B(\mathbf{X}^*, r_2)$.

A ball $B(\mathbf{X}^*, r)$ satisfying either Theorem 1 or 2 is a *ball of convergence* for \mathbf{X}^* , and any compressed minimizer that lies in $B(\mathbf{X}^*, r)$ preserves the basin of attraction of \mathbf{X}^* under Newton's method.

Notice that the given theorems require that $U_n(\mathbf{X}^*)$ be exactly zero. For computer/numerical calculations this assumption is almost never true. To adapt this condition, one can perform a comprehensive convergence test for terminating the iterates to ensure that the sequences $\{U_n(\mathbf{X}^{(k)})\}$ and $\{\mathbf{X}^{(k)}\}$ are both converging. This numerical issue is negligible when accurate minimizers are computed. All computed \mathbf{X}^* for our numerical tests are accurate to at least 10 decimal places.

While different lossy compression techniques can be implemented depending on their intended use, our work focuses on providing a theoretical bound on the maximum amount of perturbation that any lossy compression scheme can afford, in order to preserve a basin of attraction. We use significant binary digit truncation/rounding as the underlying compression technique since it is computationally cheap to implement and easy to quantify. Furthermore, the resulting relative perturbations are small and easy to control, since rounding a real number to c significant binary digits (bits) produces a maximum relative perturbation of $\frac{2^{-(c+1)}}{1 + 2^{-(c+1)}}$. For example, rounding to 4 significant bits has a maximum relative perturbation of 3.03%, and to 8 bits with more than 0.0195%, a potentially significant reduction in memory from the standard 52 significant bits of IEEE double precision. We use $\mathbf{X}^{(0),c}$ to denote \mathbf{X}^* compressed to c significant bits, which is used as the initial guess for Newton's method.

To understand the usage of the stated theory and how it may be applied at large scales, we offer a brief analysis of the 2-atom (1 free variable) LJ potential function V_2 . We apply Newton's method to the first-order conditions of (2.2) by solving $0 = U_2(\mathbf{X}) = U_2(x_2) = -12x_2^{-13} + 6x_2^{-7}$, for $x_2^* = 2^{1/6}$. To apply Theorems 1 and 2, we must choose radii r_1 and r_2 small enough to meet their corresponding Lipschitz inequality. This step may not be straightforward, however, because both L_* and L_{**} depend on r . In this one-dimensional case, we can explicitly solve for these radii because we know that the maximum values of $|U'_2|$ and $|U''_2|$ are attained at $2^{1/6} - r$ for any $0 < r < 2^{1/6}$. We have the following:

$$\begin{aligned} \left| \frac{-2184(2^{1/6} - r_1)^{-15} + 336(2^{1/6} - r_1)^{-9}}{156(2^{1/6})^{-14} - 42(2^{1/6})^{-8}} \right| &< \frac{2}{3r_1} \implies r_1 < 0.02432, \\ \left| \frac{156(2^{1/6} - r_2)^{-14} - 42(2^{1/6} - r_2)^{-8}}{r_2(156(2^{1/6})^{-14} - 42(2^{1/6})^{-8})} - \frac{1}{r_2} \right| &\leq \frac{2}{r_2} \implies r_2 \leq 0.06276. \end{aligned}$$

Therefore, given the radii needed for different truncations $r_c = |x_2^{(0),c} - x_2^*|$, either the Rheinboldt or the Wang criterion can guarantee that Newton's method converges to $x_2^* = 2^{1/6}$ for compression as far as 4 bits ($r_c = 0.0025$). Note that Newton's method succeeds beyond the theoretical guarantee for all compressed $x_2^{(0),c}$ down to 1 bit, but it does fail for $c = 0$, where $x_2^{(0),c} = 2.0$ and the Newton iterates coincide with large x_2 of nearly critical configurations.

