Tetrahedron-tiling method for crystal structure prediction

Qi-Jun Hong,^{1,*} Joseph Yasi,² and Axel van de Walle¹

¹*School of Engineering, Brown University, Providence, Rhode Island 02912, USA*

²*The MathWorks, Inc., 3 Apple Hill Drive, Natick, Massachusetts 01760, USA*

(Received 19 February 2017; published 17 July 2017)

Reliable and robust methods of predicting the crystal structure of a compound, based only on its chemical composition, is crucial to the study of materials and their applications. Despite considerable ongoing research efforts, crystal structure prediction remains a challenging problem that demands large computational resources. Here we propose an efficient approach for first-principles crystal structure prediction. The new method explores and finds crystal structures by tiling together elementary tetrahedra that are energetically favorable and geometrically matching each other. This approach has three distinguishing features: a favorable building unit, an efficient calculation of local energy, and a stochastic Monte Carlo simulation of crystal growth. By applying the method to the crystal structure prediction of various materials, we demonstrate its validity and potential as a promising alternative to current methods.

DOI: [10.1103/PhysRevMaterials.1.020801](https://doi.org/10.1103/PhysRevMaterials.1.020801)

Crystal structure prediction (CSP) [1–3] has been a topic of great interest with both fundamental importance in physical science—an ambitious goal to predict the structure of a crystal from the knowledge of its chemical composition—and enormous potential in modern materials design and discovery ranging from petrochemical industry [4] and electrical energy storage [5] to the study of phase diagrams [6]. This recognition has led to the development of a number of ingenious computational approaches, including simulated annealing [7,8], basin hopping [9], evolutionary algorithms [10–14], random structure searching [15], high-throughput data mining [16,17], etc. Over the past decade, contemporary CSP methods started to develop genuine predictive capability and they keep improving. However, structure prediction remains a challenging problem that demands large computational resources. To address this issue, we propose an alternative CSP approach, by which new crystal structures are predicted by tiling from energetically favorable and geometrically matching elementary tetrahedra. This strategy proves advantageous because (i) tetrahedra energies can be efficiently determined from a relatively small number of quantum mechanical calculations and (ii) tetrahedra form natural building blocks for arbitrary crystal structures.

We propose a tetrahedron-tiling method, which we illustrate with the flowchart shown in Fig. 1. Starting from a small training set of common crystal structures, we first calculate the local energies of the tetrahedra from these structures, based on local energy density method (EDM) [18,19] in density functional theory (DFT) [20–22] and the Delaunay triangulation [23]. These data points of known tetrahedra enable us to interpolate, at a very low cost, an empirical energy landscape for any tetrahedron with general coordinates. Based on this tetrahedron energy landscape, we run Monte

Carlo (MC) simulations to tile together tetrahedra that are favorable in energy and match each other in geometry. The simulations generate a series of stable crystal structures, which are further added to the training set to provide more data points on new tetrahedra, systematically improving the empirical tetrahedron energy landscape. After sufficiently many cycles, CSP is achieved since all important tetrahedra, hence crystal structures as well, are captured and included in the training set. The set of candidate crystal structures generated in this fashion is then further screened via DFT, to yield the final predicted low-energy structure(s) with full DFT accuracy.

This approach offers two distinct advantages.

First, slicing crossover in existing evolutionary algorithms (e.g., Ref. [11]) inevitably involves undesirable lattice and structure mismatches. When slicing crossover combines two parental structures with high fitness ranking to create offsprings, one key idea is to put together two locally favorable structures to increase the chances of producing a better chemistry. Indeed, structural formation energies are typically dominated by short-range interactions, and this wisdom is widely adopted in CSP methods. Despite being an ingenious idea, slicing crossover sometimes creates mismatches of lattice and chemical bonds, and hence the slicing plane may become an undesirable interface with an associated energy penalty. We note that there is a better way to combine local structures naturally: a deformable tetrahedron is an ideal tiling unit. Geometrically, a tetrahedron is the simplest and most elementary form of a polyhedron in three-dimensional space. Any set of three-dimensional points in general position can be uniquely decomposed into tetrahedra by the Delaunay triangulation [23]. Energetically, a tetrahedron has four vertices, so it accounts for local four-body interactions, which naturally incorporates key chemical parameters, such as bond length, bond angle, dihedral angle, solid angle, etc. Furthermore, recent development of local EDM [18,19] in DFT enables us to compute a local effective gauge-invariant energy for each atom in a crystal (see Supplemental Material S1 [24]) and assign a well-defined energy to each tetrahedron unit via a geometric construction without undergoing a large-scale fitting exercise typically associated with the construction of effective interatomic potential energy model. These capabilities open up

