Force-based gradient descent method for *ab initio* atomic structure relaxation

Yukuan Hu [1,2] Xingyu Gao [3,*] Yafan Zhao,[4] Xin Liu,[1,2,†] and Haifeng Song[3,‡]

[1]*State Key Laboratory of Scientific and Engineering Computing, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, China*
[2]*University of Chinese Academy of Sciences, Beijing 100049, China*
[3]*Laboratory of Computational Physics, Institute of Applied Physics and Computational Mathematics, Beijing 100088, China*
[4]*CAEP Software Center for High Performance Numerical Simulation, Beijing 100088, China*

Force-based algorithms for *ab initio* atomic structure relaxation, such as conjugate gradient methods, usually get stuck in the line minimization processes along search directions, where expensive *ab initio* calculations are triggered frequently to test trial positions before locating the next iterate. We present a force-based gradient descent method, WANBB, that circumvents the deficiency. At each iteration, WANBB enters the line minimization process with a trial step size capturing the local curvature of the energy surface. The exit is controlled by a nonrestrictive criterion that tends to accept early trials. These two ingredients streamline the line minimization process in WANBB. The numerical simulations on nearly 80 systems with good universality demonstrate the considerable compression of WANBB on the cost for the unaccepted trials compared with conjugate gradient methods. We also observe across the board significant and universal speedups as well as the superior robustness of WANBB over several widely used methods. The latter point is theoretically established. The implementation of WANBB is pretty simple, in that no *a priori* physical knowledge is required and only three parameters are present without tuning.

Atomic structure relaxation determines the ground-state atomic configuration by searching for the local minimum in the energy landscape. It makes the foundation and potential performance bottleneck for the search of global structure minimum [1–3] and high-throughput calculations in material design [4–8]. Conjugate gradient methods (CG) [9], the direct inversion in the iterative subspace (DIIS) [10], and limited-memory Broyden-Fletcher-Goldfarb-Shanno quasi-Newton methods (LBFGS) [11] are three widely used methods to relax the atomic structure following the density functional theory (DFT) calculation. While DIIS and LBFGS could be slightly faster in some cases, CG is more stable in the situations where the initial atomic configuration is far from equilibria. For large systems, hundreds of steps may be required to converge to the typical atomic force tolerance of 0.01 eV/Å. So, it will be tremendously helpful even if the *ab initio* atomic relaxation could be sped up to some degree.

One reason for the slow convergence of the local minimization goes to the narrowly curved energy valley, which prevents the efficient execution of the conventional CG. Thus, various preconditioners have been designed for improving search directions via approximating the Hessian matrices. Zhao *et al.* [12] notice that the Hessian matrix of large quantum dots can be approximated by that of motifs comprising merely a few atoms. Some other works construct general force fields

based on *a priori* physical knowledge. Surrogate potentials are used to either directly approximate the Hessian matrix [13–15], possibly with on-the-flight fittings [16,17], or generate the "preconditioned force" with the implicit inversion of the Hessian matrix [18].

Another reason is the inefficient execution of the line minimization (LM) to which not much attention has been paid in the previous study. Actually, gradient descent method and CG are two typical LM-based methods [9,19,20]. A flowchart of the general LM-based method for the *ab initio* atomic relaxation is depicted in Fig. 1. The *ab initio* LM process, marked out by the shaded box in Fig. 1, is carried out along search directions such as steepest descent or CG. During this process, one begins with an initial trial step size and then repeatedly invokes an LM algorithm to generate new trials until certain LM criterion is met. Involving moving atoms, each trial step is followed by *ab initio* calculation, e.g., solving the Kohn-Sham (KS) equation. But, only the last accepted trial step triggers the update of search direction. Taking the conventional CG, for instance, we test nearly 80 systems and find that each update of search direction is followed by 2.4 calls of LM algorithms on average. In other words, nearly 60% of the total computational overhead is paid for the unaccepted trial steps in the LM process. Moreover, preconditioning with the (fitted) classical force fields does not necessarily waive the cost in the LM process [17]. As a result, the computational cost on the LM process is far more than that on simply determining a search direction if the whole procedure gets frequently trapped in the inner loop.

[*]Corresponding author: gao_xingyu@iapcm.ac.cn
[†]Corresponding author: liuxin@lsec.cc.ac.cn
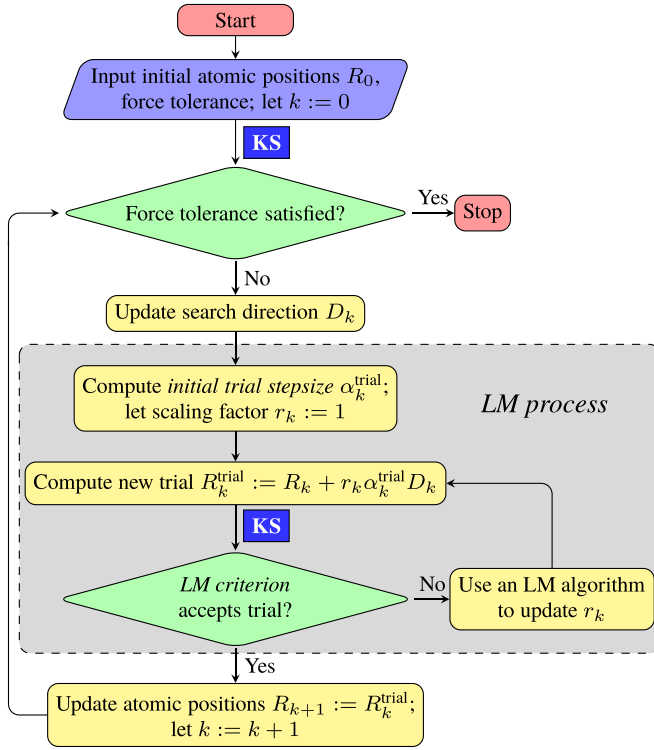[‡]Corresponding author: song_haifeng@iapcm.ac.cn

FIG. 1. A flowchart of the general LM-based method, where $N \in \mathbb{N}$ refers to the number of atoms, $R_k \in \mathbb{R}^{3 \times N}$ stands for the Cartesian atomic coordinates, $F_k \in \mathbb{R}^{3 \times N}$ denotes the atomic forces applied at R_k, and $D_k \in \mathbb{R}^{3 \times N}$ is the search direction. The white bold text "**KS**" refers to solving the corresponding Kohn-Sham equations to obtain new energies and forces whenever the atomic positions are updated. The LM process is marked out by the shaded box.

In this work, we present an LM-based gradient descent method, called WANBB, which avoids trapping in the LM process during atomic structure relaxation. We achieve this by (i) calculating an initial trial step size that captures the local curvature of potential energy surface (PES) and (ii) devising a nonrestrictive LM criterion that tends to accept early trials. Nearly 80 systems have been tested, including organic molecules, metallic systems, semiconductors, surface-molecule adsorption systems, ABX_3 perovskites, etc. The set of benchmark tests favors the universality of everyday *ab initio* atomic relaxation. The unaccepted trial steps in the LM process account for only about 1% of the total solved KS equations on average in WANBB, compared with nearly 60% in CG. The considerable compression of the LM process leads to a prominent saving on running time. The average speedup factors of WANBB over CG, DIIS, and LBFGS are about 1.5, 1.2, and 1.2, respectively. The robustness of the atomic relaxation method, i.e., convergence to equilibria regardless of initial configurations, is also investigated. While CG, DIIS, and LBFGS fail on some systems, WANBB manages to converge across the benchmark. This robustness is theoretically established and helpful for the cases with local lattice distortion or reconstruction. Last but not least, WANBB is pretty simple. Free of *a priori* physical knowledge, it can work with only three parameters present but no tuning required.