In this case, choosing Theorem 1 or 2 does not make a difference in the amount of binary digit compression. However, the ball of convergence provided by Theorem 2 can be three times more generous, given the different prefactors in the radius conditions. More specifically, if the derivative U'_n is linear, for any configurations \mathbf{X}, \mathbf{Y} , we have

$$\frac{\|U'_n(\mathbf{X}) - U'_n(\mathbf{X}^*)\|}{\|\mathbf{X} - \mathbf{X}^*\|} = \frac{\|U'_n(\mathbf{X}) - U'_n(\mathbf{Y})\|}{\|\mathbf{X} - \mathbf{Y}\|}.$$

Provided that $U'_n(\mathbf{X})$ is approximately linear near \mathbf{X}^* , then

$$\sup_{\mathbf{X} \in B(\mathbf{X}^*, r)} \frac{\|(U'_n(\mathbf{X}^*))^{-1}U'_n(\mathbf{X}) - I\|}{\|\mathbf{X} - \mathbf{X}^*\|} \approx \sup_{\mathbf{X}, \mathbf{Y} \in B(\mathbf{X}^*, r)} \frac{\|(U'_n(\mathbf{X}^*))^{-1}(U'_n(\mathbf{X}) - U'_n(\mathbf{Y}))\|}{\|\mathbf{X} - \mathbf{Y}\|}.$$

Thus $r_2 \approx 3r_1$, suggesting that Wang's theorem can offer a ball of convergence up to 3 times the size of Rheinboldt's. Furthermore, the univariate calculation of L_{**} in (3.4) is more easily *tractable* than the bivariate calculation required in (3.3), in the sense that $\mathbf{Y} = \underset{\mathbf{X} \in B(\mathbf{X}^*, r)}{\operatorname{argsup}} \left\{ \frac{\|(U'_n(\mathbf{X}^*))^{-1}U'_n(\mathbf{X}) - I\|}{\|\mathbf{X} - \mathbf{X}^*\|} \right\}$ can be described as a single configuration(s), leading to better estimation of $L_{**}(n, r)$. For these reasons, we focus on applying Wang's theorem through the rest of the paper.

Given the goal to provide a radius r_2 guaranteed to satisfy Theorem 2, so that Newton's method converges to \mathbf{X}^* for any initial guess $\mathbf{X}^{(0),c} \in B(\mathbf{X}^*, r_2)$, we propose Alg. 1 as the first contribution.

Algorithm 1 Derivation of allowable radius for compression of \mathbf{X}^*

- 1: Given a local minimizer \mathbf{X}^* and radius $r > 0$ of perturbation in lossy compression, calculate

$$L_{**}(n, r) = \sup_{\mathbf{X} \in B(\mathbf{X}^*, r)} \frac{\|(U'_n(\mathbf{X}^*))^{-1}U'_n(\mathbf{X}) - I\|}{\|\mathbf{X} - \mathbf{X}^*\|}. \quad (3.5)$$

- 2: Define the radius of the convergence ball $r_2 = \min \left\{ r, \frac{2}{L_{**}(n, r)} \right\}$.
 - 3: (*optional*) If either $r \ll \frac{2}{L_{**}(n, r)}$ or $\frac{2}{L_{**}(n, r)} \ll r$, choose \hat{r} between r and $\frac{2}{L_{**}(r)}$, and return to step 1 using \hat{r} .
 - 4: Any compressed value $\mathbf{X}^{(0),c} \in B(\mathbf{X}^*, r_2)$ converges to \mathbf{X}^* using Newton's method.
-

Figure 2 illustrates the results of applying Alg. 1 to U_6 for various radii to find an acceptable convergence ball radius. The curve for $2/L_{**}$ is due to the approximation described in the forthcoming sections. Note that an acceptable radius can always be derived after one guess of r ; but by following step 3 and refining to \hat{r} based on the results of step 1, continual improvement can be realized until the user decides to stop.




Fig. 2: Behavior of Alg. 1 for U_6 . For any r , the result is the minimum of r and $2/L_{**}(r)$, which for incremental r is shown by the black markers. The maximal r based on L_{**}^{est} is 0.0352.