*qhong@alumni.caltech.edu

FIG. 1. A flowchart showing the mechanism of crystal structure prediction we propose.

a new avenue to CSP: it becomes possible to build a structure by tiling together tetrahedra that are locally favorable and geometrically match with each other.

Second, contemporary CSP approaches are not sufficiently efficient in the context of first-principles calculations, and this becomes an obstacle to their application in high-throughput materials screening and discovery. Historically, the realization that a brute-force systematic search through all possible structures is infeasible due to the exponential complexity of the optimization problem leads to the wide adoption of more pragmatic techniques such as evolutionary algorithm and simulated annealing. While this strategy systematically improves the expected search time and increases the probability of finding the global minimum, a trial and error approach with an energy model as expensive as DFT is by no means cost-effective. For instance, evolutionary algorithms from fully first-principles demand a large amount of energy evaluations, which renders it prohibitively expensive for complex structures [13,25]. For simulated annealing, running *ab initio* molecular dynamics (MD) or MC is already a heavy task, not to mention the requirement of slow cooling and hence a long trajectory. By contrast, we devise an approach based on an effective energy model of tetrahedron, which requires only a small set of DFT calculations performed before hand. This property makes it highly efficient in the exploration of possible tetrahedron tiling.

An important aspect of the method is the mapping of a solid structure to a set of tetrahedra. In principle, a solid structure $\mathbf{R}^m$ with general coordinates ($\mathbf{R}^m$ is a $3 \times m$ matrix of atomic coordinates, m is the number of atoms) can be uniquely decomposed into tetrahedra through the Delaunay triangulation [23]. (Delaunay triangulation is a robust method widely used in scientific computing. See Supplemental Material S2 for details [24].) The tetrahedra serve as elementary building blocks that comprise the solid structure as a whole. After the

Delaunay triangulation, the solid structure $\mathbf{R}^m$ is represented by n tetrahedra $\mathbb{R} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_i, \dots, \mathbf{r}_n\}$. We therefore write the total energy of the solid $\mathbf{R}^m$ as a sum of the local energies of tetrahedra $\mathbb{R}$,

$$\mathbf{R}^m \rightarrow \mathbb{R} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_i, \dots, \mathbf{r}_n\}, \quad (1)$$

$$E_{\text{tot}}(\mathbf{R}^m) \approx \sum_{i=1}^n E(\mathbf{r}_i), \quad E(\mathbf{r}_i) = \epsilon(\mathbf{r}_i) \Omega_s(\mathbf{r}_i) / 4\pi, \quad (2)$$

where $E(\mathbf{r}_i)$ is the local energy of a tetrahedron $\mathbf{r}_i$, $\epsilon(\mathbf{r}_i)$ is the corresponding specific energy (in eV/atom), and $\Omega_s(\mathbf{r}_i)$ is the sum of its four solid angles, as $\Omega_s(\mathbf{r}_i)/4\pi$ accounts for the fractional number of atoms within a tetrahedron. The tetrahedron energy is equal to the sum of atom-specific energies found via EDM, each weighted by the solid angle that expresses the fraction of each atom actually lying inside the tetrahedron. For multicomponent systems, composition is controlled by changing chemical potentials of elements. We add a term of the form $-\mu^T x$ to the Hamiltonian (where μ and x are, respectively, the vectors of chemical potentials and compositions) and work with a grandcanonical ensemble (i.e., possible moves in our MC simulations include atom deletion or atomic species changes).

Converting a bulk solid structure into a set of *independent* tetrahedra greatly simplifies the problem. While the former has $3m$ degrees of freedom (or $3m - 6$, if excluding translation and rotation), each tetrahedron of the latter has only six, drastically reducing the complexity. In particular, it becomes feasible to explore the full energy landscape of tetrahedron $\epsilon(\mathbf{r})$, given the low dimensions.