In what follows, we are ready to deliver the algorithmic development. To facilitate narration, we collect some notations beforehand. We denote by $E(R) \in \mathbb{R}$ and $F(R) \in \mathbb{R}^{3 \times N}$, respectively, the potential energy and atomic forces evaluated at R. When describing algorithms, we use subscripts for abbreviation: e.g., both E_k and F_k are evaluated at R_k. We denote by $\mathbf{I}$ the identity mapping from $\mathbb{R}^{3 \times N}$ to $\mathbb{R}^{3 \times N}$. The operator $\langle \cdot, \cdot \rangle$ represents the Frobenius inner product of two $3 \times N$ matrices

$$\langle M, N \rangle := \mathrm{Tr}(M^\top N) \quad \forall M, N \in \mathbb{R}^{3 \times N},$$

whereas $\| \cdot \|_F$ yields the Frobenius norm as $\|M\|_F := \sqrt{\langle M, M \rangle}$ for any $M \in \mathbb{R}^{3 \times N}$. We further denote by $\|M\|_{2,\infty}$ the $\ell_{2,\infty}$ norm of matrix M, namely, the maximum among the ℓ_2 norms of columns in M. In particular, $\|F\|_{2,\infty}$ gives the maximum atomic force. Throughout iterations, we take the atomic forces F_k as the search direction D_k.

Essentially, given $D_k = F_k$ and the initial trial step size $\alpha_k^{\mathrm{trial}} > 0$, the LM process aims at inexactly solving the one-dimensional minimization

$$\min_{r > 0} \tilde{E}_k(r) := E\big(R_k + r \cdot \alpha_k^{\mathrm{trial}} F_k\big). \tag{1}$$

The LM algorithm is exactly some univariate optimization or root-finding method, initialized from $r_k = 1$. The LM criterion serves as the stopping rule for the process. From this perspective, the whole procedure can fall into the trap of the LM process owing to two factors. One is the poorly chosen initial trial step size, which degrades the starting state of the LM process. The other is the restrictive LM criterion, which acts as a stringent stopping rule. The ideal situation is that the procedure always passes the test of the LM criterion with just the original initial step size (namely, $r_k = 1$). Intuitively, we require both good initial step sizes and a loose but safe criterion to achieve this target. Here, "safe" means that the convergence property is maintained with the LM criterion.

At first sight, a good initial step size should grasp, as much as possible, the curvature of PES around the current configuration. As early as 1988, Barzilai and Borwein [21] propose the following renowned Barzilai-Borwein (BB) step sizes: for any $k \geqslant 1$,

$$\alpha_k^{\mathrm{BB1}} := \frac{\|S_{k-1}\|_F^2}{\langle S_{k-1}, Y_{k-1} \rangle}, \quad \alpha_k^{\mathrm{BB2}} := \frac{\langle S_{k-1}, Y_{k-1} \rangle}{\|Y_{k-1}\|_F^2},$$

where $S_{k-1} := R_k - R_{k-1}$ and $Y_{k-1} := F_{k-1} - F_k$ represent the difference of successive atomic positions and forces, respectively. Remarkably, these two step sizes satisfy

$$\begin{aligned} \alpha_k^{\mathrm{BB1}} &= \arg\min_\alpha \|\alpha^{-1} S_{k-1} - Y_{k-1}\|_F, \\ \alpha_k^{\mathrm{BB2}} &= \arg\min_\alpha \|\alpha Y_{k-1} - S_{k-1}\|_F. \end{aligned} \tag{2}$$

In other words, $(\alpha_k^{\mathrm{BB1}})^{-1} \mathbf{I}$ and $(\alpha_k^{\mathrm{BB2}})^{-1} \mathbf{I}$ approximate the Hessian around R_k along the last search direction. They capture the local curvature of PES with merely two $3 \times N$ matrices. The alternating BB (ABB) step size [22] is presented as follows:

$$\alpha_k^{\mathrm{ABB}} := \begin{cases} \alpha_k^{\mathrm{BB1}}, & \text{if } \mathrm{mod}(k, 2) = 1 \\ \alpha_k^{\mathrm{BB2}}, & \text{if } \mathrm{mod}(k, 2) = 0 \end{cases} \tag{3}$$
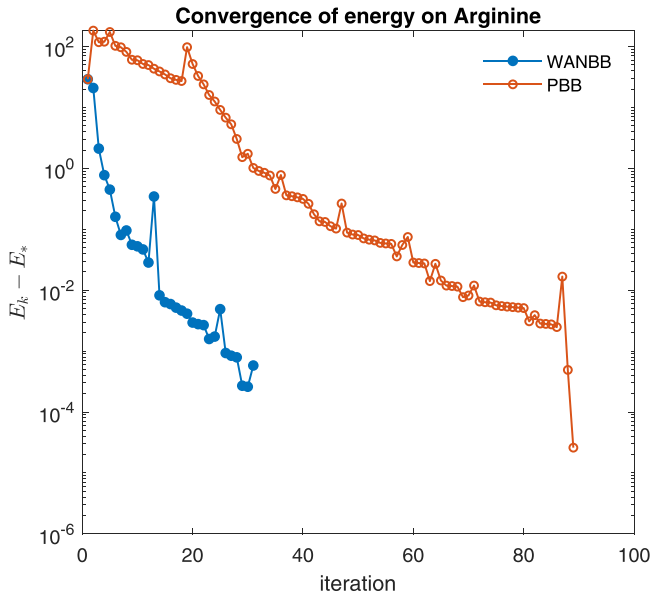
FIG. 2. The convergence of energy on arginine by gradient descent with the pure ABB step sizes (PBB) and our method (WANBB). E^* refers to the converged energy. The blue line with solid markers stands for the convergence of WANBB, while the red line with hollow markers for PBB.

which is shown to overwhelm BB1 and BB2 in many contexts (e.g., [22,23]). It is not difficult to prove that the implementation of (A)BB step sizes requires the same computational and storage complexity as computing the conjugate parameter in CG.

Despite the simplicity, the (A)BB step sizes turn out to be significant benefits in solving various optimization problems. The most related application goes to solving KS equations as orthogonality constrained optimization problems, where the (A)BB step sizes play as important local acceleration components (see, e.g., [23–27]). However, the application of the (A)BB step sizes on the atomic structure relaxation problem remains to be investigated. In this work, we take the (A)BB step sizes as the initial trial step sizes to exploit their merits.

Without LM, the energy values produced by gradient descent methods with the pure (A)BB step sizes (PBB) do not necessarily decrease monotonically, which may result in fluctuation or even divergence when starting the algorithms far from equilibria. An illustrative example is provided in Fig. 2, from which one can notice that the fluctuation of PBB leads to a lengthy convergence process. Nevertheless, the monotone LM criterion, adopted by the conventional CG, can ruin the merits of the (A)BB step sizes since frequent calls of LM algorithms are entailed to ensure the monotonic decrease of the energy and the local curvature of PES may be lost. Therefore, it is more sensible to design a *nonmonotone* LM (NLM) criterion, accepting the original (A)BB step sizes in most cases.

The existing NLM criteria primarily fall into two categories: the "max" type (MNLM) [28] and the "average" type (ANLM) [29]. In the $(k + 1)$th iteration, the MNLM criterion asks the LM algorithm to return a factor $r_k > 0$ such

that

$$E\left(R_k + r_k \alpha_k^{\text{trial}} F_k\right) \leqslant \bar{E}_k - c r_k \alpha_k^{\text{trial}} \|F_k\|_{\text{F}}^2,$$

where $\bar{E}_k := \max_{i=k-m(k)+1}^{k}\{E_i\}$, $m(k) \in [0, M]$, both c and M are given positive constants, i.e., there is a sufficient reduction with respect to (w.r.t.) the largest one of the past $m(k)$ energies. The ANLM criterion requires r_k to fulfill

$$E\left(R_k + r_k \alpha_k^{\text{trial}} F_k\right) \leqslant C_k - c r_k \alpha_k^{\text{trial}} \|F_k\|_{\text{F}}^2, \quad (4)$$

$$\left\langle F\left(R_k + r_k \alpha_k^{\text{trial}} F_k\right), F_k\right\rangle \leqslant \sigma \|F_k\|_{\text{F}}^2, \quad (5)$$

where $\sigma \in (0, 1)$, the monitoring sequence $\{C_k\}$ is updated through

$$C_{k+1} := \frac{\eta_k Q_k C_k + E_{k+1}}{\eta_k Q_k + 1}, \quad Q_{k+1} := \eta_k Q_k + 1 \quad (6)$$

with $C_0 := E_0$, $Q_0 := 1$, and $\eta_k \in [0, 1]$. Roughly in the ANLM criterion, a sufficient reduction w.r.t. C_k is obligatory. It is straightforward to deduce that C_k is a weighted average of the past energies $\{E_j\}_{j=0}^{k-1}$ and the current E_k. The construction of the monitoring sequence is completely free of *a priori* physical knowledge. Clearly, entailing reduction w.r.t. $\bar{E}_k$ and C_k, respectively, instead of E_k, both MNLM and ANLM criteria are less restrictive than the monotone criterion adopted by the conventional CG because $\bar{E}_k, C_k \geqslant E_k$. Moreover, as noted in [29], the MNLM criterion is sensitive to the choice of M in some contexts, and is often outperformed by the ANLM criterion.

