

Exploring potential energy surfaces to reach saddle points above convex regions

M. Gunde,¹ A. Jay,² M. Poberžnik,³ N. Salles,⁴ N. Richard,⁵ G. Landa,² N. Mousseau,⁶ L. Martin-Samos,⁴ and A. Hemeryck²

¹*Institute Ruder Bošković, Bijenička 54, 10000 Zagreb, Croatia*

²*LAAS-CNRS, Université de Toulouse, CNRS, 7 avenue du Colonel Roche, 31000 Toulouse, France*

³*Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia*

⁴*CNR-IOM/Democritos National Simulation Center, Istituto Officina dei Materiali, c/o SISSA, via Bonomea 265, IT-34136 Trieste, Italy*

⁵*CEA, DAM, DIF, F-91297 Arpajon, France*

⁶*Département de physique, Institut Courtois and Regroupement québécois sur les matériaux de pointe, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal (Québec), Canada, H3C 3J7*

(*Electronic mail: miha.gunde@gmail.com)

Saddle points on high-dimensional potential energy surfaces (PES) play a determining role in the activated dynamics of molecules and materials. Building on approaches dating back more than 50 years, many open-ended transition-state search methods have been developed to follow the direction of negative curvature from a local minimum to an adjacent first-order saddle point. Despite the mathematical justification, these methods can display a high failure rate: using small deformation steps, up to 80% of the explorations can end up in a convex region of the PES, where all directions of negative curvature vanish; while if the deformation is aggressive, a similar fraction of attempts lead to saddle points that are not directly connected to the initial minimum. In high-dimension PES, these reproducible failures were thought to only increase the overall computational cost, without having any effect on the methods' capacity to find all saddle points surrounding a minimum. Using the Activation-Relaxation Technique nouveau (ARTn), we characterize the nature of the PES around minima, considerably expanding on previous knowledge. We show that convex regions can lie on activation pathways, and that not exploring beyond them can introduce significant bias in the saddle-point search. We introduce an efficient approach for traversing the convex regions, almost eliminating exploration failures, while multiplying by almost 10 the number of identified unique and connected saddle points as compared with the standard ARTn, thus underlining the importance of correctly handling convex regions for completeness of saddle point explorations.

PACS numbers: 02.70.-c

I. INTRODUCTION

Saddle points (SP) on potential energy surfaces (PES) are of crucial importance in modern chemistry and materials science. These specific points correspond to the atomic structures that have the highest energy along the minimum energy paths (MEP) connecting two states. Within transition state theory¹, they represent transition states that provide the mandatory information needed to access the kinetics of an atomic system, including the energy barriers required to displace atoms. Consequently, they grant knowledge of the diffusion and reaction rates, which is crucial to comprehending the temporal evolution of an atomic system.

The PES is generally a complicated function of the configuration space: for 3D systems, it spans $3 \times N_{at}$ dimensions, where N_{at} is the number of atoms. Evaluating the PES for a given set of atomic positions is computationally demanding, as it requires knowledge of both the total energy and its first derivative (forces) that need to be computed in some way, generally using empirical potentials or quantum mechanics.

The algorithmic approach used to find SPs on a PES

depends largely on how much one already knows about the PES. The challenge can be categorized into: (i) Single-ended, or open-ended problems, for which only the initial structure is known; the PES must be explored all around the initial structure to find the adjacent SPs. (ii) Double-ended problems, for which both the initial and final structures are known; the PES exploration can then be limited to finding a continuous MEP between the two known structures, which are generally local minima of the PES. The most common algorithms used to solve double-ended problems are variations on the string method²⁻⁵. Although open-ended algorithms are used primarily to identify unexpected reaction mechanisms, they can also address double-ended problems by using a starting structure that is linearly interpolated between the initial and final ones⁶, or any other improved interpolation⁷. Today, the dominant open-ended SP search algorithms include the activation relaxation technique (ARTn)^{6,8-10}, the eigenvector-following method¹¹, the dimer method^{12,13} and the reduced gradient¹⁴. A review of algorithms that solve open-ended problems is given in Ref.¹⁵. All of them have initially been developed more than a quarter of a century ago.

The low-energy first-order SPs are the most relevant for the dynamical evolution of the system (first-order SPs are referred to simply as SPs in the rest of the paper). In fact, due to the Boltzmann factor determining their relative probability of occurrence, they are exponentially favored in thermodynamics. The regular critique of SP searching methods is that it is never certain that they are able to find all relevant SPs in a region of a PES, and that there is no straight-forward way of (dis)-proving that. Partially as an answer to that critique, it is often assumed that the regions of PES potentially inaccessible to the exploration mainly contain the high-energy SPs, which are not of much interest, and thus the ability of finding them is irrelevant. In addition, a sufficiently broad exploration should always find the low-energy SPs, assumed to reside in areas of the PES that can be described with relatively simple reaction coordinates, which the algorithm can easily explore. The hypothesis that low-energy SPs can be located in regions of the PES that are hardly or not accessible at all was therefore believed to be unlikely and was not really tested.

Within this context, the efficiency of open-ended algorithms for SP-search can be gauged based on two primary metrics: *(i)* the count of unique SPs found that are connected to the starting minimum, knowing that no mathematical theory provides the upper bound for that number in a generic system¹⁶; and *(ii)* the average number of force evaluations needed to reach a SP. As the paths to different SPs can take drastically different numbers of steps, a combination of these two metrics can also be used, to measure the total efficiency of an algorithm by counting the total number of force evaluations needed to find all the relevant SPs, *i.e.* the ones below a given energy threshold. However, such count is more dependent on the specific method selected to generate the initial displacements than on the SP search algorithm itself.

These two metrics are obviously influenced by the detailed implementation and parametrization of each algorithm, but more importantly, by the structure of the energy landscape itself. The open-ended method ARTn stops the search when, on the path to the SP, all directions that display a negative curvature of the PES vanish, because there is no promising direction to follow anymore. Mathematically, this happens when the system finds itself on a convex region (CR) of the PES. As shown in this paper, CRs can be located: *(i)* around the starting minimum, which can sometimes be reached when the algorithm relaxes back into the starting basin, or *(ii)* above the first inflection of the PES, and it is impossible to predict when they will be encountered using only local information. In some other algorithms such as the dimer method^{12,13}, the CRs are escaped by continuing to follow the direction of minimum curvature uphill until a negative eigenvalue reappears. In Section V, we show that this can lead to new SPs, but that better choices for exiting CRs are possible.

The problem *(i)* was identified a while ago¹⁰, and an elegant solution was implemented to limit its occurrence⁶.

The issue *(ii)* was long considered computationally costly to resolve and without any real impact on the SP search. In fact, it was assumed that in a high-dimensional space, it is always possible to find another path on the PES that goes around the CRs, and thus finding all SPs connected to the initial minimum would always be possible in some way.

In this article, the latter issue is critically revisited with a demonstration that these previous assumptions are incorrect. Building upon a generic analysis, we offer a detailed characterization of the features of a 2D PES, and how they affect pathways between a local minimum and the connected saddle point. We then show that this understanding is fully generalizable to higher-dimensional system; this allows us to explore various fundamental modifications to correct the major under-sampling of current open-ended methods.

Three different methods are proposed, implemented in ARTn, and compared, with the objective of properly overcoming convex regions during saddle-point searches. We find that the best solution of all three is to continue pushing the system by following a direction prescribed by a "double" random vector. The proposed method is general and can consequently be implemented in any minimum eigenmode following algorithm.

We show on a real example of solid-state physics that the features found on the 2D example, which could be regarded merely as toy model, are not only present in higher dimensions, but also of high importance for navigating the PES. We demonstrate explicitly that ignoring the CRs can result in the inability to explore a potentially significant portion of the PES, resulting in the failure to identify up to 90 per cent of truly connected saddle points.

Computationally, the proposed modifications reduce the fraction of failed attempts, which can reach 25 to 80 percent with the standard algorithms, to less than 1 per cent, considerably reducing the net number of force evaluations per successful saddle search, while multiplying by up to 10 the number of different connected saddle points found. Importantly, and contrary to previous assumptions, a significant fraction of low-energy SPs, that dominate kinetics, can be located in regions of the PES that have previously been inaccessible because of declaring the exploration as a failure once it encountered a CR.

This paper is structured as follows. Some definitions are first recalled in Section II. A 2D toy PES is analyzed to establish notions about its features in Section III. The current situation of the ARTn algorithm and its main issue are described in Section IV. The proposed improvements and the test case are detailed in section V. The results of the test case are presented in Section VI with discussions in Section VII.

II. DEFINITIONS

Throughout the paper, we use the following definitions, unless explicitly specified.

The local curvature of the PES is obtained from the Hessian matrix. It corresponds to the spectrum of eigenvalues λ_i . The eigenvector \mathbf{V}_{\min} corresponds to the lowest eigenvalue $\lambda_i = \lambda_{\min}$.