As a key component of the method, constructing this energy landscape enables us to evaluate the specific energy $\epsilon(\mathbf{r}_i)$ of any general tetrahedron $\mathbf{r}_i$. This is achieved in two steps, (1) training, i.e., rigorous DFT-EDM calculations of ϵ for a group of common structures (which form a training set), and (2) interpolation for a general tetrahedron based on the known tetrahedra obtained in the training set, i.e.,

$$\epsilon(\mathbf{r}) = \sum_{k=1}^N \epsilon(\mathbf{r}_k^0) w(\mathbf{r} - \mathbf{r}_k^0) / \sum_{k=1}^N w(\mathbf{r} - \mathbf{r}_k^0), \quad (3)$$

where w is a weighting function (that takes its maximum value at $|\mathbf{r} - \mathbf{r}_k^0| = 0$ and decays smoothly as $|\mathbf{r} - \mathbf{r}_k^0| \rightarrow \infty$), $\mathbf{r}$ is a general tetrahedron, $\mathbf{r}_k^0$ is a known tetrahedron, and N is the total number of known tetrahedra in the training set (see Supplemental Material S3 [24]).

After building a full energy landscape of tetrahedra $\epsilon(\mathbf{r})$, we can estimate energy of any structure base on $\epsilon(\mathbf{r})$. In detail, an unknown structure is first decomposed into elementary tetrahedra, whose specific energies are evaluated from $\epsilon(\mathbf{r})$. The energies of individual tetrahedra add up to the total energy of the structure according to Eqs. (1) and (2). This capability is employed to simulate tetrahedra tiling and crystal growth, with an ultimate aim to achieve CSP. The tetrahedron tiling simulation utilizes Monte Carlo methods and includes both addition and removal of an atom, as well as structural optimization (see Supplemental Material S4 [24]). We note that our tetrahedra energies represent energies in bulk material. So even though the growing cluster has a surface, there is actually

no surface energy penalty associated with it. This is a desirable property, because the surface will not bias the search for the most stable bulk structure.

The tetrahedron energy landscape effectively serves as an empirical potential in tetrahedron tiling simulations. Tetrahedra contain crucial geometric information such as bond length, bond angle, etc. Indeed, they take account of four-body interactions among first nearest neighbors, as each tetrahedron connects a set of four neighbors. Hence, these tetrahedra in principle dominate energetic properties of the solid structure. Replacing total energy of a structure by a sum of tetrahedra's local energies, to be rigorous, introduces an approximation, as this process attempts to capture real many-body interactions through local energies derived from an empirical tetrahedron energy landscape. Only interaction among the first few nearest neighbors are taken into account and long-range interactions are truncated. The real total energy will inevitably differ from the sum. For instance, the mapping from a structure $\mathbf{R}^m$ to its tetrahedron set $\mathbf{r}$ is not injective, i.e., two structures may correspond to the same set of tetrahedra, and hence two identical tetrahedra could have different local energies in different structures. Nevertheless, we note that our interpolation scheme actually handles this by smoothing: this tetrahedron is then assigned the average value for all identical tetrahedra. Furthermore, the final relative stability will be determined by DFT, while the empirical energy landscape $\epsilon(\mathbf{r})$ is employed only to narrow the search down to a few candidate stable structures. While this approach of tetrahedron energy density evaluation appears similar to interatomic potential construction and fitting, the latter process, to achieve comparable accuracy of four-body interaction, would typically require (1) specifically developed analytic functional forms and (2) an enormous number of structures in the training set, which both demand heavy human input and computer resources. The EDM approach hence far outperforms it in terms of simplicity and computer cost.

We note that there are two remaining hurdles, and we address them using robust strategies.