4 Approximation of L_{**}

The 2-atom case allows us to explicitly solve for the radius of convergence because we know the location of $\operatorname{argsup}\{L_{**}(2, r)\}$. As the system becomes large ($n \geq 3$), however, finding an analytic solution for this information is exceedingly difficult. The critical step is to find $L_{**}(n, r)$ for a given r satisfying (3.4). For any r less than the distance to the nearest critical configuration, we assume that the supremum for (3.5) is located at the boundary of the ball.⁴ Then, problem (3.5) is equivalent to finding the configuration \mathbf{X} which maximizes

$$L_{**}(n, r) = \max_{\mathbf{X} \in \partial B(\mathbf{X}^*, r)} \left\{ \frac{\|(U'_n(\mathbf{X}^*))^{-1} U'_n(\mathbf{X}) - I\|}{r} \right\}, \quad (4.6)$$

where ∂B denotes the boundary of the ball B .

It is straightforward that certain entries of $U'_n(\mathbf{X})$ approach infinity if \mathbf{X} approaches a *singular configuration* $\mathbf{X}^{(s)}$, where two or more atoms coincide and V_n returns infinite energy. Therefore, any singular configuration maximizes $\|U'_n(\mathbf{X}) - U'_n(\mathbf{X}^*)\|$. Furthermore, the shortest route to perturb a configuration to be singular is to move the two closest atoms toward each other. These observations inspire us to propose an efficient approach to approximate L_{**} in Alg. 2, which is denoted as L_{**}^{est} .

⁴ We observe this assumption to be true for all U_n through numerical experimentation.

Algorithm 2 Approximation of L_{**} via the shortest path

-
- 1: Choose a radius of perturbation r .
 - 2: Choose the closest pair of atoms in the \mathbf{X}^* configuration (with respect to the perturbation required to move the pair to singularity).
 - 3: Find $\mathbf{X}^{(p)}$ by perturbing \mathbf{X}^* with only the chosen pair of atoms from step 2, toward each other along their shortest path, in the amount of r (when $n \geq 6$ the path is linear and atom movement is equal at distance of $r/2$).
 - 4: Compute $L_{**}^{est}(n, r) = \frac{\|(U'_n(\mathbf{X}^*))^{-1}U'_n(\mathbf{X}^{(p)}) - I\|}{r}$.
-

To validate Alg. 2, we first approximate L_{**} by random sampling, a practical approach to approximate Lipschitz-type constants [22]. Specifically, given a minimizer \mathbf{X}^* and radius r , we generate $K = 10^6$ vectors $\{\mathbf{d}_k, k = 1, \dots, K\}$, where each component of \mathbf{d}_k is a uniformly chosen random number in $[-1, 1]$, and normalized so that $\|\mathbf{d}_k\| = 1$. Then we compute

$$L_{**}^{rand}(n, r) = \max_{k=1 \dots K} \left\{ \frac{\|(U'_n(\mathbf{X}^*))^{-1}U'_n(\mathbf{X}_k) - I\|}{r} : \begin{array}{l} \mathbf{X}_k = \mathbf{X}^* + r\mathbf{d}_k, \\ \|\mathbf{d}_k\| = 1 \end{array} \right\}. \quad (4.7)$$

Figure 3 tracks the configurations $\mathbf{X}^{(r)} = \operatorname{argmax} \{L_{**}^{rand}(3, r)\}$ for increasing radii, ranging from 0.01 to 1.05, incremented by 0.02. Again, we enforce the semi-fixed positions for atoms 1–3 as described earlier. \mathbf{X}^* is the equilateral triangle labeled by the stars. Each $\mathbf{X}^{(r)}$ configuration consists of a cyan marker (position of atom 3), a magenta marker (position of atom 2), and the fixed atom 1. As r increases, atom 3 of $\mathbf{X}^{(r)}$ moves along the cyan path and atom 2 of $\mathbf{X}^{(r)}$ moves along the magenta path, where for both atoms a darker shade corresponds to a configuration with larger r . Eventually, the two atoms meet at $(0.84, 0, 0)$ when $r = 1.05$, which coincides with the nearest singular configuration of \mathbf{X}^* . As hypothesized, the observed path of the $\mathbf{X}^{(r)}$ follows the shortest path as described in Alg. 2.