The parallel and perpendicular forces are defined, respectively, as: $\mathbf{F}_{\parallel} = (\mathbf{F} \cdot \mathbf{V}_{\min})\mathbf{V}_{\min}$ and $\mathbf{F}_{\perp} = \mathbf{F} - \mathbf{F}_{\parallel}$.

A specific point of the PES can be classified as either on a "valley" or "ridge", or "perpendicular hyperplane", depending on the direction of the gradient (force \mathbf{F}), and the local curvature of the PES, as follows. (i) The valleys are defined as the lines where the forces \mathbf{F} are parallel to \mathbf{V}_{\min} , *i.e.* $\mathbf{F}_{\perp} = 0$, and where there is *at most* one negative eigenvalue $\lambda_i < 0$. This negative curvature is along the direction of the valley. Therefore, when one moves from the valley line in any perpendicular direction, the potential energy increases. (ii) The ridges are defined similarly, as lines where the forces \mathbf{F} are parallel to \mathbf{V}_{\min} , however, there are *at least* two negative eigenvalues $\lambda_i < 0$. The first negative curvature is along the ridge line and the second perpendicular to it. Therefore, when one moves from the ridge line along the second negative curvature, the potential energy decreases. Similar definitions of valleys and ridges can be found in Hoffman¹⁷. (iii) The perpendicular hyperplanes (PH) are defined as hyperplanes where the forces \mathbf{F} are orthogonal to \mathbf{V}_{\min} , *i.e.* $\mathbf{F}_{\parallel} = 0$.

Any region of the PES where the lowest curvature is positive $\lambda_{\min} > 0$ is called a convex region (CR). If a CR contains the initial minimum, we call it the starting CR.

Any point where the lowest curvature is zero ($\lambda_{\min} = 0$) is called an inflection point. The inflection points together form inflection hyperplanes that are the boundaries of CRs.

Regions of the PES can be referred to as "below", or "beyond"/"above" a certain CR, with respect to the energy. The regions "below" are regions which can lead the system into a CR, while the regions "beyond" are the ones reached by starting from the boundary of a CR.

Many dynamics functions can be used to explore the PES, the ones used in this paper are the "minimization" and the "eigenmode following" SP exploration. Each of them exhibits some fixed points in the phase space, to which the dynamical evolution tends, or converges. These fixed points are called the attractors, for the minimization, they are minima points, and for SP exploration, they are the SPs. Any set of points on the PES where all the points evolve to the same attractor forms a region called the attracted region.

In higher dimensions, the valleys and ridges are still of dimension 1, whereas the perpendicular hyperplanes are of dimension $3N_{at} - 1$. The convex regions are hypervolumes of $3N_{at}$ dimensions, which are delimited by $3N_{at} - 1$ dimensional inflection hyperplanes, where the lowest eigenvalue is zero.

III. ANALYSIS OF THE PROBLEM

A 2D toy-model PES is analyzed in Fig. 1, in order to understand its characteristics, which are important from the point of view of SP search algorithms.

The region in black on Fig. 1a indicates all the points where a minimization function started from there, will end up in the minimum indicated by a larger black point, *i.e.* the black minimum is the fixed attractor point for the minimization function, and the black region is its attracted region. Throughout the paper, they are simply referred to as the minimum and its basin, respectively. This figure is generated by launching a minimization from 62 500 equally distributed points on a 250×250 2D grid and coloring the ones that relax to the black minimum. Starting the minimization from any point outside of this region will end up in a different minimum. The SPs are by definition located on the boundary of the basin (colored points on Fig. 1a), and a steepest-descent minimization from the SPs towards the minimum is a unique path connecting the SP and the minimum (colored lines in Fig. 1a), which is also called the intrinsic reaction coordinate (IRC). These SPs are defined as "connected" to the minimum.

In Fig. 1b, we change the function from minimization to a SP search algorithm that always follows \mathbf{V}_{\min} . The fixed attractor points are then the SPs indicated by colored points. The colored regions are the regions attracted to the corresponding SP, meaning that an exploration started from any point in a given region will end up at the SP of the same color. Notice that not all SPs are reachable by following the \mathbf{V}_{\min} starting from the minimum. In order to reach SPs other than yellow and magenta, the algorithm first has to move the starting position away from the minimum until it enters a given attracted region, and only then begins following the \mathbf{V}_{\min} .

In Fig. 1c, the features of the PES, namely the valleys, ridges, perpendicular hyperplanes, and CRs as defined in Sect. II, are shown in black, green, gray, and blue, respectively. Note several things from this figure as follows.

a. Reach SP by ascending a valley: Since a SP is a local maximum along a valley line, it can be reached by ascending the valley. Thus, reaching a SP by following \mathbf{V}_{\min} requires to have found that valley. Therefore, reaching all the SPs connected to a given minimum requires finding all the valleys present in the basin region of that minimum.

b. More SPs than eigenmodes – valleys must branch: In a two-dimensional system, as presented in Fig. 1, the minimum has only two eigenmodes. Following each eigenmode from the minimum in either \pm direction can thus potentially lead to 2×2 valleys in total and their corresponding four SPs (assuming that each eigenmode, even if it is not the lowest mode \mathbf{V}_{\min} , becomes a valley at some point, for which there is generally absolutely no guarantee). However, the number of SPs around the minimum in Fig. 1 is five, which means that at least one valley must necessarily "branch" into multiple valleys some-