However, preliminary simulations show that the ANLM criterion is not loose enough for the (A)BB step sizes during the atomic relaxation. We take the gradient descent direction armed with the ABB step sizes and ANLM criterion to relax the silicene surface with one adsorbed H_2O molecule. From the short relaxation history in Fig. 3, one can observe that the monitoring sequence $\{C_k\}$ (the blue dashed line with square markers) decreases fast, stays close to the energies $\{E_k\}$, and does not tolerate somewhat large increments. As a result, the combination "ABB + ANLM" solves 3 KS equations for the unaccepted trials during the short history, which are marked out by the black arrows in Fig. 3. For example, the initial trial step at the eighth iteration gives an energy value of $-99.92\,\text{eV}$, which is higher than $C_7 = -100.03\,\text{eV}$ and thus is rejected by the ANLM criterion.

To this end, we devise a reweighted ANLM (WANLM) criterion, which differs from the ANLM criterion in the weights assigned to the past and current energies. Specifically, we denote the monitoring sequence in WANLM by $\{B_k\}$ and update it via

$$B_{k+1} := \frac{B_k + \mu_k P_k E_{k+1}}{1 + \mu_k P_k}, \quad P_{k+1} := 1 + \mu_k P_k, \quad (7)$$

with $B_0 := E_0$, $P_0 := 1$, and $\mu_k \in [0, 1]$. Then WANLM needs r_k to satisfy both Eqs. (4) and (5) with C_k, Q_k, and η_k replaced by B_k, P_k, and μ_k, respectively.

The looseness of the WANLM criterion can be understood via the comparison with the ANLM criterion. Due to Eq. (6) and $\eta_0 \in [0, 1]$, the coefficient of E_1 is $1/(1 + \eta_0) \geqslant 0.5$, indicating a large portion of E_1 in C_1 and a small gap $C_1 - E_1$. Since the monitoring sequence $\{C_k\}$ decreases monotonically (see the Appendix for a rigorous proof), it would then stay
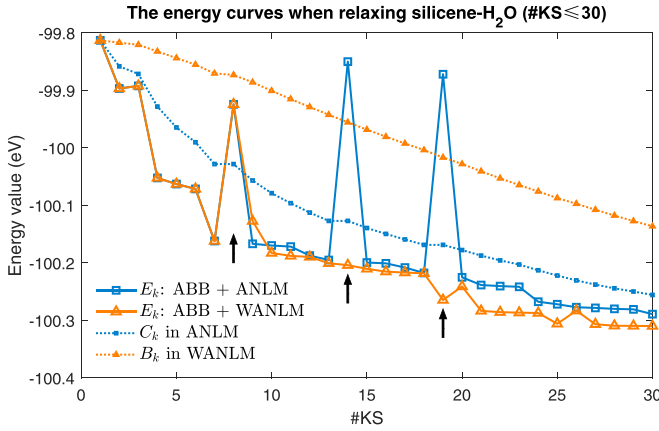
FIG. 3. The energy curves when relaxing the silicene surface with one adsorbed H_2O molecule using gradient descent method equipped with the ABB step sizes and ANLM [$\eta_k \equiv 0.85$ (this value is recommended by [29]] or WANLM ($\mu_k \equiv 0.05$) criterion. "#KS" refers to the number of solved KS equations. We plot out the history with #KS $\leqslant 30$. The blue solid line with square markers and orange line with triangle markers stand for energies produced by "ABB + ANLM" and "ABB + WANLM," respectively. The blue dashed line with square markers and orange line with triangle markers represent $\{C_k\}$ in ANLM and $\{B_k\}$ in WANLM, respectively. The black arrows indicate the *ab initio* calculations dedicated to the trials rejected by the ANLM criterion.

close to the energy values $\{E_k\}$, rendering the ANLM criterion stringent for the (A)BB step sizes as reflected in Fig. 3. In contrast, from Eq. (7) and $\mu_0 \in [0, 1]$, the coefficient of E_1 is $\mu_0/(1 + \mu_0) \leqslant 0.5$. Therefore, the monitoring sequence $\{B_k\}$ in the WANLM criterion is relatively far away from the energies $\{E_k\}$ at the beginning and gradually approaches $\{E_k\}$, leaving sufficient room to accept the (A)BB step sizes. As shown in Fig. 3, the first 7 energy values obtained by "ABB + ANLM" and "ABB + WANLM" are the same. But, one can find a much larger distance from $\{E_k\}$ to $\{B_k\}$ than $\{C_k\}$, which enables the WANLM criterion to accept the 8th initial trial step without invoking extra *ab initio*

calculations and retain the merits of the ABB step sizes. All the subsequent iterations of "ABB + WANLM" benefit from this characteristic and progress to lower-energy values faster. More comprehensive comparison between the ANLM and WANLM criteria can be found later in Fig. 5. Regarding safety, it can be verified that the sufficient reduction in terms of the atomic force norm holds for the monitoring sequence $\{B_k\}$. One can then establish the convergence to equilibria, regardless of initial configurations, after noticing the lower boundedness of $\{B_k\}$. For more details, please refer to the Appendix.

The gradient descent method combining the (A)BB step sizes with the WANLM criterion is referred to as WANBB. We set the initial trial step size to be 4.8×10^{-2} Å^2/eV at the first iteration [30], and take the alternating strategy described in Eq. (3) (i.e., $\alpha_k^{\text{trial}} = \alpha_k^{\text{ABB}}$) afterwards because it leads to better performance than BB1 and BB2 in our simulations. Before invoking an LM algorithm, we take absolute value in case that $\langle S_k, Y_k \rangle < 0$ and make truncation for safeguard:

$$\alpha_k^{\text{ABB}} := \min \left\{ \left| \alpha_k^{\text{ABB}} \right|, \max \left\{ -\log_{10}(\|F_k\|_{2,\infty}), 1.0 \right\} \right\}. \quad (8)$$