The tendency for the configurations $\mathbf{X}^{(r)}$ to track along the shortest path continues for larger systems, illustrated in Fig. 4 for U_4 (left) and U_6 (right). In each plot, 40 configurations of $\mathbf{X}^{(r)} = \operatorname{argmax} \{L_{**}^{rand}(n, r)\}$ are displayed, corresponding to $0 \leq r \leq 0.8$. As before, any $\mathbf{X}^{(r)}$ configuration consists of an atom of each color, as well as atom 1, while the star markers indicate the equilibrium positions of the atoms, which in these cases form a regular tetrahedron and a regular octahedron, respectively. With U_4 , the position of atom 2 remains near its \mathbf{X}^* position in all $\mathbf{X}^{(r)}$, while atoms 3 and 4 perturb toward each other as r increases. For U_6 , atoms 2 and 3 remain near their respective \mathbf{X}^* coordinates in all $\mathbf{X}^{(r)}$, while either the pair of atoms 4 and 5 or the pair of atoms 5 and 6 perturb toward each other as r increases. Specifically, for any r , the $\mathbf{X}^{(r)}$ configuration has atoms 5 and 6 perturb toward each other at an approximate distance of $r/2$ each, while atoms 1 to 4 remain nearly stationary, or, atoms 4 and 5 move together about $r/2$ and the remaining atoms perturb very little. This alternating behavior occurs as the nearest singular configuration to \mathbf{X}^* is




Fig. 3: For $0.01 \leq r \leq 1.05$, 53 configurations of $\mathbf{X}^{(r)}$ which maximize $L_{**}^{rand}(3, r)$. Each $\mathbf{X}^{(r)}$ configuration consists of one cyan marker (atom 3), one magenta marker (atom 2), and atom 1. Darker markers indicate an $\mathbf{X}^{(r)}$ from a greater r , while the star markers are the position of the atoms in the optimal configuration. As r increases, the $\mathbf{X}^{(r)}$ follow the shortest path toward the nearest singular configuration of \mathbf{X}^* .

not unique, since both the 4–5 and 5–6 pair singularities are equidistant from \mathbf{X}^* . Note that once the system is large enough ($n \geq 6$), the path followed by $\mathbf{X}^{(r)}$ toward the nearest singular configuration is linear, since the nearest pairs of atoms each have three free variables.




Fig. 4: For $n = 4$ (left) and $n = 6$ (right), 40 configurations of $\mathbf{X}^{(r)}$ which maximize $L_{**}^{rand}(n, r)$, for $0 \leq r \leq 0.8$. Atoms of the $\mathbf{X}^{(r)}$ are coordinated by color tone. The perturbed configurations follow the shortest path toward the nearest singular configuration $\mathbf{X}^{(s)}$, as indicated by the arrows. For U_6 , there are two $\mathbf{X}^{(s)}$, so the $\mathbf{X}^{(r)}$ follows both paths, alternating by chance.

Now, we compare the two estimations of L_{**} , from the proposed Alg. 2 and from random sampling due to (4.7). We focus on the cases $n \geq 6$ whose shortest paths are strictly linear. Taking the global minimum configurations from [5] for each \mathbf{X}^* , in Fig. 5, we compare the behavior of L_{**}^{rand} and L_{**}^{est} for U_6, U_8 , and

U_{10} , separately in each column.⁵ The top row shows for the cases when $r \leq 0.8$, which is the approximate distance to the nearest singularity for these \mathbf{X}^* . The bottom row zooms in to the cases when $r \leq 0.05$. The curve of $L = 2/r$ is drawn so that an estimate for the radius of convergence can be inferred from the plots by its intersection with the L_{**} curves. It is clear that $L_{**}^{est}(n, r) \geq L_{**}^{rand}(n, r)$ for all test cases, and therefore $L_{**}^{est}(n, r)$ is a more accurate estimate. For brevity, the cases $n = 7, 9$ are not plotted here, but consistent behavior is observed at all radii.

