ARTICLE   OPEN

Evolutionary computing and machine learning for discovering of low-energy defect configurations

Marco Arrigoni [1]✉ and Georg K. H. Madsen [1]

Density functional theory (DFT) has become a standard tool for the study of point defects in materials. However, finding the most stable defective structures remains a very challenging task as it involves the solution of a multimodal optimization problem with a high-dimensional objective function. Hitherto, the approaches most commonly used to tackle this problem have been mostly empirical, heuristic, and/or based on domain knowledge. In this contribution, we describe an approach for exploring the potential energy surface (PES) based on the covariance matrix adaptation evolution strategy (CMA-ES) and supervised and unsupervised machine learning models. The resulting algorithm depends only on a limited set of physically interpretable hyperparameters and the approach offers a systematic way for finding low-energy configurations of isolated point defects in solids. We demonstrate its applicability on different systems and show its ability to find known low-energy structures and discover additional ones as well.

INTRODUCTION

Point defects are ubiquitous in materials and their presence is able to modify the system behavior considerably. The controlled inclusion of defects in semiconductor and insulator materials can be employed to tune electronic, optoelectronic, electrochemical, and catalytical properties, and is commonly exploited in technological applications. On the other hand, the presence of point defects is not always beneficial, but can also lead to a degradation of the performance and lifetime of a device[1–9]. The desire to control the properties of defective systems requires a detailed understanding of the interaction of point defects with the host material, which has spurred the development of specific characterization methods.

First-principles calculations, based mostly on density functional theory (DFT)[10], have become a standard tool in the study of point defects in solids. In addition to their predictive power, first-principles calculations are able to provide information on several properties of interest, such as the atomic-scale geometric structure, the electronic density and wavefunction, and the defect thermochemistry, quantities that cannot be easily obtained from experiments[11–13]. As the defect concentrations most relevant for technological applications are generally very low, the first-principles study of point defects in solids has mostly been aimed at the so called dilute limit, where defect–defect interactions are considered to be negligible. To perform such calculations, one commonly employs the supercell method, in which a sufficiently large portion of the host material is taken as the simulation cell and the point defect of interest is then introduced therein. In order to minimize the spurious defect–defect interactions arising from the application of periodic boundary conditions, relatively large supercells with around 100 atoms or more might be required.

The large simulation cells, combined with defect-induced effects that strongly alter the energy landscape, make the systematic search for low-energy minima of the potential energy surface (PES) a formidable task. The approach most commonly used to tackle this problem is based on domain knowledge: the most likely low-energy defect configurations are selected according to intuition and/or the results obtained on analogs systems. Once the defect is introduced in the chosen configuration, the atomic positions are relaxed to a minimum of the PES using gradient-based optimization algorithms. However, there is no guarantee that the found minimum corresponds to the most stable defect configuration. Furthermore, a more complete understanding of the low-energy minima in the PES is also sought after, as metastable defect configurations can be responsible for phenomena that cannot be attributed to the ground state[14,15]. Searches based on previous knowledge are evidently unsatisfactory as the curse of dimensionality will inevitably make unstructured heuristic methods perform poorly.

Finding the ground state and low-energy metastable structures of a point defect is a problem of multimodal optimization. Several algorithms have been proposed to tackle such problems with evolutionary algorithms (EAs) being amongst the most popular and successful ones. EAs are a family of metaheuristic algorithms whose design is inspired by the Darwinian theory of natural evolution: a set of candidate solutions is generated and iteratively evolved through the application of genetic operators that select, recombine, and mutate the candidate solutions in order to generate new ones with higher fitness[16,17]. Although many components of EAs depend on stochastic processes and convergence to the global minimum is not guaranteed, their usefulness in optimization problems has been demonstrated in several complex problems. Genetic algorithms (GAs), the oldest, and perhaps most famous, class of EAs were adopted early on in computational materials science[18] and more recently also in combination with first-principles DFT studies[19–25]. There is however only a limited number of studies applying GAs or related methods to solids containing point defects[26–29] and, to the best of out knowledge, the only studies that have tackled the problem of finding the low-energy configurations of isolated point defects in solids have so far been based on random searches[30–32].