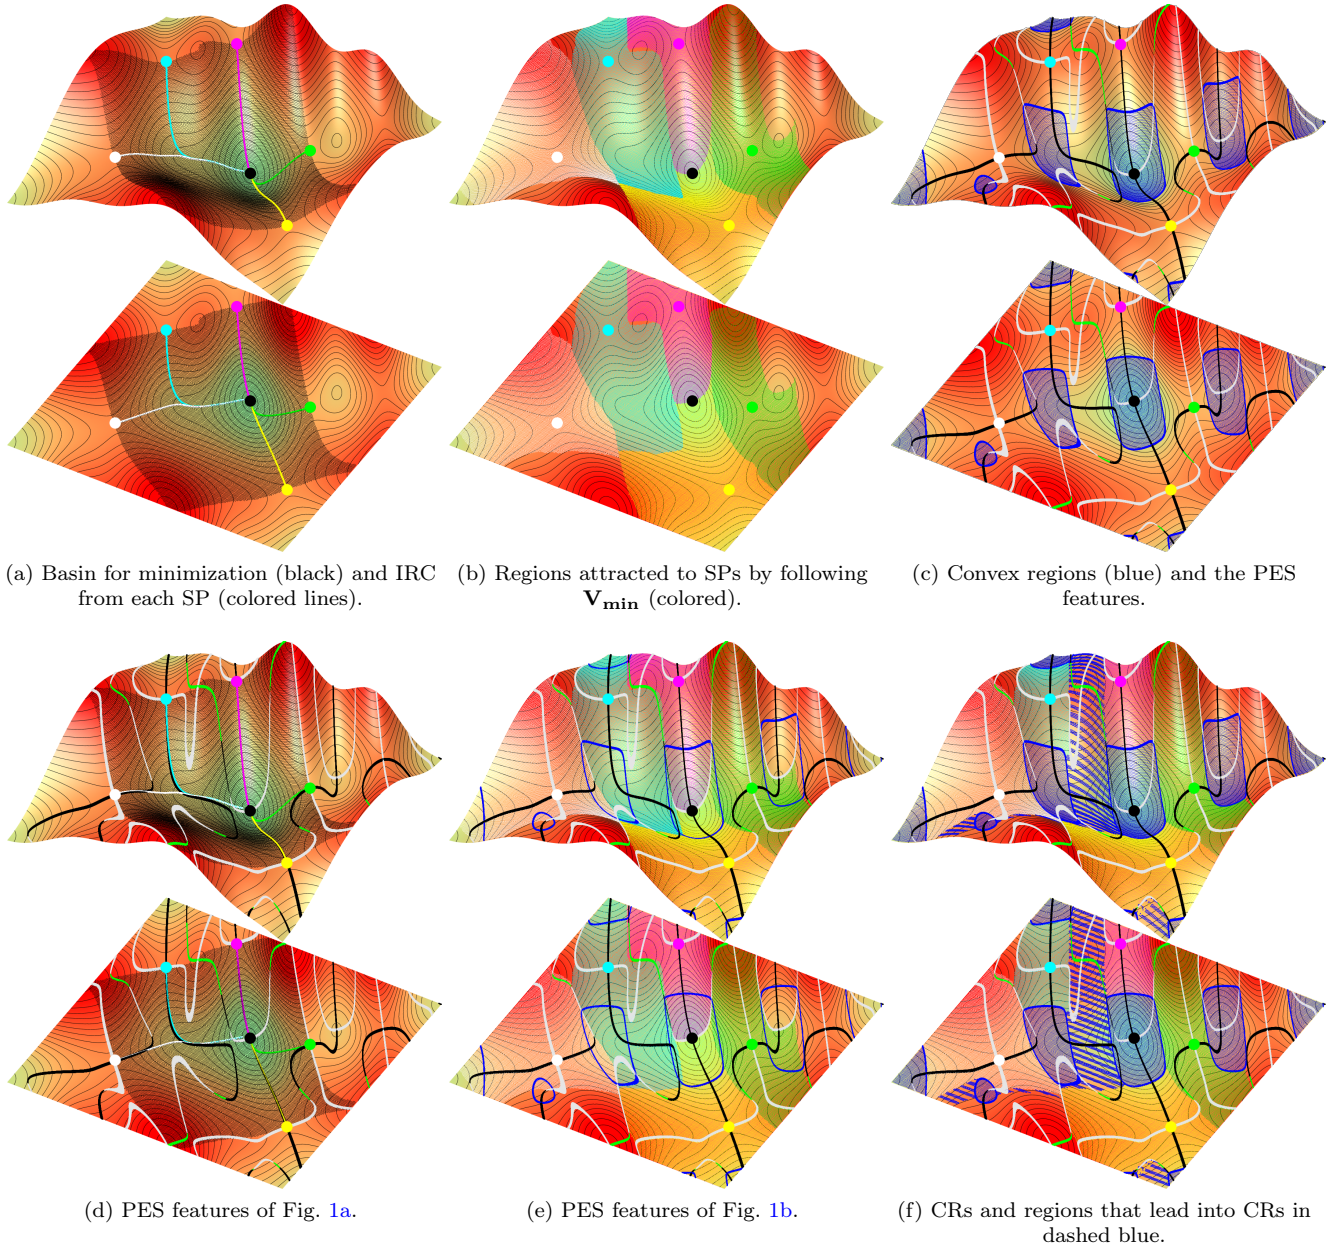


FIG. 1. Analysis of a 2D example PES where typical features can be found. The big points represent stable structures where $\mathbf{F} = 0$: minimum (black) and saddle points (colored). The regions of the PES colored in a certain way represent the set of points having the same characteristics: (1a, 1d): (black) basin for minimization, (colored lines) IRC. (1b, 1e, 1f): regions attracted to SPs of the same color. (1c, 1d, 1e, 1f): PES features, valleys (black), ridges (green), perpendicular hyperplanes (gray) and inflection hyperplanes (blue), as described in Sect. II. (1f): (dashed blue) the regions that lead to a CR by following \mathbf{V}_{\min} . The analytical value of the PES is $E(x, y) = \frac{1}{2} \cos(\frac{xy}{5}) \cos(\frac{3x}{5}) \cos(\frac{y}{2}) + \cos(x) \cos(\frac{3y}{2}) + \exp(-\frac{(x-17)^2 + (y-17)^2}{125})$ for $x \in [13; 18]$ and $y \in [14.7; 18.6]$.

where within the basin. In a realistic system with N_{at} atoms, the PES has $3N_{at}$ -dimensions and a minimum has $3N_{at}$ eigenmodes. Hence, a maximum of $2 \times 3N_{at}$ valleys and their SP can be reached when following \pm each of the eigenmodes from the minimum (limited to two SPs if following only \mathbf{V}_{\min}), while none of the "emergent"

SPs can be reached. The total number of SPs in a realistic system depends on the unknown number of valley branchings.

c. CRs other than starting CR: The minimum is located within a CR called the starting CR, but that is not the only CR present on the PES. We can observe

other CRs, where different things can happen, like a valley (black) that transforms into a perpendicular hyperplane (white) or vice versa. The point where this happens is noted by VtoPH. It corresponds to the "birth/death" of the valley.

d. Valleys branching: There are places where two valleys dissociate, which are generally called reaction path branching¹⁸, or valley bifurcations¹⁹. In the most general case where the PES is not symmetric, the two (or more) valleys of a branching do not cross by touching each other at a single point, but by only approaching each other, thus forming a forbidden crossing²⁰, also known as diabolic or exceptional points²¹. In this 2D example, a valley branching is located in the CR.

While most of the literature cited above generally speaks of bifurcation/branch *points*, we show with the simple 2D PES examples of Fig. 1c, that they possibly occur in regions larger than a single point. The crossings are indeed reduced to a single point in highly symmetric molecular systems, such as the ones studied in most associated literature. A particular example of a highly symmetric PES that exhibits a single point branching of valleys is presented in Appendix B.

e. Other significant points and regions: Valley-ridge inflection points (VRI)²², located at the contact between the black (valleys) and green (ridges) lines in Fig. 1c are important regions of the PES. Algorithms do not use infinitesimal, but finite-sized steps, which implies that ridges can sometimes be reached and also crossed. However, since the ridges are defined with at least two negative eigenvalues, they cannot be located within CRs, where all the eigenvalues are positive. As these structures are not hindering the SP search procedure, VRI are not discussed further in this article.

By definition, the boundaries of the basin can be obtained from each SP by following the hyperplanes orthogonal to the isoenergy hyperplanes (hyperplanes on which the energy is constant). As shown in Fig. 1d, in the vicinity of the SPs, this is equivalent to following the perpendicular hyperplane. In Fig. 1e, these PES features are superposed above the regions attracted to the SPs. One can observe that the boundaries of these regions are mainly perpendicular hyperplanes that act like uncrossable barriers if the size of the displacement used by the algorithm is small enough.

f. Implications for SP search algorithms: The main goal of SP exploration algorithms is to find as many SPs as possible for a given basin. Thus, it can be concluded from the preceding observations that the main focus is to find all the valleys present in the basin, no matter where they originate. This implies first moving away from the initial minimum. The choice of the algorithm is then to specify when it is "far enough" to start following \mathbf{V}_{\min} . The algorithm must be able to reach the regions where valley branchings occur, in the form of a forbidden crossing or otherwise. Cases where a branching occurs below a CR should not be problematic, since consecutive searches with different initial directions can easily access

all its branches. However, branching can occur within CRs, or above. Therefore, to reach these new valleys, it should be important to pass through a CR when it is encountered and explore the regions beyond it.

These observations are not specific to the particular region of the example 2D PES analyzed above. A different region of the PES with similar features is shown in Appendix A.

It might not be obvious or straightforward that observations made on a 2D example PES can be generalized to a realistic system with $3N_{at}$ dimensions. However, as shown in Sections V and VI, realistic systems actually do exhibit the described properties of the PES, and some notions from the 2D example can be applied in higher dimensions.

IV. THE SITUATION IN ARTN

The original ARTn method is discussed in more detail elsewhere, for instance, in Refs.^{6,10,23,25}. In the present section, we highlight its features with respect to the insights from Sec. III. A schematic is given in Fig. 2, where the improvements present in this paper are highlighted in blue and detailed in Section V A.

A. Original algorithm

The ARTn algorithm can start the PES exploration from any structure (point on the PES), not only strictly from a minimum structure⁶. However, most applications require starting in a minimum, therefore we specify the description to that case. The overall algorithm does not change when the starting structure is not in a minimum.

A saddle point is a maximum along one dimension (its valley) and a minimum along all other dimensions (its perpendicular hyperplane defined in Section II). Therefore, to reach a SP, the ARTn algorithm is based on the following three-step procedure:

1. Compute λ_{min} , and the corresponding \mathbf{V}_{\min} at the current position. If $\lambda_{min} < 0$, set $\mathbf{V}_{\text{push}} = \mathbf{V}_{\min}$.
2. Push the system uphill along \mathbf{V}_{push} .
3. Relax \mathbf{F}_{\perp} , the forces perpendicular to \mathbf{V}_{push} .

At first, the starting structure is deformed away from the minimum in a random (or constructed) direction $\mathbf{V}_{\text{push}} = \mathbf{V}_{\text{rinit}}$ until a negative eigenvalue λ_{min} is found, indicating that the boundary of the starting CR has been passed, beyond which $\mathbf{V}_{\text{push}} = \mathbf{V}_{\min}$ is followed. This initial random deformation already avoids biasing the search to the valleys directly connected to the initial minimum. Note, however, that despite this initial random deformation, the white SPs is still unreachable in Fig. 1e.

In the starting CR, the perpendicular relaxations only serve to avoid atomic collisions and non-physical configurations that can arise from the random deformation.




FIG. 2. The 3 steps of the ARTn algorithm (red) taken from Refs.^{23,24}. The new part that allows to escape the CRs is in blue. \mathbf{V}_{CR} is the vector used to push the system when it is in a CR: it can be $\mathbf{V}_{\text{rinit}}$, \mathbf{V}_{new} , $\mathbf{V}_{\text{rinit}} + \mathbf{V}_{\text{new}}$ or \mathbf{V}_{min} , depending on the method. ϵ is the user-defined force threshold that defines the saddle point.

The number of relaxation steps is then reduced to 1 or 2. Above the inflection, it serves to fall down into the valley and a higher number of relaxation steps must be performed. To avoid perpendicular relaxation that brings the system back to the starting CR, the switch of the \mathbf{V}_{push} direction from $\mathbf{V}_{\text{rinit}}$ to \mathbf{V}_{min} can be smoothed in a small number of steps⁶, which helps steer the structure away from the starting CR. The number of perpendicular relaxation steps can also be reduced in the vicinity of the starting CR and progressively increased.