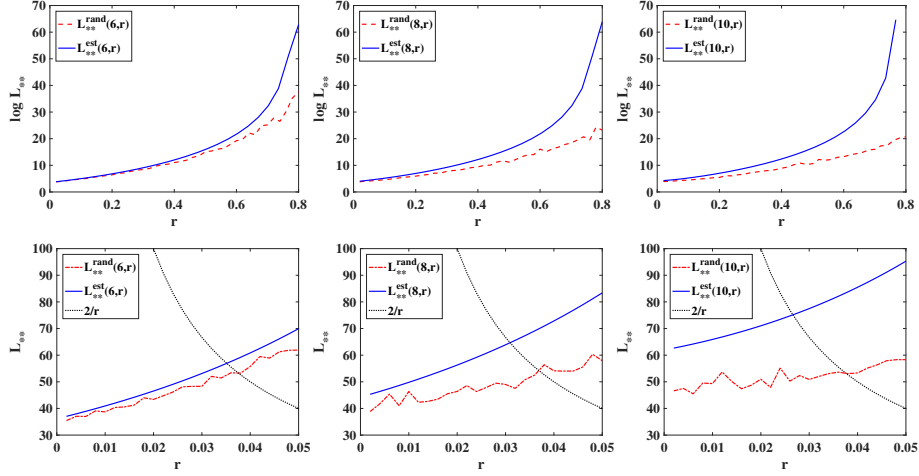


Fig. 5: Comparison of L_{**}^{est} with L_{**}^{rand} for U_6, U_8, U_{10} , showing L_{**}^{est} is a better estimate for L_{**} . The curve of $2/r$ is included since $L_{**} \leq 2/r$ determines the radius of the ball of convergence.

Next, we compute the corresponding binary digit compression levels for systems of size $n = 6, \dots, 10$, and we test the convergence of the compressed minimizers using Newton's method. Table 1 catalogs these results. Its first column shows the system size by number of atoms; the second column shows the returned radius of the ball of convergence (to the nearest 0.002) after applying the proposed method, and the third column shows the corresponding compression levels in terms of number of bits. The last column shows that, in practice, Newton's method still converges to the same \mathbf{X}^* with even higher compression levels than the theoretical guarantee.

5 Numerical Results in Higher Dimensions

In this section, we apply the proposed method to larger atomic systems to show its scalability. For the optimal configuration chosen, we use the global minimums

⁵ Similar results hold for other local minimizers.

Table 1: For $n = 6, \dots, 10$, the radii of convergence, maximal compression ability based on L_{**}^{est} , and convergence results for the compressed minimizers $\mathbf{X}^{(0),c}$ using Newton’s method. The theoretical c_t obtained by our method provides a guarantee for compression so that $c \geq c_t$ implies that $\mathbf{X}^{(0),c}$ converges to \mathbf{X}^* with Newton’s method.

| n | $\max r : r \leq 2/L_{**}^{est}$ | theoretical c_t (in bits): $\ \mathbf{X}^{(0),c} - \mathbf{X}^*\ \leq r$ | practical c_p (in bits): $\mathbf{X}^{(0),c} \rightarrow \mathbf{X}^*$ |
|-----|----------------------------------|--|---|
| 6 | 0.034 | 5 | 2:24 |
| 7 | 0.034 | 6 | 2:24 |
| 8 | 0.030 | 6 | 4:24 |
| 9 | 0.026 | 7 | 4:24 |
| 10 | 0.026 | 6 | 5:24 |