As is shown in the Appendix, the convergence of WANBB to equilibria can be attributed to both Eqs. (4) and (5). For the sake of implementation, it is more efficient to fulfill only Eq. (4). In WANLM, we set $\mu_k \equiv 0.05$, $c = 10^{-4}$ [31]. The LM algorithm works by forming quadratic or cubic approximation m_k to $\tilde{E}_k$ [defined in Eq. (1)]. We explain this by taking quadratic approximation as example. Suppose the last scaling factor $\bar{r} > 0$ is rejected by the WANLM criterion. Since F_k is a descent direction, we know that one acceptable scaling factor must exist in $[0, \bar{r}]$. Note that already we have energies $\tilde{E}_k(0)$, $\tilde{E}_k(\bar{r})$ and forces F_k. The derivative of $\tilde{E}_k$ at $r = 0$ is then $\tilde{E}_k'(0) = \alpha_k^{\text{trial}} \|F_k\|_{\text{F}}^2$. Based upon these, we can uniquely determine a quadratic function m_k which satisfies $m_k(0) = \tilde{E}_k(0)$, $m_k(\bar{r}) = \tilde{E}_k(\bar{r})$, and $m_k'(0) = \tilde{E}_k'(0)$. The quadratic function m_k can be regarded as a nice local approximation to $\tilde{E}_k$. It is natural to take its minimizer in $[0, \bar{r}]$, which is easy to compute, and generate the next trial configuration. Analogously, we are able to construct cubic
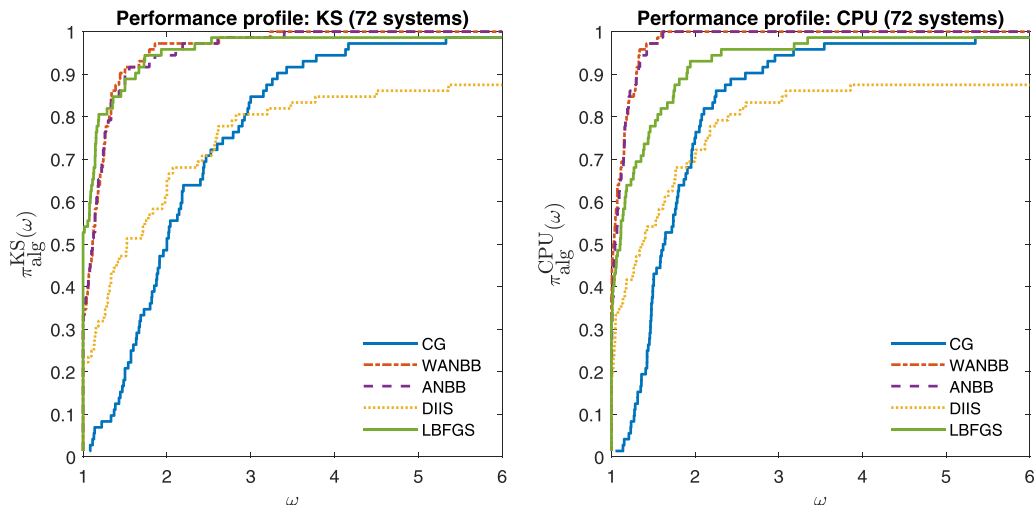

FIG. 4. The #KS (left) and CPU (right) performance profiles of CG, WANBB, ANBB, DIIS, and LBFGS.
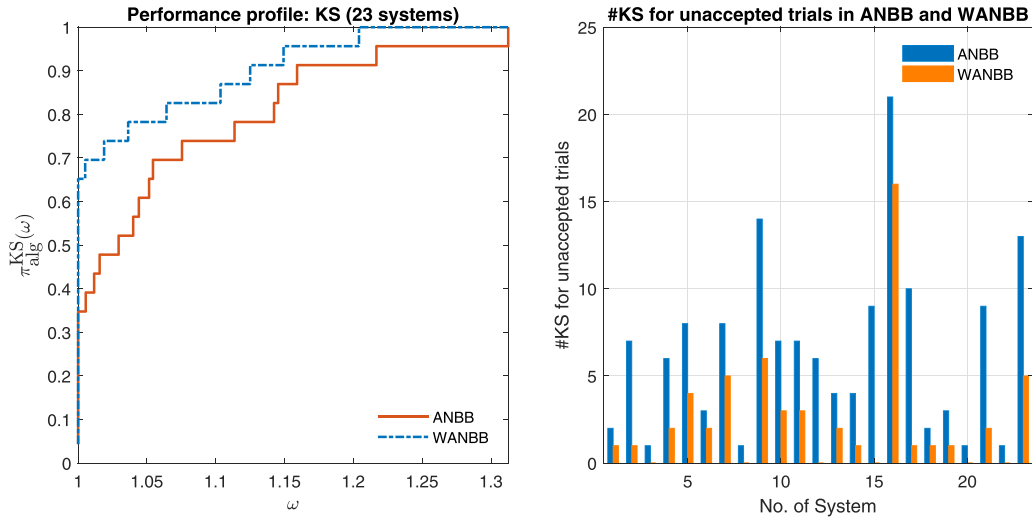
FIG. 5. Left: the performance profile of ANBB and WANBB on 23 systems where their #KS differ; right: #KS for unaccepted trials in ANBB and WANBB when relaxing the 23 systems.

models once more information about PES is available. For details, please see [19, Chap. 3].

We have implemented WANBB in the in-house plane-wave code CESSP [32–35]. The exchange-correlation energy is described by the generalized gradient approximation [36]. Electron-ion interactions are treated with the projector augmented-wave (PAW) potentials based on the open-source ABINIT Jollet-Torrent-Holzwarth data set library in the PAW-XML format [37]. The total energies are calculated using the Monkhorst-Pack mesh [38] with the k-mesh spacing of 0.2 Å^{-1}. The typical plane-wave energy cutoffs are around 500–600 eV. The KS equations corresponding to the updated atomic positions are solved by the preconditioned self-consistent field (SCF) iteration [34].

In addition to WANBB, we test the performances of other relaxation methods as well, including the conventional CG [9,19], LBFGS [11], and DIIS [10]. By the way, DIIS is the so-called quasi-Newton method in some software such as VASP [39,40], which can be derived from Broyden's method [41]. The implementations of CG and DIIS follow those in [11,19], and LBFGS is realized by the transition state tool VTST [42]. CG employs Brent's method [43] as the LM algorithm. We also conduct numerical comparison between WANBB and gradient descent combining ABB step sizes with the ANLM criterion (ANBB), where $\eta_k \equiv 0.85$ [29], $c = 10^{-4}$.

The benchmark for performance test contains nearly 80 systems from various categories, including organic molecules, metallic systems, semiconductors, surface-molecule adsorption systems, and ABX_3 perovskites. The number of atoms ranges from 2 to 429. Some of them are available online, e.g., in Materials Project [44] and Organic Materials Database [45]. In the beginning, these systems are at their ideal crystal positions for defects, heterostructures, and substitutional alloys, or simply place a molecule on top of a surface. These are the typical initial configurations used by researchers in their simulations. More information on the benchmark can be found in the Supplemental Material [46]. The atomic structure relaxation is terminated if either the maximum atomic force falls below 0.01 eV/Å or the number of the solved KS equa-

tions (#KS) arrives at 1000. The convergence criterion for the SCF iteration is 10^{-5} eV.

We evaluate the performances of the relaxation methods in two ways. One is the number of the solved KS equations (#KS) needed to reach the 0.01 eV/Å force tolerance. The other is the running time (CPU) of the atomic structure relaxation. We adopt Dolan and Moré's performance profile [47] for an overall comparison among CG, WANBB, ANBB, DIIS, and LBFGS. To this end, for any algorithm indicated by alg $\in$ {CG, WANBB, ANBB, DIIS, LBFGS}, we denote its #KS and CPU for relaxing the nth system by #KS$_{n,\mathrm{alg}}$, CPU$_{n,\mathrm{alg}}$, respectively [48], define ratios

$$r_{n,\mathrm{alg}}^{\mathrm{KS}} := \frac{\#\mathrm{KS}_{n,\mathrm{alg}}}{\min_{\mathrm{alg}} \#\mathrm{KS}_{n,\mathrm{alg}}}, \quad r_{n,\mathrm{alg}}^{\mathrm{CPU}} := \frac{\mathrm{CPU}_{n,\mathrm{alg}}}{\min_{\mathrm{alg}} \mathrm{CPU}_{n,\mathrm{alg}}},$$

and for any $\omega \geqslant 0$, let

$$\pi_{\mathrm{alg}}^{\mathrm{KS}}(\omega) := \frac{\text{the number of systems satisfying } r_{n,\mathrm{alg}}^{\mathrm{KS}} \leqslant \omega}{\text{the number of systems}},$$

$$\pi_{\mathrm{alg}}^{\mathrm{CPU}}(\omega) := \frac{\text{the number of systems satisfying } r_{n,\mathrm{alg}}^{\mathrm{CPU}} \leqslant \omega}{\text{the number of systems}}.$$