In this contribution, we present a computational approach that aims at finding low-energy configurations of isolated point defects in solid materials. Our approach is based on the covariance matrix adaptation evolution strategy (CMA-ES) algorithm[33]. ESs are

[1]Institute of Materials Chemistry, TU Wien, Vienna, Austria. ✉email: marco.arrigoni@tuwien.ac.at

mainly employed for solving continuous-parameter optimization tasks and, although the distinction between the two classes has become more blurred in recent years, the main difference between ESs and GAs is the presence of self-adapting endogenous strategy parameters in the former. Such strategy parameters determine the distribution of the candidate solutions and are adapted during the evolutionary process according to the individuals' fitnesses. ESs possess some advantages that make them a promising alternative to the more popular GAs. In particular, the flexibility of GAs is endowed by the existence of a wide variety of genetic operators and while this might be a benefit for maximizing the performance of a specific algorithm for a specific problem, it might also require extensive fine-tuning, as each kind of genetic operator depends on some hyperparameters, whose optimal values generally have to be obtained heuristically. The choice of genetic operators in ESs is typically much more restricted facilitating the systematic optimization of the hyperparameters. In the present work we demonstrate how the original CMA-ES can be modified to suit DFT studies of point defects in the dilute limit, leading to an algorithm which depends on a limited set of physically interpretable hyperparameters.

Characteristic of all EAs is the generation of several structures before a converged solution is obtained, making methods able to make use of this information particularly useful. We propose two methods based on machine learning (ML) to do this. First of all, in order to discover additional low-energy minima, we propose a method employing an unsupervised learning post-processing step which is able to exploit the configurations explored by the EA and discover efficiently defective structures leading to low-energy minima which do not coincide with the solution found by the EA itself. Furthermore, we address one of the main drawbacks common to all EAs, namely the need of several evaluations of the fitness function before an optimal solution is found, which make their use in first-principles studies computationally very demanding. The recent development of accurate supervised ML methods[34–38] has allowed the use of ML metamodels with the ability of approximating accurately the DFT PES in conjunction with EAs[39,40]. We investigate the augmentation of our proposed EA by the inclusion of a ML metamodel based on Gaussian processes (GPs) in order to substantially reduce the computational burden.

The rest of the paper is organized as follows: in section Results we firstly illustrate all the components of the proposed computational approach in detail and analyzes their contribution to its overall performance. In particular, subsection Evolutionary Algorithm is concerned with the description of the EA, subsection Unsupervised Machine Learning Model describes the unsupervised post-processing step and subsection Supervised Machine Learning Model the ML regressor metamodel. The capabilities of the specific method outlined in each subsection is shown by considering intrinsic impurities in silicon as test cases. Finally, in the last part of section Results the proposed approach is applied for the study of the uncharged oxygen vacancy, $\square_O$, in TiO_2 anatase. We show that the method described in this work is able to find an additional low-energy configuration which is stable at the level of hybrid functionals.

RESULTS
Evolutionary algorithm
The CMA-ES[33,41–43] has been successfully applied in many complex problems ranging from engineering to artificial intelligence[44]. One of the main improvements of the CMA-ES over previous ESs was the use of evolution paths[33], which contain, not only information on a single generation, but on the whole evolutionary process and can be used to effectively adapt the strategy parameters. Other characteristics of the CMA-ES that make the use of this algorithm more advantageous and require

less problem-dependent tuning are its invariance with respect to any strictly monotonic transformation of the fitness function, and invariance with respect to translations and both proper and improper rotations of the coordinate system[42].

In the following, we start from the implementation presented in reference[43] and adapt it in order to solve the particular optimization problem of interest for this work: finding the stable structures of an isolated point defect in a supercell. If not stated otherwise, the hyperparameter values are set to the default ones proposed in that study. We denote as "cost function" the objective function to be minimized. This is simply the negative of the "fitness function" discussed so far, which one tries to maximize. Given the one-to-one relationship between these two functions, we use both terms interchangeably. The aim of the CMA-ES is to find the vector that minimizes a cost function f:

$$\mathbf{x}^{\star} = \underset{\mathbf{x} \in \mathbb{R}^d}{\operatorname{argmin}} \; f(\mathbf{x}). \tag{1}$$

Here $\mathbf{x}$ is a vector in the genotype space, $\mathbb{R}^d$, and represents an individual. Considering the particular problem of interest in this work, we can consider an individual as a particular defect-containing supercell, whose cell parameters are fixed to those of the host material, but the atomic positions are free to vary. We chose the natural genetic representation of $\mathbf{x}$ being the vector of the atomic Cartesian coordinates, hence, $d = 3N$, with N being the number of atoms in the supercell. Moreover, the natural choice for the cost function is the potential energy of the system. In order to search the genotype space, the CMA-ES samples new individuals from a multivariate Gaussian distribution, which parameters are adapted during the evolutionary process. At generation g a population made of λ individuals, $\{\mathbf{x}_k^{(g)}\}_{k=1,\ldots,\lambda}$, is sampled as:

$$\mathbf{x}_k^{(g)} \sim \mathsf{N}\left(\mathbf{m}^{(g)}, (\sigma^{(g)})^2 \mathbf{C}^{(g)}\right). \tag{2}$$

where $\mathbf{m}^{(g)} \in \mathbb{R}^d$ is the mean of the normal distribution and $(\sigma^{(g)})^2 \mathbf{C}^{(g)}$ its covariance matrix. Specifically $\mathbf{C}^{(g)} \in \mathbb{R}^{d \times d}$ is a symmetric and positive semidefinite matrix and $\sigma^{(g)} \in \mathbb{R}$ a global step-size. All these quantities form the strategy parameters of the algorithm and their values are updated during the evolutionary process.