tabulated in [5] and first published in [7, 11]. Because of the partially fixed positions of atoms 1, 2, and 3, we observe that some permutations of a minimizer \mathbf{X}^* (in terms of interchanging atom numbering) can cause the spectral radius of $(U'_n(\mathbf{X}^*))^{-1}$ to be multiple orders of magnitude larger than others, resulting in significantly larger $L_{**}(r)$ than other permutations. This observation is particularly true when atom 3 is nearly on the x -axis. Therefore, we choose permutations of these minimizers so that the matrix $(U'_n(\mathbf{X}^*))^{-1}$ has lower spectral radius.

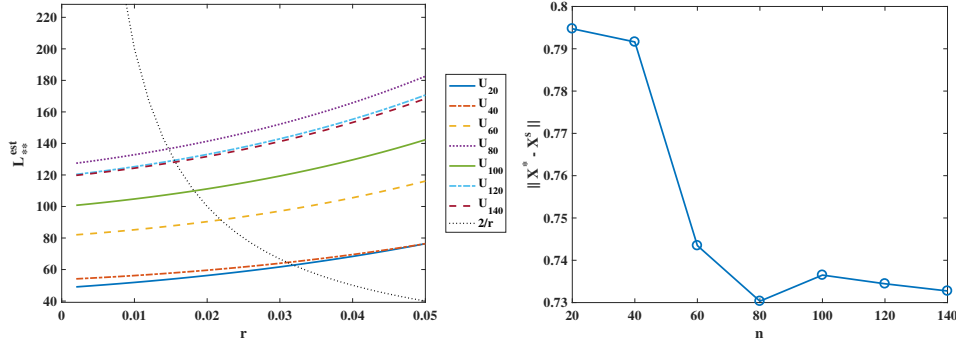


Fig. 6: Left: $L_{**}^{est}(r)$ values for the U_{20} to U_{140} functions. The curve of $2/r$ is included because $L_{**} \leq 2/r$ determines the radius of the ball of convergence. Right: distances to nearest singularity $\|\mathbf{X}^* - \mathbf{X}^{(s)}\|$ for each \mathbf{X}^* . An increase in $L_{**}(n, r)$ corresponds to a decrease in $\|\mathbf{X}^* - \mathbf{X}^{(s)}\|$.

For systems with $n = 20, 40, 60, 80, 100, 120$, and 140 atoms, Fig. 6 (left) shows the shortest-path estimates $L_{**}^{est}(n, r)$, respectively. For a fixed radius r , the corresponding $L_{**}^{est}(n, r)$ are increasing in n for $n \leq 80$, but the values plateau

Table 2: Compression ability based on L_{**}^{est} and corresponding convergence results of Newton’s method for U_{20} through U_{140} functions.

| n | $\max r : r \leq 2/L_{**}^{est}$ | $L_{**}^{est}(r)$ | $L_{**}^{rand}(r)$ | theoretical c_t (in bits): $\ \mathbf{X}^{(0),c} - \mathbf{X}^*\ \leq r$ | practical c_p (in bits): $\mathbf{X}^{(0),c} \rightarrow \mathbf{X}^*$ |
|-----|----------------------------------|-------------------|--------------------|--|---|
| 20 | 0.032 | 62.89 | 30.22 | 6 | 5:24 |
| 40 | 0.030 | 63.96 | 25.28 | 7 | 5:24 |
| 60 | 0.020 | 90.39 | 25.08 | 8 | 5:24 |
| 80 | 0.014 | 136.0 | 30.01 | 10 | 5:24 |
| 100 | 0.018 | 109.7 | 23.42 | 9 | 5:24 |
| 120 | 0.014 | 128.1 | 23.87 | 10 | 6:24 |
| 140 | 0.014 | 127.0 | 23.18 | 10 | 6:24 |