The quantities $\pi_{\mathrm{alg}}^{\mathrm{KS}}(\omega)$ and $\pi_{\mathrm{alg}}^{\mathrm{CPU}}(\omega)$ estimate, respectively, the probabilities of $r_{n,\mathrm{alg}}^{\mathrm{KS}} \leqslant \omega$ and $r_{n,\mathrm{alg}}^{\mathrm{CPU}} \leqslant \omega$ over the benchmark. For example, $\pi_{\mathrm{WANBB}}^{\mathrm{CPU}}(1)$ yields the portion of the systems where WANBB is the fastest. And $\pi_{\mathrm{WANBB}}^{\mathrm{CPU}}(2)$ means the estimated probability that, for one system, the running time of WANBB is not more than twice the shortest running time needed by all the five algorithms. In terms of the #KS, $\pi_{\mathrm{alg}}^{\mathrm{KS}}(\omega)$ can be interpreted in a similar way. We plot $\pi_{\mathrm{alg}}^{\mathrm{KS}}(\omega)$ and $\pi_{\mathrm{alg}}^{\mathrm{CPU}}(\omega)$ as a function of ω for each method in Fig. 4 [49]. Basically, the larger the area under the curve, the better the overall performance of the method. It is easy to see that the overall performance of WANBB ranks the best. More specifically, one can tell from the Supplementa Material Tables I and II [46] that the average CPU speedup factors of WANBB over CG, DIIS, and LBFGS are 1.51, 1.21, and 1.16, respectively.
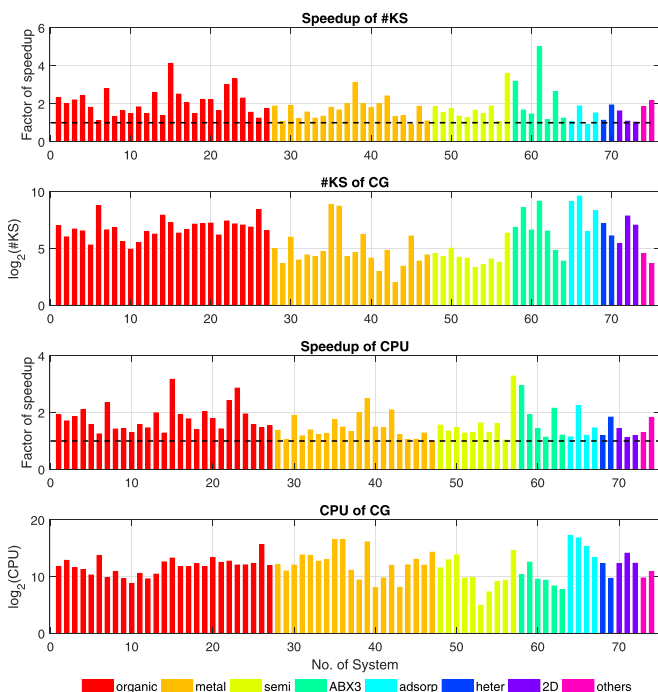
FIG. 6. The speedup factors of #KS and the #KS of CG (the upper two), the speedup factors of CPU and the CPU of CG (the lower two). The 1.0 baselines are marked out by bold black dashed lines. The system where CG fails is not included.

FIG. 7. The frequency distributions of the LM step percentages of WANBB and CG for unaccepted trials. The system where CG fails is not included.

The comparison between WANBB and ANBB is not evident from Fig. 4 because the ABB step sizes are highly efficient so that the NLM criteria are inactive in most cases. To this end, we pick out systems on which the #KS of these two methods differ (23 systems in total), and conduct further comparison between ANBB and WANBB on these systems (see Fig. 5). One can observe that WANBB is superior overall to ANBB and spends less computational overhead for unaccepted trials. This advantage should be credited to the nonrestrictive WANLM criterion.

As shown in Fig. 4, though sometimes requiring solving more KS equations than LBFGS, WANBB turns out to be faster than LBFGS in most tested systems. LBFGS may make larger atomic displacements due to constant step size, which results in underestimating the quality of initial wave functions and slowing down the SCF iteration. Later, we will explain it in more detail using the relaxation of the PuO_2 surface with one adsorbed H_2O molecule (see Fig. 11 later).

In addition to the efficiency, the robustness is also worth the whistle. As shown in the Supplemental Material Tables I and II [46], WANBB converges in all the benchmark systems, while CG fails in one case due to the breakdown of Brent's method, DIIS and LBFGS fail in eight cases and one case, respectively, due to divergence. This kind of robustness is theoretically guaranteed by the convergence property of WANBB proved in the Appendix.

Taking the performance of CG as the baseline, we illustrate the acceleration obtained by WANBB for each system, as depicted in Fig. 6. We have an average CPU speedup factor around 1.51, and a factor above 2.0 on about 20% of systems. According to the Supplemental Material Tables I and
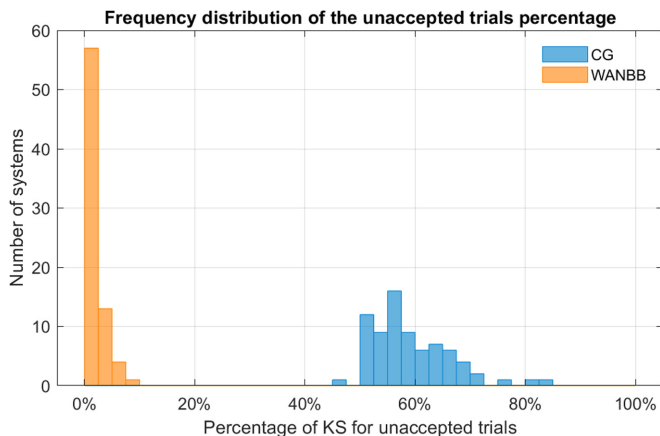
II [46], CG takes on average 117.68 calls of solving the KS equations to reach the force tolerance. Thus, a speedup factor of 1.5–2 can mean a significant saving of computational overhead. This saving is mainly attributed to the enhanced *ab initio* LM process of WANBB since the distribution of #KS speedup is consistent with that of CPU speedup in Fig. 6. Moreover, we also count the percentage of the LM steps invoked by WANBB and CG for unaccepted trials and draw the frequency distribution over the benchmark in Fig. 7. The average LM step percentages for unaccepted trials are 59.09% and 1.47% for CG and WANBB, respectively. It is obvious from these statistics that WANBB escapes from the LM process quickly, usually without invoking LM algorithms at all.

The following discussion is concentrated on the performance comparison between CG and WANBB for each category. In Fig. 8, we plot the mean speedup factors by category and the number of atoms, respectively. All description below is supported by Tables I and II in the Supplemental Material [46]. Whenever speaking of "speedup" below, we are talking about the CPU speedup if not specified.

Organic molecules. This class of systems includes various alkanes and amino acids. The amino acids contain common organic atomic types, e.g., oxygen, nitrogen, sulfur, etc. The number of atoms in this class ranges from 26 to 106. Most of the initial configurations are available in Organic Materials Database [45]. Compared with CG, WANBB achieves a mean speedup factor of 1.81 for this class and speedup factors greater than 2 on over 20% of systems. In the case of glutamate, the acceleration factor even reaches 3.18, which means a savings of nearly 70% computational overhead.

Metallic systems. We consider Ag_{16} cluster, Cu(111) surface with adsorbed organic molecules, and some alloys including Ni_3Al-based intermetallic compound as well as high-entropy alloys (HEAs) such as FeCrNiCoAl. The speedup factors on these systems w.r.t. CG are mostly around 1.5–3. In particular, for Cu(111) surface with adsorbed organic molecules, CG needs to solve over 400 KS equations for atomic relaxations, while WANBB cuts the required #KS roughly in half. It is also worth mentioning that WANBB can efficiently treat HEAs with potential strong local lattice
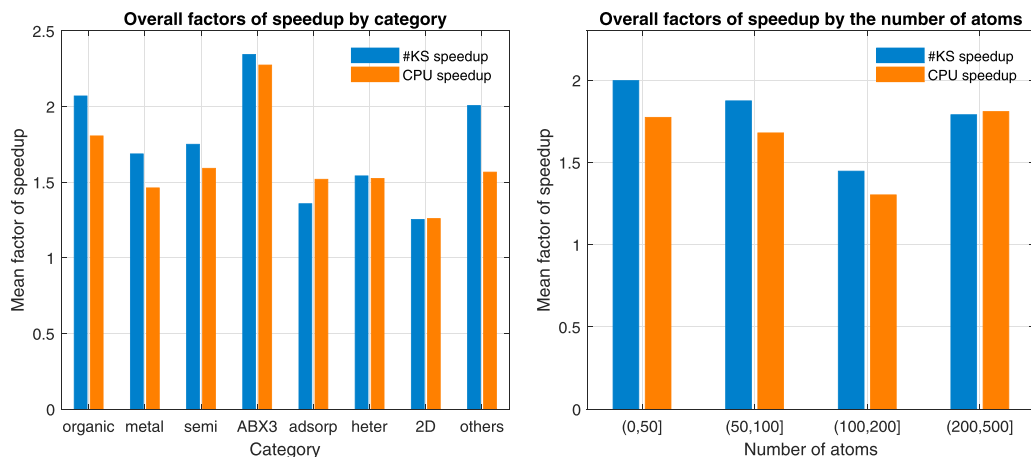
FIG. 8. Mean factors of speedup by category (left) and the number of atoms (right). The system where CG fails is not included.