The first problem lies in the process of tetrahedron tiling. Since the tetrahedron energy landscape $\epsilon(\mathbf{r})$ is an approximation to DFT interactions, the structure built by tetrahedron tiling is not necessarily the most stable one in the DFT aspect. Furthermore, the simulation may undesirably generate a known structure that is already in the training set, thus achieving little more. This problem is avoided via a simple but powerful exploration algorithm: we ask the simulation to produce a series of distinct stable structures, and let DFT decide their final relative stability. With this strategy, tetrahedron tiling is less likely to miss a stable structure. To find a structure that is less stable according to the effective energy model, we adopt an idea that discourages the simulation to explore structures already found. The idea is similar to the one employed in metadynamics [26] (However, a collective variable is not required here, since the tetrahedra are already described by a low-dimensional vector and we do not need to recover the potential energy surface. In other words, this method does not suffer from the drawback of metadynamics). For each stable structure found, positive Gaussian potentials, located at the tetrahedra that compose the structure, are added to the tetrahedron energy landscape $\epsilon(\mathbf{r})$ (see Supplemental

TABLE I. A summary of applications of the tetrahedron tiling method.

Materials	Training set	Structures predicted
Ti	fcc, bcc	<i>hP3</i> (omega)
Na	fcc, bcc, <i>oP8</i>	<i>cI16</i>
NaCl	<i>cP2</i> (CsCl)	<i>cF8</i> (rocksalt)
Ir-W	bcc, fcc for Ir, W, and IrW	IrW: <i>oP4</i> (<i>B19</i>) Ir ₃ W: <i>hP8</i>

Material S5 [24]). When relaunching tetrahedron tiling with the new perturbed tetrahedron energy landscape $\epsilon(\mathbf{r})$, these Gaussian potentials artificially destabilize the tetrahedra, thus discouraging the system to visit known structures. Hence, the system tends to explore a broader landscape and to build new stable structures.

The other problem is associated with the completeness of the training set. If the number of structures in the training set is very limited, the known tetrahedra (data points) may be too few to give an effective interpolation, i.e., data points are too far away from each other and important data (tetrahedra) are missing, so that the tetrahedron energy landscape is not fully captured. We introduce an iterative strategy to systematically improve the training set and the energy landscape. We note that EDM calculations and tetrahedron tiling simulations can be carried out in an iterative and complementary manner. While EDM calculations provide energetic input to tetrahedron tiling simulations, new crystal structures predicted by tetrahedron tiling, in return, serve as new samples, which augment the training set. As the number of loops increases, the accuracy of the tetrahedron energy field gradually improves and is better able to describe the system energetics and ground states.

To illustrate the validity of this approach, we provide several examples of application of the method to the crystal structure prediction of titanium, sodium (at 120 GPa), sodium chloride, and iridium-tungsten alloys. Titanium, a simple metal with three major allotropes, allows us to quickly test out the new idea, while sodium has a wide variety of interesting complex crystal structures at high pressure [27] and is thus representative of challenging CSP problems one may encounter in practice. Sodium chloride, a typical ionic compound, serves as a prototype of multicomponent system with long-range interaction, and the success in this material suggests the method's potential in this respect. Finally, we apply the method to iridium-tungsten (Ir-W) alloys, which combine elements that each favors a different lattice: Ir is face-centered cubic (fcc) and W is body-centered cubic (bcc). This test case most reflects practical settings where the structures of interest are multicomponent ordered alloys. We summarize in Table I and Fig. 2 the crystal structures predicted in these systems. We note that these structures are deliberately excluded from training sets, and finding them through the tetrahedron tiling method most directly demonstrates the method's "out-of-sample" predictive ability.

We first perform DFT-EDM calculations on a training set and construct a tetrahedron energy landscape for each material. The DFT calculations are performed with the Vienna *ab initio* simulation package (VASP) code [28] implementing the projector-augmented-wave (PAW) [29] method, with the

FIG. 2. Crystal structures predicted by tetrahedron tiling method. (a) Ti omega, (b) Na *cI16* (a perturbation of bcc phase), (c) NaCl rocksalt (yellow, Na; green, Cl), (d) IrW *oP4* (yellow, Ir; gray, W; this structure is called *B19*), and (e) Ir₃W *hP8*. These structures are deliberately excluded from training sets.

generalized gradient approximation (GGA) for the exchange-correlation energy, in the form known as Perdew-Burke-Ernzerhof (PBE) [30]. Based on the tetrahedron energy, we then run tetrahedron-tiling simulations to grow crystal nuclei. As the size of the crystal increases, we detect periodicity and extract periodic unit cells, which we further optimize using DFT and run structure analysis. We sort these structures by energy and include new structures in the training set for the next iteration. In the end, structure search is completed if no new stable structures are produced (see Supplemental Material and movies [24]).