It has been observed that the convergence of ARTn to an SP is faster when the rate of convergence in the parallel and perpendicular directions is approximately equal. Therefore, each ARTn step outside of any CR performs the perpendicular relaxation for a desired number of steps, or until $\mathbf{F}_{\perp} < \mathbf{F}_{\parallel}$.

ARTn steps are repeated until an SP is reached, which is characterized by a negative curvature $\lambda_{\text{min}} < 0$ and a total force of the system $\mathbf{F} = 0$, or operationally lower than a given threshold. At a first-order SP, λ_{min} is the unique negative eigenvalue and its corresponding eigenvector \mathbf{V}_{min} lies along the valley direction.

Because only the minimum eigenvalue λ_{min} and the eigenvector \mathbf{V}_{min} are sought, the full Hessian matrix is never constructed or diagonalized, but only a subset of the space is generated using the Lanczos algorithm²⁶. The direction of \mathbf{V}_{push} is always opposite to \mathbf{F}_{\parallel} . Due to this, a structure that is pushed slightly through a boundary of the basin can still find its way to the nearby SP. In fact, as seen by comparing Figs. 1a and 1b, the regions attracted by the SPs extend beyond the boundary of the basin.

Finally, to confirm that the SP found is on the hyperplane marking the boundary between two basins, the initial and the other one, an energy minimization is performed from the SP after a slight displacement away from it in both the \pm directions of \mathbf{V}_{min} . If one of the two end points of this minimization corresponds to the initial structure, the SP is considered connected.

B. The main issue

The main issue with the current ARTn algorithm is that it considers any CR other than the starting CR as a dead end. As a result, any exploration directed to a region of the PES where ARTn leads the structure into a CR is useless, since the algorithm will stop and no SP will be found going there. For the 2D toy-model, such regions are marked by dashed blue in Fig. 1f. It means that all the $\mathbf{V}_{\text{rinit}}$ directions emerging from the starting CR toward the indicated area will end up in the second CR and result in a failed SP search attempt. This situation is not anecdotal: for the 2D model, failed searches represent about 35% of the possible initial directions.

The historical decision to stop the exploration when losing the negative curvature was due to two false hypotheses. It was assumed that (i) the fact that some SPs could not be reached by following \mathbf{V}_{min} from the edge of the starting CR is limited to low-dimensional PES, and (ii) it is always possible to find another path on the PES that goes around the CRs. However, contrary to what was expected, we show in Section VI that, in higher dimensions, the lost SP searches associated with such features of the PES can represent up to more than 80% of the explorations. High dimensions can bring more complexity: for some systems, a significant number of CRs must be crossed before reaching an SP, which implies that, depending on the systems, a sizable fraction of the PES and its SPs are not reachable by the current ARTn algorithm. Therefore, simply bringing the system beyond the starting CR before following \mathbf{V}_{min} is not sufficient to reach the SPs located above further CRs, such as the cyan and white SPs shown in Fig. 1f. To mitigate the issue and ensure the complete sampling of SPs, further changes of the algorithm are required.

In the following, to facilitate reading, we qualify SPs as "direct" when they can be reached from the starting CR without crossing any additional CR, and as "indirect" when at least one other CR has to be crossed.




FIG. 3. SPs reached with ARTn starting from the same minimum, as a function of the vector \mathbf{V}_{CR} used to escape the CRs encountered. Starting from the minimum, 200 searches are launched in uniformly distributed $\mathbf{V}_{\text{rinit}}$ directions, forming a circle in 2D. Each of the 200 slices of the circle around the initial minimum is colored according to the color of the SP reached, with red indicating directions leading to failed explorations (no SP reached due to trajectory cycling in a loop). The $\mathbf{V}_{\text{CR}} = \mathbf{V}_{\text{rnew}}$ case is not shown, as it gives the same kind of results as Fig. 3c, but with a larger number of force evaluations. A video showing the complete exploration leading to these results is given in the Supplementary Material with its gnuplot script.

V. METHODS AND PROCEDURES

As has been described in the previous sections, CRs should not be simply discarded as a “failed search”, because this failure blocks the access to the portions of the PES that are above them and consequently to its SPs. Ideally, every time a CR is encountered, the PES should be explored to find all the merging valleys. Each valley should then be followed, as they all lead to an SP or to a valley-ridge inflection point²⁷. As it is *a priori* not possible to distinguish the location of the branchings by simply looking at the local eigenvalue spectrum, a method to find all of them could be to start a number of new fully random explorations in multiple independent directions. However, such a method would be computationally demanding and potentially inefficient. In the following, we propose and explore some lower-cost alternative procedures to continue the ARTn SP search when encountering a CR. Such a procedure must avoid terminating ARTn exploration in cases where $\mathbf{V}_{\text{rinit}}$ is directed from the initial minimum toward the dashed blue area in Fig. 1f. It must be able to reach the white and cyan attracted regions, and to reach the SPs that are above CRs.

A. Proposed procedures

A CR is considered detected when the lowest eigenvalue λ_{min} becomes positive while previously negative ($\lambda_{\text{min,prev}} < 0$). When this happens, the current ARTn algorithm schematized in Fig. 2 aborts its SP search, as \mathbf{V}_{min} in CR does not contain particular information with respect to a connected SP.

To prevent exploration failure and in the absence of local information, the push vector in CR, \mathbf{V}_{push} , is no longer determined by \mathbf{V}_{min} and must be assigned a new direction $\mathbf{V}_{\text{push}} = \mathbf{V}_{\text{CR}}$. This modification is highlighted in blue in Fig. 2. The vector \mathbf{V}_{CR} is chosen so that it allows reaching any of the valleys that potentially emerge from the CR or are located beyond it.

Operationally, when a CR is detected, \mathbf{V}_{CR} is defined as the new push vector until a negative eigenvalue λ_{min} is detected. \mathbf{V}_{CR} is further updated if a new CR is detected. Inside the CR, the norm of the pushing vector along \mathbf{V}_{CR} is set to constant, because the goal is to rapidly pass the inflection, not to converge somewhere. This choice is discussed in Section. VII C). Here, four definitions for \mathbf{V}_{CR} are tested:

- $\mathbf{V}_{\text{CR}} = \mathbf{V}_{\text{min}}$, where the lowest eigenvector \mathbf{V}_{min} is followed regardless of the sign of the convexity of the PES. This is the simplest and most obvious choice. However, doing this in a CR containing a forbidden crossing of two valleys always leads to the

same valley (the one most aligned with the \mathbf{V}_{\min} calculated when entering the CR), and will consequently always ascend to the same SP. In the 2D example shown in Fig. 1f, following the valley that enters the CR containing the branching will always lead to the cyan SP, while the second valley merging from the convex region will never be followed. Consequently, the white SP will never be reached. This problem is illustrated in Fig. 3a, where exploration of the PES starting from the minimum is performed in all possible initial directions $\mathbf{V}_{\mathbf{r}_{\text{init}}}$ up to a negative curvature, from which point the eigenvector \mathbf{V}_{\min} is followed even if the ARTn algorithm encounters a CR. Each $\mathbf{V}_{\mathbf{r}_{\text{init}}}$ direction is then colored with the color of the SP reached. Only four (yellow, green, magenta, and cyan) of the five saddle points are reached with this approach (not white). Therefore, the proposed modification $\mathbf{V}_{\text{CR}} = \mathbf{V}_{\min}$ cannot meet the objectives described above.

- $\mathbf{V}_{\text{CR}} = \mathbf{V}_{\mathbf{r}_{\text{init}}}$, where the starting random vector is reused. This ensures that the system does not move back towards the direction of the initial minimum. The 2D example using this vector in Fig. 3b works well, but it can generate some cycling issues represented by the red directions (cycling issues are also observed in higher-dimensional atomistic systems, see Section VI B).
- $\mathbf{V}_{\text{CR}} = \mathbf{V}_{\mathbf{r}_{\text{new}}}$, where $\mathbf{V}_{\mathbf{r}_{\text{new}}}$ is a completely new random vector. As the positions of the valleys within the CR are unknown, all directions should be considered with equal probability. With such an approach, two different ARTn searches that enter the same CR will leave it following two different directions and potentially reach two different valleys. To be efficient, an isotropic random vector is first generated and then multiplied by $\mathbf{V}_{\mathbf{r}_{\text{init}}}$ resulting in a vector $\mathbf{V}_{\mathbf{r}_{\text{new}}}$ that is specifically localized on the atoms of interest. Multiplying $\mathbf{V}_{\mathbf{r}_{\text{new}}}$ by the local forces has also been tried in this study; this approach, however, often leads to SPs that are not localized around the initial central atom. As suggested in Ref.²⁸, a force threshold could have been used to confine the displacement on the atoms that have the highest force, but this represents an additional user-defined parameter to finally get quite the displacements of $\mathbf{V}_{\mathbf{r}_{\text{init}}}$.
- $\mathbf{V}_{\text{CR}} = (1 - \alpha)\mathbf{V}_{\mathbf{r}_{\text{init}}} + \alpha\mathbf{V}_{\mathbf{r}_{\text{new}}}$, where $\alpha \in [0; 1]$ is a mixing parameter (see the Appendix for its optimization). This procedure generates a wide variety of directions originating from the same initial $\mathbf{V}_{\mathbf{r}_{\text{init}}}$, while preventing pushing the system back to the minimum. In 2D, the possible directions that can be generated are represented by the red area in Fig. 4. This effectively combines the advantages of $\mathbf{V}_{\text{CR}} = \mathbf{V}_{\mathbf{r}_{\text{init}}}$ and $\mathbf{V}_{\text{CR}} = \mathbf{V}_{\mathbf{r}_{\text{new}}}$ procedures. An additional benefit is that this procedure avoids the

looping trajectory discussed in Section VI B. This approach offers the best results on the 2D model system, as seen in Fig. 3c, and all SPs are reached without any failure.