distortion [50]. Actually, local lattice distortion is an essential factor in studying HEAs [51], and can be described by the supercell method with *ab initio* atomic structure relaxation. In Fig. 9, the local lattice distortion obtained by relaxing the body-centered-cubic supercell of FeCrNiCoAl is illustrated by the smearing radial distribution function (RDF), in sharp contrast to the peaks at coordination shells in the case of the ideal lattice. On average in our benchmark, WANBB achieves a speedup of 1.54 over CG for HEAs.

Semiconductors and heterojunction systems. We test semiconductors (CdSe, GaAs, Si, etc.) and heterojunction systems (GaAs-InAs). Moreover, we consider semiconductors with defects by constructing the bulk GaAs and Si supercells with some vacancies. In most of these systems, the acceleration factors are above 1.5, and sometimes exceed 3. We also investigate the performances of WANBB compared with CG on Si supercells with increasing system sizes (see Fig. 10). As the system size increases, WANBB is more stable in terms of #KS, and its CPU advantage becomes more prominent.
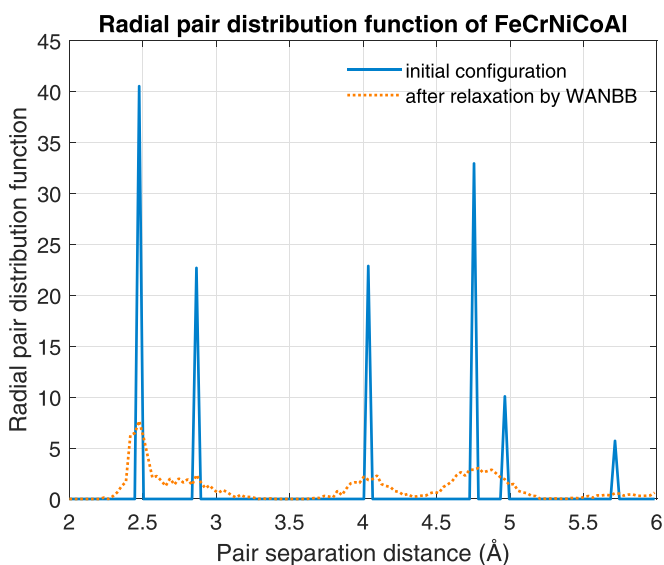

FIG. 9. RDF of FeCrNiCoAl with 6-Å cutoff.

ABX$_3$ perovskite systems. This type of system is fairly popular in recent solar cell studies. We consider the cases where $A = CH_3NH_3$, $CH(NH_2)_2$, or Cs, $B = Pb$ or Sn, $X = Br$, Cl, or I. Their initial configurations come from Materials Project [44]. Due to the existence of the organic molecules and their relative weak interaction with the inorganic framework, the atomic relaxations of such systems are quite challenging. Taking $CH_3NH_3SnI_3$, for instance, CG solves 581 KS equations, while WANBB entails only about one fifth of the expenditure.

Surface-molecule adsorption systems. Several surface-molecule adsorption systems have been tested, including a two-dimensional (2D) silicene surface with one adsorbed H_2O molecule (total 21 atoms), a 2D MoS_2 surface with one adsorbed NH_3 molecule (total 16 atoms), 2D PbS(001) and (111) surfaces with one adsorbed oleic acid molecule [(001) total 254 atoms, (111) total 429 atoms], and a 2D PuO_2 surface with one adsorbed H_2O molecule (total 195 atoms). The results show that CG often undergoes quite a long way until convergence in this category, and even fails in one case due to the breakdown of Brent's method. On the contrary, WANBB converges in all cases and usually costs less time; it is even twice as fast as CG on PbS(001) surface with one adsorbed oleic acid molecule. One interesting observation is that WANBB occasionally takes equal or a few more #KS than CG, even though it is less time costly. This is also observed when comparing WANBB and LBFGS. To explain, we plot the cumulative SCF iteration numbers of CG, WANBB, and LBFGS when relaxing the PuO_2 surface with one adsorbed H_2O molecule (see Fig. 11). One could observe that CG and LBFGS often consume more SCF iterations per KS equation than WANBB. In fact, CG and LBFGS make larger atomic displacements, resulting in degrading initial wave functions and slowing down the SCF iteration. On the other hand, Fig. 11 also reveals that CG and LBFGS produce more efficient search directions, yielding fewer iteration numbers. This weakness, however, does not prevent WANBB from being a more favorable alternative in the context of *ab initio* simulations because the cut down on the LM process obviously brings much more benefits.

2D and other systems. The 2D materials in our benchmark include C_{96}, Si_{32}, and germanene on hexagonal boron-nitride. The average speedup factors of WANBB over CG on these
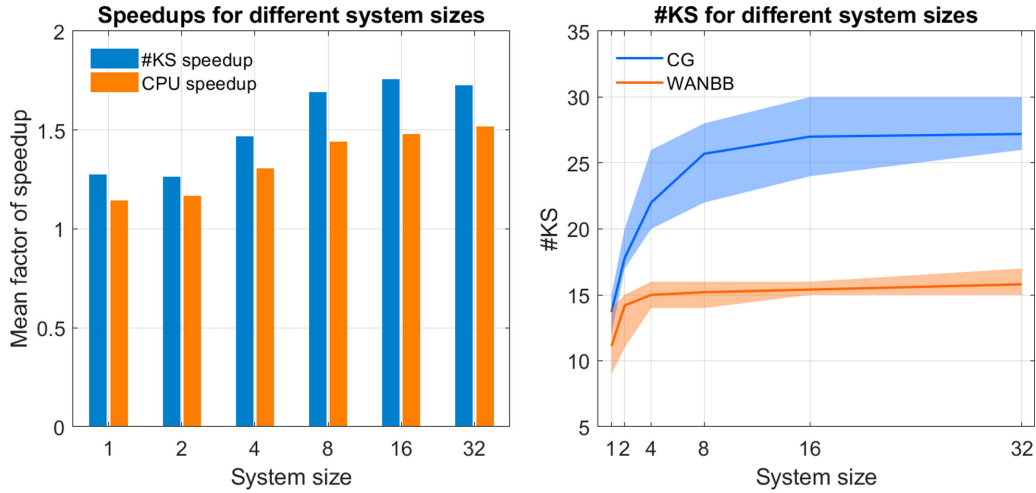
FIG. 10. Left: the average #KS and CPU speedup factors of WANBB over CG on $N \times 1 \times 1$ Si supercells containing from 8 to 256 atoms, where each size has 10 random samples and the shift on each atom away from equilibrium is no larger than 0.1 Å. Right: the #KS of WANBB and CG on $N \times 1 \times 1$ Si supercells, where the solid lines represent the average values, and the upper and lower boundaries of the shaded area stand for the maximum and minimum #KS, respectively.

systems are 1.26. We also include systems of other types, e.g., cerium-iron oxide and potassium chloride. The former one is of practical usage for NO_x reduction. WANBB reduces the numbers of solved KS equations by half on both of them.