Tetrahedron-tiling simulations find new crystal structures in three progressive stages. Based on the tetrahedron energy function, we first expect tetrahedron tiling simulation to tile and recover, self-consistently, the most stable structure so far in the training set. Even though this exercise does not constitute a CSP, per se, it is instructive to illustrate the construction process, and it demonstrates the method's capability to locate a minimum and tile a structure effectively. In the second stage, we modify the energy landscape and penalize the tetrahedra belonging to the minimum, making it less favorable. This strategy, destabilizing the minimum found in the previous step, guides the simulation to other minima and iteratively generates a list of stable structures other than the tiled minimum. This capability allows the simulation to explore and build a diverse range of structures. In the final stage, we scrutinize the predictive power of the approach by examining whether the simulation produces new unknown structures that are even more stable than the current minimum. Our tetrahedron-tiling simulations confirm this capability. Figure 3 shows the periodic cells captured in the polymorphic crystal nuclei from tetrahedron tiling simulations. In addition to known structures, the simulation predicts new global minima (shown in red), which were deliberately excluded from the training sets.

We highlight two interesting features of tetrahedron tiling that reflect the method's robustness. First, all types of tetrahedra are accounted for, i.e., any tetrahedron of any shape in principle can be added to the surface of a growing crystal nucleus during tetrahedron tiling. This feature makes tiling highly flexible, and the simulation assesses and compares the stabilities of a tremendous number of possible combinations in the stochastic exploration. Indeed, tetrahedra from other phases transiently appear, but they give way to the global minimum in the end, as shown in Fig. 4(a). Second, the simulation faces a common challenge in tiling: the most favorable structure is built not necessarily from the most stable tetrahedra. Instead, it requires collaboration among a group of matching tetrahedra, aiming to minimize the total energy. If

FIG. 3. New global minima (red) predicted by the tetrahedron tiling method. These global minima are deliberately excluded from the training sets, so they are *a priori* unknown. After tetrahedron tiling simulation builds a structure nucleus (a cluster of atoms, which is usually polymorphic), we extract possible periodic cells (label them with structure number, *x* axis), and we further optimize them with DFT to nearby local minima. This figure shows the energetics of the periodic cells after DFT structure optimization, with the number of atoms in each periodic cell labeled above each bar. The calculations reveal new global minima (red), which demonstrates the method's capability to predict new stable structures.

two tetrahedra do not match each other, they will not form a stable crystal even though each individual tetrahedron is highly stable. One typical example is the Ti omega phase, which requires a particular way to put together two different types of tetrahedra. As shown in Fig. 4(a), while one tetrahedron (omega-1, in orange, see Supplemental Material Table 1 [24]) is the most stable tetrahedra in the training set, it cannot fill the whole space by itself due to structural limitations and hence requires other tetrahedra to collaborate with. It turns out that the chosen tetrahedron (omega-2, in red) is in fact energetically less stable than hcp and bcc, but it best matches omega-1 in geometry, filling the holes left in space and facilitating the stacking. Indeed, the transient appearance of the hcp tetrahedra (green and cyan) in Fig. 4(a) suggests that the simulation attempts to incorporate these stable tetrahedra into the crystal, but, due to structural mismatch, they merely serve as surface defects, poisoning local facet growth, before they finally give way to omega-2. The fact that the simulation builds the desired omega structure demonstrates that the approach has the foresight to account for the overall energy and it is capable of identifying the global minimum despite complexity.