FIG. 4. Black vector $\mathbf{V}_{\mathbf{r}_{\text{init}}}$: starting random vector pointing away from the initial minimum. Blue vector $\mathbf{V}_{\mathbf{r}_{\text{new}}}$: new random vector reinitialized each time a CR is detected; it can point anywhere. Red vector: $\mathbf{V}_{\text{CR}} = (1 - \alpha)\mathbf{V}_{\mathbf{r}_{\text{init}}} + \alpha\mathbf{V}_{\mathbf{r}_{\text{new}}}$ is the new pushing direction used in CRs. All vectors are normalized. The red area corresponds to the possible directions for \mathbf{V}_{CR} vector as a function of $\mathbf{V}_{\mathbf{r}_{\text{new}}}$ orientation for a given α . $\alpha = 0.42$ is used in this schematic representation.

B. Applications to complex atomistic systems

The proposed methods have been tested on many different systems to validate their generality. We consider here an amorphous system that represents the problem described in Section III and a highly degenerate case is discussed in the appendix B. The amorphous system is chosen because the initial minimum is connected to a very large number of different SPs, and because many CRs need to be crossed to get them. This is one of the worst situations for SP exploration algorithms. Explorations of the PES have been performed with the ARTn code in its fast plugin version²⁵ coupled with LAMMPS²⁹.

We start with a simulation box of 1000 silicon atoms in a highly stable amorphous configuration created as described in Ref.³⁰ and stabilized following the procedure of Ref.³¹. Interatomic forces are calculated using a modified Stillinger-Weber empirical potential³². For each of the four \mathbf{V}_{CR} procedures proposed in Section V A, 3000 explorations are performed, all starting from the same stable structure that represents a local minimum of the PES. Each exploration begins with a push along a randomly selected vector $\mathbf{V}_{\mathbf{r}_{\text{init}}}$ away from the initial minimum. This random vector is localized on a given atom and all its

neighbors contained within a 3.5 \AA radius. The central atom and its four neighbors remain the same throughout the 3000 explorations, to allow detailed sampling of the 15-dimensional space of initial random deformations. The displacement step norm is set to 0.2 \AA when λ_{min} is positive and becomes proportional to the forces above the inflection (see Fig. 2). To compare the efficiency of the four \mathbf{V}_{CR} procedures described above, the 3000 different starting random vectors $\mathbf{V}_{r_{init}}$, the starting Lanczos vector, and the random seed are identical for each of the four sets. To limit the computational cost of cycling trajectories, all SP searches are stopped (returns a failure) when the number of force calculations exceeds 4000 or when the number of encountered CRs is larger than 30.

The force threshold used to determine the convergence to an SP is set to 10^{-5} eV/\AA . This very stringent threshold is needed to ensure that new and already visited SPs are properly identified, as in amorphous materials, atoms can be easily displaced without necessarily changing the energy significantly; thus, a precise convergence is essential for obtaining similar atomic positions in different calculations of the same SP. Two SPs are considered identical if the absolute differences in their energy and the distance between structures are less than 10^{-2} eV and 10^{-1} \AA , respectively.

VI. RESULTS


A. No failures and new saddle points

The first remarkable effect of continuing to push beyond the CRs is that regardless of the selected \mathbf{V}_{CR} method, the number of saddle points found is drastically increased as new portions of the PES become accessible to the algorithm. This is shown in Fig. 5: only 23 different direct saddle points (big red squares) are generated using the previous implementation of ARTn, whereas hundreds of indirect SPs are found with the modifications proposed here. Not surprisingly, the SP energy increases with the number of CRs crossed during the SP search. However, even if most of the new SPs have a high energy, some of these energies are relatively low, comparable to those of the direct SPs. This is the main message of the article: some thermodynamically important SPs are missed without crossing CR. Interestingly, the SPs that are most often found, with all methods, are those close, in distance, to the initial structure. As shown in Fig. 6, these SPs are reached, for the most part, either directly, without crossing a CR, or quasi-directly, crossing no more than 2-3 convex regions.

When following only the eigenvector $\mathbf{V}_{CR} = \mathbf{V}_{min}$, the number of different saddle points found is about half that of the other proposed vectors, regardless of the number of CRs crossed, as shown by the numbers at the top of Fig. 6. Most of them are not localized on the five atoms of interest, those that have been moved by $\mathbf{V}_{r_{init}}$ to escape the starting CR.

\mathbf{V}_{CR}	Prev	\mathbf{V}_{min}	$\mathbf{V}_{r_{init}}$	$\mathbf{V}_{r_{init}} + \mathbf{V}_{r_{new}}$	$\mathbf{V}_{r_{new}}$
Total SP	712	2793	2967	3000	3000
All CSP	708	2095	2591	2652	2711
Unique CSP	23	127	248	237	225

(a) Total number of SPs found with each method. *All CSP*: counting only the connected SP. *Unique CSP*: counting CSPs reached several times only once. *Prev*: previous ARTn approach that stops in CR.



(b) As a function of the SPs energy relatively to the starting minimum: atomic displacement from the initial minimum (top panel), occurrence of each SP (middle panel) and cumulative number of unique saddle points (bottom panel). For clarity, only connected SPs below 5.5 eV and 4.5 Å are plotted, and \mathbf{V}_{min} results are not shown.

FIG. 5. SPs connected to the initial minimum (CSP) found in an amorphous silicon cell after 3000 ARTn searches as a function of the method used to handle CRs (\mathbf{V}_{CR}). \mathbf{V}_{min} , $\mathbf{V}_{r_{init}}$, $\mathbf{V}_{r_{init}} + \mathbf{V}_{r_{new}}$ and $\mathbf{V}_{r_{new}}$ are described in the text.

The second remarkable effect of pushing through a CR is that ARTn almost always finds a saddle point, drastically increasing its efficiency. As shown in Tab. 5a, the ratio of explorations that lead to a direct saddle point does not exceed 25%. With the optimal \mathbf{V}_{CR} , except for some cases described in the next section, 100% of the explorations lead to an SP. Furthermore, 75% of the SPs found do not need to cross more than 3 CRs. For the other 25%, it is hard to precisely affirm that all the CRs encountered are really separated: it is possible that the algorithm revisits previously encountered large CRs, for example.

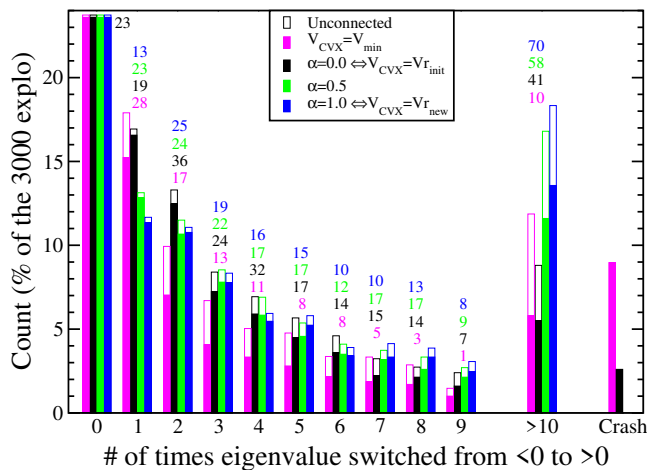


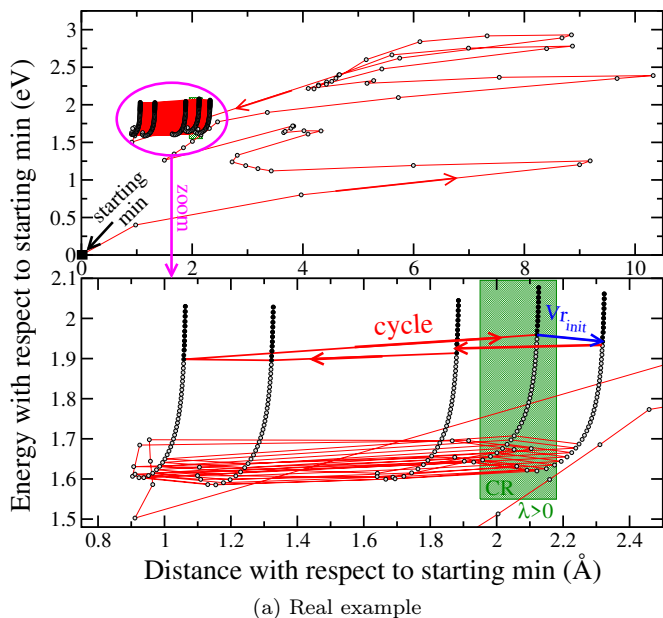
FIG. 6. Histogram of number of times an eigenvalue switches from negative to positive (detecting a CR) over 3000 SP searches for each of the four proposed vectors. "Crash" indicates that the research has not stopped before 30 CRs (cycling), or that the number of force calculations is greater than 4000. The numbers in the plot indicate how many new connected SPs have been found due to the increase of CRs crossed. Note that the 23 direct SPs are found in all simulation sets since V_{CR} modifies only the treatment of indirect ones, with pathways that go through a CR.