We recognize that the inefficiency of the LM process is an important obstacle to the high-performance execution of *ab initio* atomic structure relaxation. With the most widely used CG, nearly 60% of the computational overhead on average is attributed to the unaccepted trials in the LM process. To this end, we propose a force-based gradient descent method called WANBB, where the initial trial (A)BB step sizes grasp well the local curvature of PES and the nonrestrictive WANLM criterion always accepts the initial trials without calling extra *ab initio* calculations. The robustness of WANBB, i.e., the convergence to equilibria regardless of initial configurations, is theoretically guaranteed. The numerical simulations on the benchmark containing nearly 80 various systems reveal the average 1.47% computational cost for unaccepted trials in

WANBB and demonstrate its universal speedups as well as superior robustness over CG, DIIS, and LBFGS.

The speedup of WANBB over other methods should mainly be credited to the usage of both (A)BB trial step sizes for capturing local curvatures of PES and the WANLM criterion for stabilizing the convergence. It is also worth mentioning that the (A)BB step sizes are well suited for steepest descent direction due to resemblance to Newton iteration in the sense of Eq. (2).

Our work concentrates on the LM process. Nevertheless, efficient search directions, such as conjugate gradient, together with subtly designed LM can be another great boon. The choices of initial trial step sizes and LM criterion may differ vastly from WANBB. This topic merits future study.

FIG. 11. Cumulative SCF iteration numbers of CG, WANBB, and LBFGS on PuO_2 surface with one adsorbed H_2O molecule.

APPENDIX

In this Appendix, we analyze the convergence property of WANBB. Before moving on, we define the level set

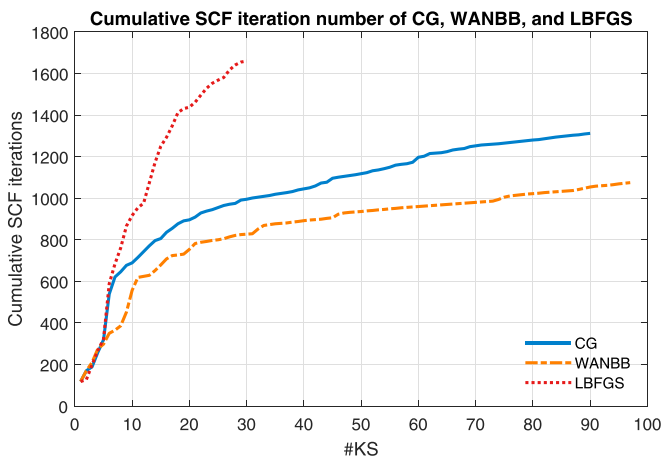$$\mathcal{L} := \{R : E(R) \leqslant E_0\}.$$

For the ease of reference, we restate the nonmonotone conditions in WANLM as follows: in the $(k + 1)$th iteration, R_{k+1} satisfies

$$E_{k+1} \leqslant B_k - c\alpha_k \|F_k\|_{\mathrm{F}}^2, \tag{A1}$$

$$\langle F_{k+1}, F_k \rangle \leqslant \sigma \|F_k\|_{\mathrm{F}}^2. \tag{A2}$$

By the arguments in [29, Lemma 1.1], we know that the step size α_k satisfying both Eqs. (A1) and (A2) always exists and that B_k is a convex combination of $\{E_i\}_{i=0}^k$.

We first show that the atomic position sequence generated by WANBB remains in $\mathcal{L}$.

Lemma 1. Let $\{R_k\}$ be the atomic position sequence generated by WANBB with $\mu_k \in [\mu_{\min}, \mu_{\max}] \subseteq [0, 1]$. Then we have $\{R_k\} \subseteq \mathcal{L}$.

Proof. We prove by contradiction. Suppose there exists $l_1 > 0$ such that $E_{l_1} > E_0$. By the condition (A1),

$$E_{l_1} \leqslant B_{l_1-1} - c\alpha_{l_1-1}\|F_{l_1-1}\|_{\mathrm{F}}^2 \leqslant B_{l_1-1}.$$

Since B_{l_1-1} is a convex combination of $\{E_0, \ldots, E_{l_1-1}\}$, there must exist some $l_2 \leqslant l_1 - 1$ for which $E_{l_1} \leqslant E_{l_2}$ holds. We can repeat the above arguments to obtain a non-negative sequence $\{l_i\}$ which is finite for its strict monotonicity and ends up with 0. Therefore, we get $E_0 < E_{l_1} \leqslant E_{l_2} \leqslant \cdots \leqslant E_0$, a contradiction. ∎

From the update formula (7), one could obtain that the surrogate sequence $\{B_k\}$ is always bounded from below by the potential energy sequence $\{E_k\}$.

Lemma 2. Let $\{R_k\}$ be the atomic position sequence generated by WANBB with $\mu_k \in [\mu_{\min}, \mu_{\max}] \subseteq [0, 1]$. Then we have $E_k \leqslant B_k$ for any $k \geqslant 0$.

Proof. The claim is true for $k = 0$ by definition. For $k \geqslant 1$, by the update formula (7) of B_{k+1}, one has

$$B_k = \frac{B_{k-1} + \mu_{k-1}Q_{k-1}E_k}{1 + \mu_{k-1}Q_{k-1}} \geqslant \frac{E_k + \mu_{k-1}Q_{k-1}E_k}{1 + \mu_{k-1}Q_{k-1}} = E_k,$$

where the inequality follows from the condition (A1). The proof is complete. ∎

With Lemmas 1 and 2 in place, we are ready to show the convergence of WANBB. The proof relies on the conditions (A1) and (A2) as well as the update formula (7). Specifically, the condition (A2) ensures a uniform positive lower bound of step sizes, while Eqs. (A1) and (7) guide us to a sufficient reduction over $\{B_k\}$.

Theorem 1. Suppose the potential energy is bounded from below and F is Lipschitz continuous, with modulus $L > 0$, on $\mathcal{L}$, namely,

$$\|F(R) - F(\bar{R})\|_{\mathrm{F}} \leqslant L\|R - \bar{R}\|_{\mathrm{F}}, \ \forall \ R, \bar{R} \in \mathcal{L}.$$

Let $\{R_k\}$ be the atomic position sequence generated by WANBB with $\mu_k \in [\mu_{\min}, \mu_{\max}] \subseteq (0, 1]$. Then we have

$$\lim_{k \to \infty} \|F_k\|_{\mathrm{F}} = 0.$$

That is to say, the atoms will approach equilibria, regardless of initial configurations.

Proof. Subtracting $\|F_k\|_{\mathrm{F}}^2$ from both sides of Eq. (A2), we have

$$\langle F_k - F_{k+1}, F_k \rangle \geqslant (1 - \sigma)\|F_k\|_{\mathrm{F}}^2.$$

Combining Lemma 1, the Lipschitz continuity of F in $\mathcal{L}$, and the fact $\sigma \in (0, 1)$, we can further derive

$$\begin{aligned} (1 - \sigma)\|F_k\|_{\mathrm{F}}^2 &\leqslant \|F_k - F_{k+1}\|_{\mathrm{F}} \cdot \|F_k\|_{\mathrm{F}} \\ &\leqslant L\|R_k - R_{k+1}\| \cdot \|F_k\| \\ &= \alpha_k L\|F_k\|_{\mathrm{F}}^2, \end{aligned}$$

and hence a lower bound for step size whenever F_k does not vanish:

$$\alpha_k \geqslant \frac{1 - \sigma}{L}. \tag{A3}$$

Plugging Eq. (A3) into condition (A1), we obtain

$$E_{k+1} \leqslant B_k - c\alpha_k\|F_k\|_{\mathrm{F}}^2 \leqslant B_k - \frac{c(1 - \sigma)}{L}\|F_k\|_{\mathrm{F}}^2. \tag{A4}$$

Let $\beta := c(1 - \sigma)/L$. By the update formula (7) of B_{k+1}, we obtain a recursion for the surrogate sequence using Eq. (A4):

$$\begin{aligned} B_{k+1} &= \frac{B_k + \mu_k Q_k E_{k+1}}{1 + \mu_k Q_k} \\ &\leqslant \frac{B_k + \mu_k Q_k(B_k - \beta\|F_k\|_{\mathrm{F}}^2)}{Q_{k+1}} \\ &= B_k - \beta\frac{\mu_k Q_k}{Q_{k+1}}\|F_k\|_{\mathrm{F}}^2. \end{aligned} \tag{A5}$$