We note that the tetrahedron tiling scheme requires little computational cost and thus is highly efficient, since expensive DFT-EDM calculations are carried out on a very limited number of structures either initially in the training set or found by tetrahedron tiling, and the rest of the energy field is approximated through interpolation. While the simulation of tetrahedron tiling demands a huge amount of energy evaluation on structures, these calculations are based on the empirical tetrahedron energy landscape, hence having a negligibly small cost. Indeed, the runtime of our algorithm is competitive compared to other CSP methods. Though accurate runtime will

FIG. 4. Tetrahedron composition during tetrahedron tiling simulations. The colored areas show the fractional compositions (compositions sum up to 1) of different types of tetrahedra in the process of tetrahedron tiling. The type is determined as the closest known tetrahedron. The major six known tetrahedra (see Supplemental Material Table 1 [24]) are marked with different colors. The size of the growing cluster is also plotted for reference, as the black (y axis on left) and gray (y axis on right) curves, respectively, represent the number of tetrahedra and atoms in the clusters. (a) With the omega phase in the training set, tetrahedron tiling simulation builds an omega phase. (b) Excluding the omega phase from the training set, the simulation builds a polymorphic nucleus, which leads to the discovery of the omega phase.

depend on specific material to study, we here provide a general estimate of computer cost. A tetrahedron tiling simulation typically completes in ~ 10 – 100 h on a single core. The simulation usually generates 10–100 periodic structures that require DFT optimization. We repeat this iteration (tetrahedron tiling simulation followed by DFT structure optimization) until the global minima are captured. In the systems we studied, the global minima are usually discovered within three iterations.

In the case of multicomponent systems, it can be fruitful to explore various possible ordering of the atomic species on the candidate structures found via the tetrahedron tiling method using methods specifically designed for handling ordering problems on a known lattice, such as cluster expansion methods [6,31]. In this combination, the two methods are perfectly complementary: tetrahedron tiling can predict the lattice geometry that the cluster expansion is unable to autonomously find while the cluster expansion can handle long-range interactions not accounted for in tetrahedron tiling. (In fact, we have tried this approach for the Ir-W system using the ATAT software [6,32,33] and the $hP8$ structure was rapidly confirmed as a ground state on the hcp lattice after about 30 small-cell DFT calculations.)

To summarize, we propose a tetrahedron tiling method as an alternative approach to current CSP methods. This approach employs tetrahedra to combine locally favorable structures into a crystal, a simple and natural way that outperforms slicing crossover. The method is highly cost-effective, with a low requirement of DFT calculations. We apply the approach to the crystal structure prediction of titanium, sodium (at 120 GPa), sodium chloride, and iridium-tungsten alloys. It is demonstrated that the approach is capable of (1) finding the global minimum with nonobvious ground states, (2) generating a series of stable structures with proper perturbation, and (3) predicting new structures outside the training set.

We thank Drs. Min Yu and Dallas Trinkle for providing EDM code for DFT-EDM calculations. This research was supported by NSF under Grant No. DMR-1154895, by ONR under Grants No. N00014-14-1-0055 and No. N00014-17-1-2202, and by Brown University through the use of the facilities at its Center for Computation and Visualization. This work uses the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation Grant No. ACI-1548562.

- [1] J. Maddox, Crystals from first principles, *Nature (London)* **335**, 201 (1988).
- [2] G. R. Desiraju, Cryptic crystallography, *Nat. Mater.* **1**, 77 (2002).
- [3] S. M. Woodley and R. Catlow, Crystal structure prediction from first principles, *Nat. Mater.* **7**, 937 (2008).
- [4] J. K. Nørskov, T. Bligaard, J. Rossmeisl, and C. H. Christensen, Towards the computational design of solid catalysts, *Nat. Chem.* **1**, 37 (2009).
- [5] M. M. Thackeray, C. Wolverton, and E. D. Isaacs, Electrical energy storage for transportation—approaching the limits of, and going beyond, lithium-ion batteries, *Energy Environ. Sci.* **5**, 7854 (2012).
- [6] A. van de Walle and G. Ceder, Automating first-principles phase diagram calculations, *J. Phase Equilib.* **23**, 348 (2002).
- [7] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
- [8] V. Cerny, Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm, *J. Optim. Theory Appl.* **45**, 41 (1985).
- [9] D. J. Wales and H. A. Scheraga, Global optimization of clusters, crystals, and biomolecules, *Science* **285**, 1368 (1999).
- [10] D. A. Coley, *An Introduction to Genetic Algorithms for Scientists and Engineers* (World Scientific, Singapore, 1999).
- [11] A. R. Oganov and C. W. Glass, Crystal structure prediction using *ab initio* evolutionary techniques: Principles and applications, *J. Chem. Phys.* **124**, 244704 (2006).
- [12] M. d'Avezac and A. Zunger, Identifying the minimum-energy atomic configuration on a lattice: Lamarckian twist on Darwinian evolution, *Phys. Rev. B* **78**, 064102 (2008).
- [13] Y. Wang, J. Lv, L. Zhu, and Y. Ma, Crystal structure prediction via particle-swarm optimization, *Phys. Rev. B* **82**, 094116 (2010).
- [14] W. Zhang, A. R. Oganov, A. F. Goncharov, Q. Zhu, S. E. Boulfelfel, A. O. Lyakhov, E. Stavrou, M. Somayazulu, V. B. Prakapenka, and Z. Konôpková, Unexpected stable stoichiometries of sodium chlorides, *Science* **342**, 1502 (2013).
- [15] C. J. Pickard and R. J. Needs, *Ab initio* random structure searching, *J. Phys.: Condens. Matter* **23**, 053201 (2011).