for $80 \leq n \leq 140$. These changes in $L_{**}^{est}(n, r)$ correspond to a decrease and leveling off of the distances to nearest singularity $\|\mathbf{X}^* - \mathbf{X}^{(s)}\|$ for each \mathbf{X}^* , which are plotted in Fig. 6 (right). It follows from Hoare [6] that the smallest inter-atom distances in any n -atom structure cannot fall far below the equilibrium distance of $2^{1/6}$, so that these observed distances to singularity have approached a lower bound. Indeed, for the 1000 atom minimizer recorded in [23], the minimal distance to singularity is 0.733. Further, we have $L_{**}^{est}(1000, 0.014) = 135.5$, a corresponding maximum compression level $c_t = 13$, and Newton’s method succeeds for all $\mathbf{X}^{(0),c}$ for $c \geq 9$. As both this minimal distance to singularity and L_{**}^{est} are consistent with the 80-140 atom cases, this is an indication that our proposed method scales to much larger atom systems.

Now, we further emphasize the accuracy of L_{**}^{est} by comparing it to the random sampling approximation from (4.7), again with $K = 10^6$ guesses. The comparison is done at the radius that we deduced is optimal using Alg. 1 with L_{**}^{est} . As shown in Table 2, the shortest-path estimates are significantly higher than the random sampling estimates in all test cases. The corresponding compression levels deduced from L_{**}^{est} are reported in the fifth column. We apply Newton’s method on $\mathbf{X}^{(0),c}$ from all compression levels, where the last column shows that Newton’s method converges to the same \mathbf{X}^* for compression levels exceeding the theoretical guarantee.

6 Conclusions

In this work, we propose a framework to guide the amount of compression that can be applied to an optimal atom configuration with the Lennard-Jones potential function, so that the compressed value converges to the same optimal configuration under Newton’s method. Furthermore, we develop a shortest-path estimate by exploring the topology of the Lennard-Jones function to provide the necessary information required by the framework. We show that the shortest-path estimate outperforms the estimate from random sampling. We also demonstrate the reliability and stability of the proposed framework for systems of

varying size, so that the provided compression level guarantees the preservation of the basins of attraction.

Our proposed framework also can be applied in a wider range of applications. One potential application is the compression of the large neural nets in deep learning. As the trained neural net is moved from a resource-rich setting to different model architectures with restrictive computation constraints such as memory and speed, the large-scale model with billions of trained weights and hyperparameters needs to be compressed. In this case, one can study the compression accuracy given the proposed framework so that certain accuracy can be preserved from the perspective of optimization.

Some gaps between the theory and numerical approximation need to be studied that are beyond the scope of this work. For example, a rigorous proof for the optimality of the proposed shortest-path estimate is desirable, to offer a truly definitive convergence guarantee. If such a result is unavailable, an error bound could be developed that measures the gap between the true L_{**} and estimate L_{**}^{est} , which is guaranteed to be only a lower bound for L_{**} . A similar error analysis should investigate the effect of using approximations for the inverse Hessian $(U'(\mathbf{X}^*))^{-1}$ and operator norm $\|\cdot\|_2$ in the computation of L_{**}^{est} , as in practice either are likely to be precisely computed. Furthermore, one can extend the proposed framework to local convergence theory for quasi-Newton methods (e.g., [4]) to preserve the corresponding basins of attraction.

Acknowledgments

We thank Florian Potra and Stefan Wild for their discussions and insights on this work. This work was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation’s exascale computing imperative. The material was also based in part on work supported by the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357.