Interestingly, $V_{CR} = V_{r_{new}}$ results in a larger number of pathways crossing more than 10 CRs (blue bars Fig. 6): the fully random direction used here sometimes pushes the system towards the initial minimum. This implies that a subsequent push outside the CR ($V_{push} = V_{min}$) will most likely bring the system back into the same CR, thus increasing the number of encountered CRs without helping the algorithm to find SPs. This shows the importance of preserving a part of $V_{r_{init}}$ in V_{CR} .


B. Avoid cycling

In some rare cases, which represent less than 1% to 3% of the search directions, the eigenvalue alternates between positive and negative a large number of times. As shown in Fig. 6 (column Crash), this cycling behavior is only observed when using $V_{CR} = V_{r_{init}}$. Looking into the atomic structures visited with ARTn, we see that after some standard steps, the algorithm becomes trapped in a looping pattern on the PES, which forces the structure to escape and relax back into the same CR(s) indefinitely. This cycling behavior is presented in Fig. 7 where the loop is composed of one CR (green area), but loops composed of up to three different CRs have been observed. A video of the cycling behavior on the 2D PES is available in the Supplementary Materials. Applying a double random vector $V_{CR} = (1 - \alpha)V_{r_{init}} + \alpha V_{r_{new}}$ avoids this problem, since the random part added to $V_{r_{init}}$ each time a CR is detected ensures that the escape direction is never repeated, so infinite cycling is

prevented.



(a) Real example



(b) 2D toy model

FIG. 7. Cycling behavior observed in 1% to 3% of the ARTn explorations with $V_{CR} = V_{r_{init}}$. Gray and black dots: older and last intermediate structures. Blue arrow: $V_{r_{init}}$. (7a): Red arrows indicate the chronological order of the intermediate structures. Top panel: global view of the algorithm. Bottom panel: Zoom around the region where the algorithm is trapped. Green area: approximate CR. (7b): The green, cyan and pink arrows show the cycle occurring using the $V_{CR} = V_{r_{init}}$ each time it enters the CR. The direction between two black points in CR and that of V_{CR} are not the same due to the perpendicular relaxation after each push.

C. Unconnected paths

The ratio of SPs connected to the initial minimum to the total number of SPs generated is approximately 90% in all cases (see the top table in Fig. 5). This particularly good score depends on the PES and the set of ARTn parameters, mainly the maximum size of the displacements. However, the ratio of unconnected to total number of SPs increases with the number of crossed CRs (see Fig. 6). The unconnected SPs are reached because a boundary of the basin has been crossed, probably due to a too large push step while emerging from a CR, in which the step size is fixed. Disconnected SPs are also more frequent with $\mathbf{V}_{\text{CR}}=\mathbf{V}_{\text{min}}$: in physical systems, \mathbf{V}_{min} inside the CRs generally corresponds to a highly delocalized mode, such as quasi-acoustic modes. Following that direction then induces highly delocalized displacements, favoring disconnection with the basin. Note that leaving the basin is obviously a mandatory condition to reach an unconnected SP, but is not sufficient. In fact, the regions attracted by the connected SP pass over the basin's boundaries (see the difference between Figs. 1a and 1b), thus connected SPs can be reached from outside the basin. The sufficient condition is to pass over the regions attracted by the connected SPs, which are *a priori* unknown.

D. Optimization of the mixing parameter

The best mixing parameter to define \mathbf{V}_{CR} is obviously PES-dependent. As shown in Fig. 8 for the system studied here, while a very low α leads to a lower number of SPs reached and a lower fraction of SPs connected to the initial minimum, a $0.25 \leq \alpha \leq 0.4$ generates the highest number of different but connected SPs, a solid measure of sampling quality. As α increases, the number of lost SPs falls to zero at the cost of a slight decrease (15%) in the number of different connected SPs.

The results of Fig. 8 confirm that each $\mathbf{V}_{\mathbf{r}_{\text{init}}}$ represents overall the correct direction to follow, and the random part serves mostly to avoid cycles and provide access to the merging valleys that have turned towards a direction drastically different from $\mathbf{V}_{\mathbf{r}_{\text{init}}}$.

VII. IMPROVING THE EFFICIENCY OF THE METHOD

Several additional improvements to ARTn are possible, as discussed in the section.

A. Detecting the basin boundaries

An unconnected SP is reached when one of the hyperplanes that binds the basin is crossed during the SP search. Looking at the regions around the SPs in Fig. 1d,




FIG. 8. Count of the number of SP reached using different \mathbf{V}_{CR} in the convex region.

these boundaries correspond to the hyperplanes orthogonal to \mathbf{V}_{min} . Crossing them requires the system to be close enough to one of those and to undergo a sufficient push along \mathbf{V}_{min} to go over those. However, in the "minimum mode following" algorithms, the norm of the push along \mathbf{V}_{min} is proportional to the parallel force \mathbf{F}_{\parallel} , which is reduced the closer the system gets to an orthogonal hyperplane (where $\mathbf{F}_{\parallel} = 0$). At a hyperplane, the algorithm only performs orthogonal relaxations down to the SP (as the size of the push across the hyperplane is zero, due to the force parallel to \mathbf{V}_{min} being zero). The hyperplanes orthogonal to \mathbf{V}_{min} therefore act as uncrossable barriers. On the contrary, it is possible to observe in Fig. 1d that far from the SPs, the basin is not bounded everywhere by the hyperplanes orthogonal to \mathbf{V}_{min} . There are paths in which the algorithm can potentially leave the initial minimum attractor. This behavior can typically occur when the size of the orthogonal relaxation steps, or of the pushes along \mathbf{V}_{min} are too large. To avoid crossing these boundaries, a simple solution is to reduce the maximum size of the displacements, but this is at the price of a higher computational cost for all the other explorations. Therefore, one of the possible improvements should be a detection of the minimum attractor boundaries. This detection should automatically reduce the size of the push and put the algorithm back to its previous position on the PES, which allows the use of step sizes that do not limit the other explorations.

B. Smooth switching of push vector

One of the previously implemented improvements of ARTn consisted of a way to smoothly change the direction of the push from the random initial $\mathbf{V}_{\mathbf{r}_{\text{init}}}$ to the eigenvector \mathbf{V}_{min} in a series of a few steps, when escaping the starting basin⁶. The initial goal of this option was to avoid driving the system back into the starting basin at the first perpendicular relaxation step above

the inflection. However, we have observed that applying the smoothing option whenever exiting any CR helps to push the system sufficiently far away from the CR, before recommencing the regular algorithm, *i.e.* following the minimum eigenmode. It therefore represents an excellent option to help navigate the system away from the CR that has just been crossed. Unfortunately, performing the smoothing in too many steps can increase the number of unconnected saddle points. Typically, this occurs when a boundary of the basin is close enough to the crossed CR, such that this boundary can be crossed by a smoothed push vector that has a constant size. Therefore, to use the smoothing option to its full potential, it is recommended to limit the smoothing procedure to a single step when leaving the CR, *i.e.*, to perform a single intermediate step with $\mathbf{V}_{\text{push}} = \mathbf{V}_{\text{CR}}/2 + \mathbf{V}_{\text{min}}/2$

C. Choice of $\mathbf{V}_{\text{rinit}}$ amplitude

There are two strategies to escape the starting CR in a random direction: *(i)* Progressively push the system along $\mathbf{V}_{\text{rinit}}$ using small amplitudes up to the inflection hyperplane; each push followed by a few perpendicular relaxation steps to reduce the forces and avoid atomic collisions. *(ii)* Drastically push the system along $\mathbf{V}_{\text{rinit}}$ with a large amplitude with the expectation to immediately step over the first inflection hyperplane surrounding the minimum. The step amplitude can also be randomized with a Gaussian distribution in order to make some bigger deformations, which hopefully bring the system beyond the places where the valleys are "created" (valley branchings). Even if this second strategy requires fewer force evaluations to escape the starting CR, it has three main problems: *(i)* The step size and its standard deviation must be defined by the user, requiring *a priori* knowledge of the PES. *(ii)* Steps that are too large can lead to unconnected paths or miss SPs close to the starting minimum. *(iii)* Not doing intermediate relaxations can lead to non-physical configurations for which the forces are abnormally high. Although the second strategy can be applied in ARTn with appropriate preprocessing, the first is recommended. This discussion concerning the amplitude of $\mathbf{V}_{\text{rinit}}$ in the starting CR can be generalized to the amplitude of \mathbf{V}_{CR} used to escape the other CR.