Once the fitness function has been calculated for each individual, the algorithm employs the truncation selection operator for selecting the parent individuals which will be used to generate the offspring of the next generation. In truncation selection, μ parents are selected deterministically according to their fitness. For this reason, it is convenient to rank all the individuals according to this property; in particular, with $\{\mathbf{x}_{(k)}^{(g)}\}_{k=1,\ldots,\lambda}$ we denote the individuals of generation g rearranged in order of descending fitness. The population size λ generally is chosen according to:

$$\lambda = 4 + \lfloor 3\ln d \rfloor, \tag{3}$$

as suggested in the original work[33]. This quantity is hence increasing very slowly with the search-space dimension. This is particularly convenient for high-dimensional problems as it entails a reduced computational burden when the evaluation of the fitness function is expensive.

In defect calculations, relative large supercells are needed; however, the typical displacement field produced in a material by the defect-induced distortions subsides quickly as the distance from the defect increases. This local nature of the structural perturbation suggests to turn the original optimization problem into a constrained one with a lower dimensionality. Given that $\mathbf{x}(i)$ indicates the coordinates of atom i in the supercell, one can thus introduce a hard cutoff radius c_n such that only the atoms with $\|\mathbf{x}(i) - \mathbf{x}_d\|_2 < c_n$ contribute to the total number of degrees of freedom. Here $\mathbf{x}_d$ indicates the coordinates of the defect center

and the distances are calculated taking the periodic boundary conditions into account. This has the advantage of reducing the dimensionality of the problem from $3N$ to $3s$, where s is the number of atoms within the cutoff. The choice of the value c_n is critical, as in the constrained optimization the algorithm will converge to the global minimum $\mathbf{x}^{\ddagger} \in \mathrm{IR}^d$ of the constrained PES. A too small value would impose too strong constraints, altering considerably the PES and ultimately making $\mathbf{x}^{\ddagger}$ substantially different from the true global minimum $\mathbf{x}^{\star}$, so that, even if the atomic positions of the solution configuration are relaxed without any constraint, the relaxed configuration does not coincide with $\mathbf{x}^{\star}$. An example of this phenomenon is illustrated in the Supplementary Fig. 1, which shows that the relative stability of two low-energy minima of the the $\square_O$ in TiO_2 anatase can be noticeably affected by the value of c_n. In our experience, the value of c_n should at least comprise three to four coordination shells around the defect position. However, the optimal value of c_n is ultimately system-dependent, nevertheless, systematic improvements can be achieved by starting with a small value of c_n and increasing it according to the obtained results.

The localized nature of many point defects suggests another improvement on the algorithm. Normally, the matrix $\mathbf{C}$ is initialized to the identity matrix; however for the problem of interest, this initialization will cause the algorithm to initially sample highly disordered configurations with a very high energy. This is especially so because the initial global step-size, $\sigma^{(0)}$, should be large enough in order to ensure that the PES is sampled thoroughly. To avoid this, we initialize $\mathbf{C}$ in such a way that displacements for the atoms close to the defective site are endowed with a larger variance. This permits an exhaustive sampling of the most relevant part of the PES and avoids the generation of high-energy structures. Specifically, the covariance matrix is initialized as:

$$\mathbf{C}^{(0)} = \mathbf{I}_{3s} + \sum_{j=1}^{3s} c_r^2 (\pi_3(j)) \mathbf{e}_j \mathbf{e}_j^T. \tag{4}$$

The vectors of the standard basis of IR^{3s} are represented by $\mathbf{e}_j$. The function $\pi_3: \mathrm{IN} \to \mathrm{IN}$ maps indexes running over components of $\mathbf{x}$ onto indexes running over the s atoms within the cutoff according to:

$$\pi_3(n) = (n \text{ div } 3) + 1, \tag{5}$$

where div indicates the integer division. The coefficients $c_r(i)$ are given by:

$$c_r(i) = \frac{c_r}{\left(1 + ||\mathbf{x}^{(0)}(i) - \mathbf{x}_d||_2\right)^2}, \qquad i = 1, \dots, s; \tag{6}$$

The addition of 1 Å to the distance from the defect avoids the appearance of singularities when the defect is not a vacancy. In any case, $\mathbf{x}_d$ does not need to coincide with an actual atom in the supercell. For example, for point defect complexes, $\mathbf{x}_d$ could be chosen as the centroid of the defective sites forming the complex. The summation term on the right-hand-side of equation (4) is a matrix of rank $3s$ and c_r is the hyperparameter defining the relative importance of this matrix in the initialization of $\mathbf{C}$.