- [16] S. Curtarolo, D. Morgan, K. Persson, J. Rodgers, and G. Ceder, Predicting Crystal Structures with Data Mining of Quantum Calculations, *Phys. Rev. Lett.* **91**, 135503 (2003).
- [17] Y. Wang, S. Curtarolo, C. Jiang, R. Arroyave, T. Wang, G. Ceder, L.-Q. Chen, and Z.-K. Liu, *Ab initio* lattice stability in comparison with CALPHAD lattice stability, *Calphad* **28**, 79 (2004).
- [18] N. Chetty and R. M. Martin, First-principles energy density and its applications to selected polar surfaces, *Phys. Rev. B* **45**, 6074 (1992).
- [19] M. Yu, D. R. Trinkle, and R. M. Martin, Energy density in density functional theory: Application to crystalline defects and surfaces, *Phys. Rev. B* **83**, 115113 (2011).
- [20] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
- [21] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, *Phys. Rev.* **140**, A1133 (1965).
- [22] R. Jones and O. Gunnarsson, The density functional formalism, its applications and prospects, *Rev. Mod. Phys.* **61**, 689 (1989).
- [23] B. Delaunay, On the empty sphere (in the memory of Georges Voronoi), *Bull. USSR Acad. Sci. Class Math. Natural Sci.* **6**, 793 (1934).
- [24] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevMaterials.1.020801> for details on the energy density method, Delaunay triangulation, tetrahedron-tiling method and its applications.
- [25] Current limitations of USPEX, retrieved from <http://uspex.stonybrook.edu/uspex.html>. Accessed June 5, 2017.
- [26] A. Barducci, M. Bonomi, and M. Parrinello, Metadynamics, *WIREs Comput. Mol. Sci.* **1**, 826 (2011).
- [27] E. Gregoryanz, L. F. Lundegaard, M. I. McMahon, C. Guillaume, R. J. Nelmes, and M. Mezouar, Structural diversity of sodium, *Science* **320**, 1054 (2008).
- [28] G. Kresse and J. Furthmüller, Efficiency of *ab initio* total energy calculations for metals and semiconductors using a plane-wave basis set, *Comput. Mater. Sci.* **6**, 15 (1996).
- [29] P. Blöchl, Projector augmented-wave method, *Phys. Rev. B* **50**, 17953 (1994).
- [30] J. P. Perdew, K. Burke, and M. Ernzerhof, Generalized Gradient Approximation Made simple, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [31] J. M. Sanchez, F. Ducastelle, and D. Gratias, Generalized cluster description of multicomponent systems, *Physica A* **128**, 334 (1984).
- [32] A. van de Walle, M. Asta, and G. Ceder, The alloy theoretic automated toolkit: A user guide, *Calphad* **26**, 539 (2002).
- [33] A. van de Walle and M. Asta, Self-driven lattice-model Monte Carlo simulations of alloy thermodynamic properties and phase diagrams, *Model. Simul. Mater. Sci. Eng.* **10**, 521 (2002).