D. Choice of $\mathbf{V}_{\text{rinit}}$ direction

For activated processes in solid-state systems, connected SPs are generally localized in a region of the configuration space in which a small number of atoms (a dozen or so) are displaced from the initial minimum: the ones for which the chemical bond is broken or created and their neighbors. Therefore, the optimal initial random deformation $\mathbf{V}_{\text{rinit}}$ consists of moving one atom of interest and its nearest neighbors (up to the first or second shell) in a random direction. When too many atoms are

displaced, multiple disconnected valleys can arise, representing competing activated mechanisms. This can result in an oversampling of a small set of SPs, the ones that are the easiest to reach, *i.e.* that have the largest attracted region, or in the generation of unconnected SPs, as the Hessian spectrum becomes denser, resulting in disconnected pathways. An example of such disconnected pathway would be a simultaneous diffusion of atoms at two different locations in the material.

It is formally possible to fully sample the PES around a basin by systemically choosing deformation directions from a set of equidistant points on the hypersphere of possible deformations, as performed by a pre-processing routine of the dimer code³³. Although this results in a faster sampling of SPs in low dimensions (see adaptive_HSphere results on the OptBench saddle search benchmark test³⁴), this approach becomes unmanageable in higher dimensions, where the number of equidistant points needed to ensure proper sampling becomes prohibitively large.

VIII. CONCLUSION

The identification of saddle points on potential energy surfaces (PES) is crucial for the characterization of the kinetics of chemical reactions and of the evolution of materials. Over the last 50 years, multiple algorithms have been developed to find these saddle points, based on a relatively simplistic view of the features of PES. Building on a 2D model, we show that these methods underestimated the importance of convex regions (CR) standing between the local minimum and the connected saddle point, greatly reducing their ability to find saddle points of the PES. This causes a general under-sampling of connected saddle points, while increasing the overall computational costs of the open-ended methods.

The much-improved understanding of the PES features leads us to explore various solutions to reach the different valleys beyond convex regions applying ARTn to an amorphous silicon model. Pushing the system through a CR along a modified direction, given by the double random vector $\mathbf{V}_{\text{rinit}} + \mathbf{V}_{\text{new}}$, we show that *(i)* we can identify up to 10 times more unique connected saddle points than with the standard ARTn, while *(ii)* reducing the proportion of failed explorations to well below one percent.

These results are of critical importance for the characterization of the PES for complex materials as well as the use of open-ended methods such as ARTn within kinetic algorithms^{35,36}. Further work will focus on these directions.

CODE AND DATA AVAILABILITY

The algorithm presented has been implemented in the pARTn code²⁵ and is freely downloadable from the

GitLab repository³⁷. It can be plugged into force engines using empirical potentials (LAMMPS²⁹) and density functional theory (presently Quantum Espresso³⁸, soon Siesta³⁹, VASP⁴⁰ and abinit⁴¹). The input data used in the amorphous silicon study is available in the examples directory of the pARTn GitLab repository.

SUPPLEMENTARY MATERIALS

Supplementary Materials are available for this article. There are two mp4 video files, one gnuplot script, and one pdf with detailed description of each.

AUTHOR DECLARATIONS

Conflict of interest

The authors have no conflicts to disclose.

ACKNOWLEDGMENTS

We thank Hannes Jónsson for fruitful discussions. This work was carried out using HPC resources from GENCI-A0090911942 and CALMIP-Grant P1555. NM acknowledges partial support from NSERC through a Discovery Grant. The authors are active members of the Multiscale And Multi-Model ApproacheS for Materials In Applied Science consortium (MAMMASMIAS consortium), and acknowledge the efforts of the consortium in fostering scientific collaboration.

Appendix A: Example of saddle points reached by crossing several convex regions

It has been shown in Fig. 6 that the number of CRs that need to be crossed to reach a SP can be particularly high. We show here that this observation is also found in lower dimensions, namely on a 2D toy-model PES with the same analytical function as in Sec. III, albeit in a different region. It is shown on Fig. 9, where the VtoPH are all inside some CR, and additionally, the CR contain three forbidden crossings. The first forbidden crossing makes the SP marked in white unreachable without adding a random part to \mathbf{V}_{CR} .

Appendix B: Example of valleys branching without forbidden crossing

Branching/bifurcation usually occurs far from the initial minimum, where the structural symmetries (symmetries of the PES) are generally broken. This means that the lowest eigenvalue is generally not degenerate and, consequently, that there is only one single eigenvector




FIG. 9. Same PES and colors as Figs. 1 and 3, but region between $x \in [17.1; 20.2]$ and $y \in [25.25; 28.45]$. Two CRs need to be crossed to reach the magenta SP from the minimum.

\mathbf{V}_{min} . This idea becomes obviously false in highly symmetric systems, since the lowest eigenvalue in that case is degenerate, and thus multiple eigenvectors \mathbf{V}_{min} exist. If the PES keeps its symmetry out of the minima, which is typically the case during the diffusion of atoms in crystalline systems, then the valleys are mathematically allowed to cross each other at a single point.

An example of a symmetric case is shown in Fig. 10 which represents the diffusion of a substitutional Cu atom in a 128-atom bulk aluminum simulation box. SPs have been found using ARTn in its plugin version with Quantum Espresso V7.3³⁸. Total energy and forces have been calculated within density functional theory in the Local Density Approximation, using pseudo-potentials and a plane waves approach. The Brillouin zone is sampled at zone center only and a 37 Ry cutoff energy is used to restrict the number of plane waves summed to form the wave functions. This example has been chosen because one of the valleys (pink balls in Fig. 10a) coming from the starting minimum is suddenly split into four different valleys at the branching point (Fig. 10b), each of them leading to a different saddle point (only one is represented in Fig. 10c).

Interestingly, the lowest eigenvalue at the branching point (Fig. 10b) is still negative, as the valley branching

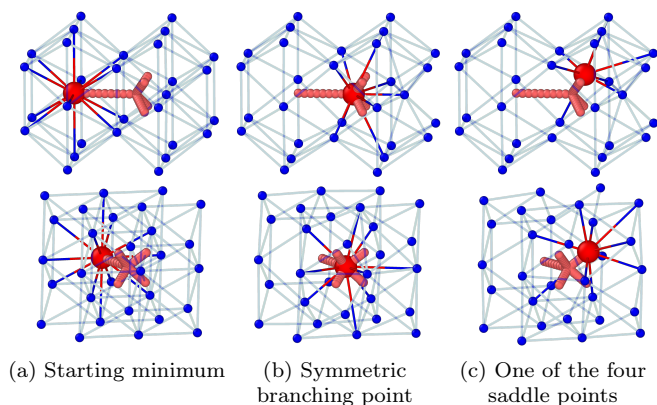


FIG. 10. Diffusion of a substitutional Cu atom (red ball) in cubic Aluminium (blue balls) by a pushing mechanism. The small pink balls represent the positions of the Cu atom in the valleys linking the starting minimum to the 4 saddle points. (10a): The starting Cu position is a substitutional site. (10c): Only one saddle point is represented. At the branching point (Fig. 10b), which is not a saddle point, the lowest eigenvalue is negative and four times degenerate. The top and bottom panels represent the same structure from two different points of view.

does not necessarily induce a convex region. This means that this point cannot be detected by a change of inflection but only by a symmetry analysis. In this case, the symmetry is a rotation of order 4 around the axis defined by the starting valley. This symmetry induces the degeneracy of lowest eigenvalues at the branching point, and that any linear combination of their eigenvector is an eigenvector. Hopefully, when one eigenvector \mathbf{V}_{\min} has been chosen by the Lanczos algorithm and the system has been pushed along it, the system is no longer at the branching point, the PES symmetries are broken, and the lowest eigenvalue is then uniquely defined.

However, during a PES exploration, a problem occurs when the system reaches such single point branchings. In fact, after the push along an eigenvector, during the orthogonal relaxation step, the system can potentially fall down into one of the 3 other surrounding valleys or around the branching point where the eigenvectors are still quite degenerate. Hence, at the next Lanczos step, the eigenvector \mathbf{V}_{\min} will be modified to be in another valley. This eigenvector modification will occur again and again at each perpendicular relaxation, such that the system will oscillate around the branching point by constantly changing the valley it follows. This does not occur when the branching is a forbidden crossing because the valleys are far enough from each other, and the perpendicular relaxation cannot be large enough to permit the system to change the valley. To avoid this phenomenon, a simple solution is to drastically reduce the number of orthogonal relaxation iterations that are performed after the first few pushes following the detection of a symmetric point. With this method, the orthogonal relaxation

becomes efficient only when the system is far enough from the branching point. Another solution is to increase the size of the push along \mathbf{V}_{\min} when a symmetry is detected, with the risk of crossing the basin boundaries.