The last modification is due to the fact that forces in first-principles calculations can be conveniently obtained from the Hellmann–Feynman theorem. So, while the CMA-ES was initially proposed for black-box optimization problems, where one can only evaluate $f(\mathbf{x})$, the information given by $\nabla f(\mathbf{x})$ is generally available in DFT calculations, and it can be used in the algorithm. In particular, we update the mean of the distribution according to:

$$\mathbf{m}^{(g+1)} = \mathbf{m}^{(g)} + c_m \sum_{k=1}^{\mu} w_k \left[\left(\mathbf{x}_{(k)}^{(g)} - \mathbf{m}^{(g)} \right) - c_a^{(g+1)} \nabla f(\mathbf{x}_{(k)}^{(g)}) \right], \tag{7}$$

where c_m is a learning rate which has been set to one in this work. The weights $w_k > 0$ are characterized by a magnitude that is not

increasing with k and such that $\sum_{k=1}^{\mu} w_k = 1$. In addition to the first term, which is also present in the original implementation of the CMA-ES, equation (7) also includes the term depending on the cost function gradient. This update not only forces the distribution mean to move more towards the basins where high-fitness individuals are located, but takes into account the basins' slope. The hyperparameter c_a determines the relative importance of the gradient term in updating this mean. Ideally, its value should be small far from the minimum; otherwise the atomic forces might be very large, which in turn would shift the distribution mean towards regions of the PES far from any minimum. As a consequence, individuals generated in the following generation would be characterized by very high energies and the convergence rate of the algorithm would be reduced considerably or fail completely. Conversely, in order to speed-up the convergence, one would want a large value of c_a close to the minimum and/or in flat regions of the PES where the cost function is almost constant. This suggests to adapt the value of c_a iteratively. In order to satisfies the former two conditions, we employed an update in the form:

$$c_a^{(g+1)} = c_a^{(g)} \exp\left(c_\beta \left[1 - \frac{\sigma^{(g)} || \sum_{k=1}^{\mu} w_k \nabla f(\mathbf{x}_{(k)}^{(g)}) ||_2}{\mathrm{E}\left[||Z||_2\right]} \right] \right), \tag{8}$$

where $Z \sim \mathrm{N}(\mathbf{0}, \mathbf{I}_{3s})$ and E indicates the expectation of a random variable. The hyperparameter c_β has been set to $0.1\sigma^{(0)}$ eV^{-1} throughout the present work.

While the algorithm is characterized by the presence of several hyperparameters, in practice the default values suggested in reference[43] have been found to be adequate. This makes $\sigma^{(0)}$, $\mathbf{m}^{(0)}$, and c_n, c_r, and $c_a^{(0)}$ the only problem-dependent quantities. This is a small set of parameters, all of which have a straightforward physical interpretation, thus making their tuning an uncomplicated task. Additionally, the fact that these hyperparameters only assume numerical values, greatly helps in understanding the effect of the hyperparameter value on the performance of the algorithm. A systematic improvement of the search is then also possible by increasing or decreasing their values in a structured way. For example, some defects might be characterized by long-ranged defect–defect interactions. In this case one can run a couple of calculations with increasing values of c_n and c_r and check whether larger values of these hyperparameters lead to a new global minimum. In some cases, even if defect–defect interaction are long-ranged, as in the cases of charged defects, still rather small values of c_n and c_r can be used. This is illustrated in the Supplementary Fig. 2, where a charged defect is considered. Moreover, it is natural to initialize $\mathbf{m}^{(0)}$ with the coordinates of an intuitive initial defect configuration and to let $\mathbf{x}_k^{(0)} = \mathbf{m}^{(0)}$ for all the individuals of the initial generation. For example, for a system with a vacancy, such initial configuration would consists in the supercell with the removed atom, before any optimization of the atomic positions takes place. Similarly, for an interstitial impurity, the initial configuration would be the unrelaxed supercell with the defect placed in some reasonable interstitial site. We call such an initial defect configuration the population founder.