References

1. Alder, B.J., Wainwright, T.E.: Studies in molecular dynamics, I: general method. *The Journal of Chemical Physics* **31**(2), 459–466 (1959)
2. Argyros, I.K., George, S.: Ball convergence comparison between three iterative methods in banach space under hypothese only on the first derivative. *Applied Mathematics and Computation* **266**, 1031–1037 (2015)
3. Asenjo, D., Stevenson, J.D., Wales, D.J., Frenkel, D.: Visualizing basins of attraction for different minimization algorithms. *The Journal of Physical Chemistry B* **117**(42), 12717–12723 (2013)
4. Dembo, R.S., Eisenstat, S.C., Steihaug, T.: Inexact Newton methods. *SIAM Journal on Numerical Analysis* **19**(2), 400–408 (1982)

5. Doye, J.: Table of Lennard-Jones cluster global minima, <http://www-wales.ch.cam.ac.uk/~jon/structures/LJ/tables.150.html>
6. Hoare, M.R.: Structure and Dynamics of Simple Microclusters, pp. 49–135. John Wiley & Sons, Ltd (2007)
7. Hoare, M., Pal, P.: Physical cluster mechanics: Statics and energy surfaces for monatomic systems. *Advances in Physics* **20**(84), 161–196 (1971)
8. Jones, J.E.: On the determination of molecular fields, II: from the equation of state of a gas. *Proc. R. Soc. Lond. A* **106**(738), 463–477 (1924)
9. Kantorovich, L.V.: Functional analysis and applied mathematics. *Uspekhi Matematicheskikh Nauk* **3**(6), 89–185 (1948)
10. Maranas, C.D., Floudas, C.A.: A global optimization approach for Lennard–Jones microclusters. *The Journal of Chemical Physics* **97**(10), 7667–7678 (1992)
11. Northby, J.: Structure and binding of Lennard–Jones clusters: $13 \leq N \leq 147$. *The Journal of Chemical Physics* **87**(10), 6166–6177 (1987)
12. Perez, D., Cubuk, E.D., Waterland, A., Kaxiras, E., Voter, A.F.: Long-time dynamics through parallel trajectory splicing. *Journal of Chemical Theory and Computation* **12**(1), 18–28 (2015)
13. Potra, F.: A superquadratic variant of Newton’s method. *SIAM Journal on Numerical Analysis* **55**(6), 2863–2884 (2017)
14. Rahman, A.: Correlations in the motion of atoms in liquid argon. *Physical Review* **136**(2A), A405 (1964)
15. Sørensen, M.R., Voter, A.F.: Temperature-accelerated dynamics for simulation of infrequent events. *The Journal of Chemical Physics* **112**(21), 9599–9606 (2000)
16. Traub, J.F., Woźniakowski, H.: Convergence and complexity of Newton iteration for operator equations. Technical report, Carnegie Mellon University, Department of Computer Science, Pittsburgh, PA (1977)
17. Uppenbrink, J., Wales, D.J.: Packing schemes for Lennard–Jones clusters of 13 to 150 atoms: minima, transition states and rearrangement mechanisms. *J. Chem. Soc., Faraday Trans.* **87**, 215–222 (1991)
18. Voter, A.F.: Hyperdynamics: Accelerated molecular dynamics of infrequent events. *Physical Review Letters* **78**(20), 3908 (1997)
19. Voter, A.F.: Parallel replica method for dynamics of infrequent events. *Physical Review B* **57**(22), R13985 (1998)
20. Wang, X.: Convergence of Newton’s method and uniqueness of the solution of equations in Banach space. *IMA Journal of Numerical Analysis* **20**(1), 123–134 (2000)
21. Wille, L.T., Vennik, J.: Computational complexity of the ground-state determination of atomic clusters. *Journal of Physics A: Mathematical and General* **18**(8), L419 (1985)
22. Wood, G., Zhang, B.: Estimation of the lipschitz constant of a function. *Journal of Global Optimization* **8**(1), 91–103 (1996)
23. Xiang, Y., Cheng, L., Cai, W., Shao, X.: Structural distribution of lennard-jones clusters containing 562 to 1000 atoms. *The Journal of Physical Chemistry A* **108**(44), 9516–9520 (2004)
24. Yamamoto, T.: Historical developments in convergence analysis for Newton’s and Newton-like methods. In: *Numerical Analysis: Historical Developments in the 20th Century*, pp. 241–263. Elsevier (2001)