- ¹G. H. Vineyard, *J. Phys. Chem. Sol.* **3**, 121 (1957).
- ²G. Mills, H. Jonsson, and G. K. Schenter, *Surf. Sci.* **324**, 305 (1995).
- ³G. Henkelman, B. P. Uberuaga, and H. Jonsson, *J. Chem. Phys.* **113**, 9901 (2000).
- ⁴E. L. Kolsberg, M. N. Groves, and B. Hammer, *J. Chem. Phys.* **145**, 094107 (2016).
- ⁵N. R. Mathiesen, H. Jónsson, T. Vegge, and J. M. G. Lastra, *J. Chem. Theory Comput.* **15**, 32155 (2019).
- ⁶A. Jay, C. Huet, N. Salles, M. Gunde, L. Martin-Samos, N. Richard, G. Landa, V. Goiffon, S. De Gironcoli, A. Hémercyck, and N. Mousseau, *J. Chem. Theory Comput.* **16**, 6726 (2020).
- ⁷S. Smidstrup, A. Pedersen, K. Stokbro, and H. Jónsson, *J. Chem. Phys.* **140**, 214106 (2014).
- ⁸G. T. Barkema and N. Mousseau, *Phys. Rev. Lett.* **77**, 4358 (1996).
- ⁹R. Malek and N. Mousseau, *Phys. Rev. E* **62**, 7723 (2000).
- ¹⁰N. Mousseau, L. K. Béland, P. Brommer, J. F. Joly, F. El-Mellouhi, E. Machado-Charry, M. C. Marinica, and P. Pochet, *J. At. Mol. Opt. Phys.* **2012**, 925278 (2012).
- ¹¹J. Doye and D. Wales, *Z. Phys. D - At. Mol. and Clust.* **40**, 194 (1997).
- ¹²G. Henkelman and H. Jónsson, *J. Chem. Phys.* **111**, 7010 (1999).
- ¹³S. T. Chill, M. Welborn, R. Terrell, L. Zhang, J.-C. Berthet, A. Pedersen, H. Jónsson, and G. Henkelman, *Mod. Sim. Mat. Sci. Eng.* **22**, 055002 (2014).
- ¹⁴W. Quapp, M. Hirsch, O. Imig, and D. Heidrich, *J. Comp. Chem.* **19**, 1087 (1998).
- ¹⁵R. A. Olsen, G. J. Kroes, G. Henkelman, A. Arnaldsson, and H. Jónsson, *J. Phys. Chem.* **121**, 9776 (2004).
- ¹⁶D. Wales, *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses* (Cambridge University Press, 2001).
- ¹⁷D. K. Hoffman, R. S. Nord, and K. Ruedenberg, *Theor. Chem. Acc.* **69**, 265 (1986).
- ¹⁸W. Quapp, M. Hirsch, and D. Heidrich, *Theor. Chem. Acc.* **112**, 40 (2004).
- ¹⁹R. Palmeiro and O. Castaño, *Inter. J. Quant. Chem.* **107**, 2687 (2007).
- ²⁰J. von Neumann and E. Wigner, *Phys. Zeit* **30**, 467 (1929).
- ²¹A. Seyranian, O. Kirillov, and A. Mailybaev, *J. Phys. A: Math. Gen.* **38**, 1723 (2005).
- ²²W. Quapp, M. Hirsch, and D. Heidrich, *Theor. Chem. Acc.* **100**, 285 (1998).
- ²³A. Jay, M. Gunde, N. Salles, M. Poberžnik, L. Martin-Samos, N. Richard, S. deGironcoli, N. Mousseau, and A. Hémercyck, *Comp. Mat. Sci.* **209**, 111363 (2022).
- ²⁴A. Jay, N. Mousseau, N. Salles, M. Gunde, M. Poberžnik, M. Brut, L. Martin-Samos, N. Richard, and A. Hémercyck, *Tech. Ing.* (2023).
- ²⁵M. Poberžnik, M. Gunde, N. Salles, A. Jay, A. Hemeryck, N. Richard, N. Mousseau, and L. Martin-Samos, *Comp. Phys. Comm.* **295**, 108961 (2024).
- ²⁶C. Lanczos, *J. Res. Nat. Bur. St.* **45**, 255 (1950).
- ²⁷N. T. Bijoy K. Dey, Taylor Kapsch and R. Fick, *Mol. Phys.* **112**, 937 (2014).
- ²⁸A. Pedersen and M. Luise, *J. Chem. Phys.* **141**, 024109 (2014).
- ²⁹A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, *Comp. Phys. Comm.* **271**, 108171 (2022).
- ³⁰R. Lot, L. Martin-Samos, S. de Gironcoli, and A. Hémercyck, *IEEE 16th Nanotech. Mat. Dev. Conf.* , 1 (2021).
- ³¹A. Jay, M. Raine, N. Richard, N. Mousseau, V. Goiffon,

- A. Hémerlyck, and P. Magnan, *IEEE Trans. Nucl. Sc.* **64**, 141 (2017).
- ³²R. Vink, G. Barkema, W. van der Weg, and N. Mousseau, *J. Non-Cryst. Sol.* **282**, 248 (2001).
- ³³M. Plasencia Gutiérrez, C. Argáez, and H. Jónsson, *J. Chem. Theory Comput.* **13**, 125 (2017).
- ³⁴<https://optbench.org/saddle-search.html>.
- ³⁵F. El-Mellouhi, N. Mousseau, and L. Lewis, *Phys. Rev. B* **78**, 153202 (2008).
- ³⁶L. Béland, P. Brommer, F. El-Mellouhi, J. Joly, and N. Mousseau, *Phys. Rev. E* **84**, 046704 (2011).
- ³⁷<https://gitlab.com/mammasmias/artn-plugin/>.
- ³⁸P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M. B. Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, N. Colonna, I. Carnimeo, A. D. Corso, S. de Gironcoli, P. Delugas, R. A. DiStasio, A. Ferretti, A. Floris, G. Fratesi, G. Fugallo, R. Gebauer, U. Gerstmann, F. Giustino, T. Gorni, J. Jia, M. Kawamura, H.-Y. Ko, A. Kokalj, E. Küçükbenli, M. Lazzeri, M. Marsili, N. Marzari, F. Mauri, N. L. Nguyen, H.-V. Nguyen, A. O. de-la Roza, L. Paulatto, S. Poncé, D. Rocca, R. Sabatini, B. Santra, M. Schlipf, A. P. Seitsonen, A. Smogunov, I. Timrov, T. Thonhauser, P. Umari, N. Vast, X. Wu, and S. Baroni, *J. Phys.: Cond. Matter.* **29**, 465901 (2017).
- ³⁹A. García, N. Papior, A. Akhtar, E. Artacho, V. Blum, E. Bosoni, P. Brandimarte, M. Brandbyge, J. I. Cerdá, F. Corsetti, R. Cuadrado, V. Dikan, J. Ferrer, J. Gale, P. García-Fernández, V. M. García-Suárez, S. García, G. Huhs, S. Illera, R. Korytár, P. Koval, I. Lebedeva, L. Lin, P. López-Tarifa, S. G. Mayo, S. Mohr, P. Ordejón, A. Postnikov, Y. Pouillon, M. Pruneda, R. Robles, D. Sánchez-Portal, J. M. Soler, R. Ullah, V. W.-z. Yu, and J. Junquera, *J. Chem. Phys.* **152**, 204108 (2020).
- ⁴⁰G. Kresse and D. Joubert, *Phys. Rev. B* **59**, 1758 (1999).
- ⁴¹X. Gonze, B. Amadon, G. Antonius, F. Arnardi, L. Baguet, J.-M. Beuken, J. Bieder, F. Bottin, J. Bouchet, E. Bousquet, N. Brouwer, F. Bruneval, G. Brunin, T. Cavignac, J.-B. Charraud, W. Chen, M. Côté, S. Cottenier, J. Denier, G. Geneste, P. Ghosez, M. Giantomassi, Y. Gillet, O. Gingras, D. R. Hamann, G. Hautier, X. He, N. Helbig, N. Holzwarth, Y. Jia, F. Jollet, W. Lafargue-Dit-Hauret, K. Lejaeghere, M. A. Marques, A. Martin, C. Martins, H. P. Miranda, F. Naccarato, K. Persson, G. Petretto, V. Planes, Y. Pouillon, S. Prokhorenko, F. Ricci, G.-M. Rignanese, A. H. Romero, M. M. Schmitt, M. Torrent, M. J. van Setten, B. Van Troeye, M. J. Verstraete, G. Zerah, and J. W. Zwanziger, *Comp. Phys. Comm.* **248**, 107042 (2020).