A by-product of Eq. (A5) is the monotonicity of $\{B_k\}$. Since E is bounded from below on $\mathcal{L}$ and $B_k \geqslant E_k$ for any $k \geqslant 0$ (by Lemma 2), $\{B_k\}$ is bounded from below as well. Consequently, by Eq. (A5),

$$\sum_{k=0}^{\infty} \frac{\mu_k Q_k\|F_k\|_{\mathrm{F}}^2}{Q_{k+1}} \leqslant \frac{1}{\beta}\sum_{k=0}^{\infty}(B_k - B_{k+1}) < +\infty. \tag{A6}$$

Following from the update formula of $\{Q_k\}$ and the fact $\mu_{\max} \geqslant \mu_k \geqslant \mu_{\min} > 0$, $Q_k \geqslant 1$, we can derive

$$\frac{\mu_k Q_k}{Q_{k+1}} \geqslant \frac{\mu_{\min}Q_k}{\mu_k Q_k + 1} \geqslant \frac{\mu_{\min}}{\mu_k + 1/Q_k} \geqslant \frac{\mu_{\min}}{\mu_{\max} + 1},$$

which, together with Eq. (A6), yields the summability of $\|F_k\|$. Hence, $\|F_k\| \to 0$ as $k \to \infty$ as desired. ∎

[1] A. R. Oganov and C. W. Glass, J. Chem. Phys. **124**, 244704 (2006).

[2] Y. Wang, J. Lv, L. Zhu, and Y. Ma, Phys. Rev. B **82**, 094116 (2010).

[3] Z. Chen, W. Jia, X. Jiang, S. Li, and L.-W. Wang, Comput. Phys. Commun. **219**, 35 (2017).

[4] S. Goedecker, W. Hellmann, and T. Lenosky, Phys. Rev. Lett. **95**, 055501 (2005).

[5] L. B. Vilhelmsen and B. Hammer, Phys. Rev. Lett. **108**, 126101 (2012).

[6] C. H. Hu, A. R. Oganov, Q. Zhu, G. R. Qian, G. Frapper, A. O. Lyakhov, and H. Y. Zhou, Phys. Rev. Lett. **110**, 165504 (2013).

[7] Q. Li, D. Zhou, W. Zheng, Y. Ma, and C. Chen, Phys. Rev. Lett. **110**, 136403 (2013).

[8] L. Zhang, J. Luo, A. Saraiva, B. Koiller, and A. Zunger, Nat. Commun. **4**, 2396 (2013).

[9] J. Shewchuk, An introduction to the conjugate gradient method without the agonizing pain, Carnegie Mellon University, Technical Report (1994).

[10] P. Pulay, Chem. Phys. Lett. **73**, 393 (1980).

[11] D. Liu and J. Nocedal, Math. Program. **45**, 503 (1989).

[12] Z. Zhao, L.-W. Wang, and J. Meza, Phys. Rev. B **73**, 193309 (2006).

[13] S. Goedecker, F. Lançon, and T. Deutsch, Phys. Rev. B **64**, 161102(R) (2001).

[14] M. V. Fernández-Serra, E. Artacho, and J. M. Soler, Phys. Rev. B **67**, 100101(R) (2003).

[15] D. Packwood, J. Kermode, L. Mones, N. Bernstein, J. Woolley, N. Gould, C. Ortner, and G. Csányi, J. Chem. Phys. **144**, 164109 (2016).

[16] Z. Chen, J. Li, S. Li, and L.-W. Wang, Phys. Rev. B **89**, 144110 (2014).

[17] Z. Chen, J. Li, S. Li, and L.-W. Wang, Phys. Rev. B **96**, 115141 (2017).

[18] L. Liu, Z. Chen, Y. Yao, and L.-W. Wang, Phys. Rev. B **98**, 144109 (2018).

[19] J. Nocedal and S. Wright, *Numerical Optimization,* 2nd ed. (Springer, New York, 2006).

[20] W. Press, H. William, S. Teukolsky, W. Vetterling, A. Saul, and B. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing* (Cambridge University Press, Cambridge, 2007).

[21] J. Barzilai and J. Borwein, IMA J. Numer. Anal. **8**, 141 (1988).

[22] Y. Dai and R. Fletcher, Numer. Math. **100**, 21 (2005).

[23] B. Gao, X. Liu, and Y. Yuan, SIAM J. Sci. Comput. **41**, A1949 (2019).

[24] B. Gao, G. Hu, Y. Kuang, and X. Liu, SIAM J. Sci. Comput. **44**, B723 (2022).

[25] B. Gao, X. Liu, L. Wang, N. Xiao, and Y. Yuan, STOP–a toolbox for Stiefel manifold optimization.

[26] N. Xiao, X. Liu, and Y. Yuan, Optim. Methods Software, (2020).

[27] N. Xiao, X. Liu, and Y. Yuan, arXiv:2103.03514.

[28] L. Grippo, F. Lampariello, and S. Lucidi, SIAM J. Numer. Anal. **23**, 707 (1986).

[29] H. Zhang and W. Hager, SIAM J. Optim. **14**, 1043 (2004).

[30] We choose this specific value because the initial step size of CG in the software CESSP is 4.8×10^{-2} Å^2/eV. The numerical results of CG and WANBB are thus more comparable. Its value is not so critical.

[31] We find through numerical simulations that $\mu_k \equiv 0.05$ helps the WANLM criterion achieve an overall balance of stability and efficiency. The value of c follows the literature and is not so critical either.

[32] J. Fang, X. Gao, H. Song, and H. Wang, J. Chem. Phys. **144**, 244103 (2016).

[33] X. Gao, Z. Mo, J. Fang, H. Song, and H. Wang, Comput. Phys. Commun. **211**, 54 (2017).

[34] Y. Zhou, H. Wang, Y. Liu, X. Gao, and H. Song, Phys. Rev. E **97**, 033305 (2018).

[35] J. Fang, X. Gao, and H. Song, Commun. Comput. Phys. **26**, 1196 (2019).

[36] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).

[37] F. Jollet, M. Torrent, and N. Holzwarth, Comput. Phys. Commun. **185**, 1246 (2014).

[38] H. J. Monkhorst and J. D. Pack, Phys. Rev. B **13**, 5188 (1976).

[39] G. Kresse and J. Furthmüller, Comput. Mater. Sci. **6**, 15 (1996).

[40] G. Kresse and J. Furthmüller, Phys. Rev. B **54**, 11169 (1996).

[41] D. D. Johnson, Phys. Rev. B **38**, 12807 (1988).

[42] Jónsson group and Henkelman group, Transition state tools for VASP.

[43] R. P. Brent, *Algorithms for Minimization without Derivatives* (Prentice-Hall, Englewood Cliffs, NJ, 1973).

[44] A. Jain, S. Ong, G. Hautier, W. Chen, W. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. Persson, APL Mater. **1**, 011002 (2013).

[45] S. Borysov, R. Geilhufe, and A. Balatsky, PLoS One **12**, e0171501 (2017).

[46] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevB.106.104101 for information on the benchmark set and detailed numerical results.

[47] E. Dolan and J. Moré, Math. Program. **91**, 201 (2002).

[48] For visualization, we reset #KS$_{n,\text{alg}} = 10^5$ and CPU$_{n,\text{alg}} = 8.64 \times 10^7$ when the algorithm fails on the n-th system.

[49] To give a fair and expressive numerical comparison, we (i) filter out the the systems where the differences between the maximum and the minimum of the converged energies per atom by the four solvers are larger than 1 meV; (ii) retain the systems at which certain solvers fail. The final number of systems used for Fig. 4 is 72.

[50] J. Wu, Z. Yang, J. Xian, X. Gao, D. Lin, and H. Song, Front. Mater. **7**, 590143 (2020).

[51] H. Song, F. Tian, Q. M. Hu, L. Vitos, Y. Wang, J. Shen, and N. Chen, Phys. Rev. Materials **1**, 023404 (2017).