Once the EA has converged to a single structure, this is optimized locally, after removing all constraints on the atomic positions, by means of some gradient-based optimization method.

Performance: the Si$_i$ in Si with an analytic potential

To investigate the performance of the modified algorithm for the optimization problem of finding low-energy configurations of point defects in solids, we first considered the Si interstitial in bulk silicon using a supercell containing a total of 217 atoms. Given the large number of runs needed to perform a statistical analysis, we used the Tersoff-like potential proposed in reference[45] for silicon, since it is able to produce different local minima for the Si
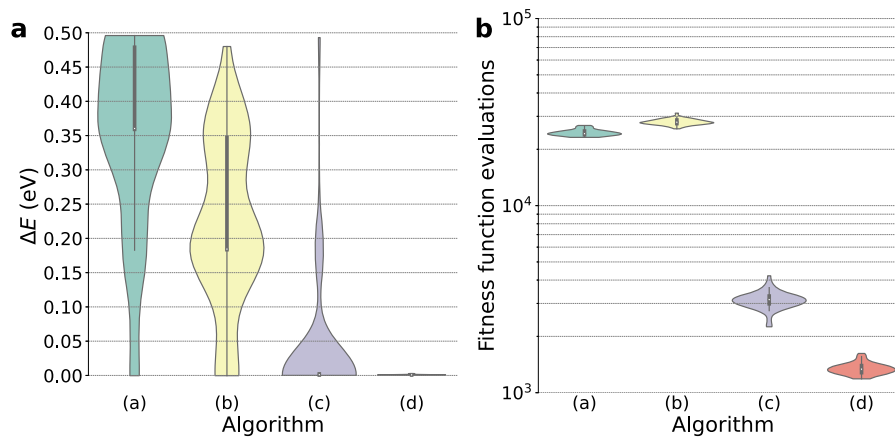
Fig. 1 Comparison of the performance of four different implementations of the CMA-ES. For each of the four implementations 50 independent runs of the algorithm were performed. **a** The distribution of total energies of the 50 final solutions with respect to the global minimum. **b** The distribution of the number of fitness evaluations necessary before such converged solutions are reached. Labels: (a) Original CMA-ES. (b) Initialization of the covariance matrix as per equation (4), without using a hard cutoff. (c) Initialization of the covariance matrix as per equation (4), using a hard cutoff. (d) The final version of the algorithm, as proposed in this work including gradients in the update of the distribution mean. In each violin plot, the white dot represents the median of the distribution, the bold black bar the interquartile range and the thin black lines represents the range of data points comprising the interquartile range extended by a factor of 1.5.

interstitial defect. For all cases, as the population founder, we took a configuration that gradient-based optimization methods relax to a minimum of the PES with an energy around 0.5 eV larger than the global minimum. The other parameters were set as follow: $\sigma^{(0)} = 0.08$ Å, $c_n = 4$ Å, $c_r = 4$ Å, $c_a = 0.6$ Å^2 eV^{-1}. Different variations of the CMA-ES algorithm are then applied in order to find the global minimum. For each algorithm, 50 independent runs are performed. A run is considered to be converged when the standard deviation of the energy of the individuals within a generation remains below 0.01 eV for at least ten consecutive generations. The converged solution is then relaxed, without any constraint on the atomic positions, applying a gradient-based optimization method.

The performance of the algorithms is compared in Fig. 1, where the distribution of the energies of the converged solution (with respect to the energy of the global minimum) and the distribution of the number of fitness-function evaluations needed to reach the solution are illustrated. Initializing the covariance matrix according to equation (4) without introducing a hard cutoff (label (b)), already helps the algorithm to converge to the global minimum more often. A drastic improvement is obtained however when the hard cutoff is also enforced (label (c)), illustrating the detrimental role of a high search-space dimensionality on the overall performances. As one can see, the algorithm proposed in this work (label (d)) is both the most robust and requires the least number of fitness-function evaluations, leading to much lower computational costs. As mentioned above, updating the mean using the gradient of the fitness-function might prevent the convergence of the EA, especially when large $\sigma^{(0)}$ and c_r are used. In this example, only four runs do not converge due to this issue. On the other hand, in all other cases, the EA always converge to some solution. As shown by this discussion, the modifications applied to the algorithm lead to fewer evaluations of the fitness function and a more robust implementation than the original one (label (a)).

Case study: the Si$_i$ in Si with DFT

The short-ranged nature of interactions in the empirical potential used in the previous section produces a too complex PES, with several false minima, as one can notice from Fig. 1a. We will now illustrate the algorithm performances on the Si$_i$ defect employing DFT at a level of accuracy that allows to reproduce the defect configurations observed in o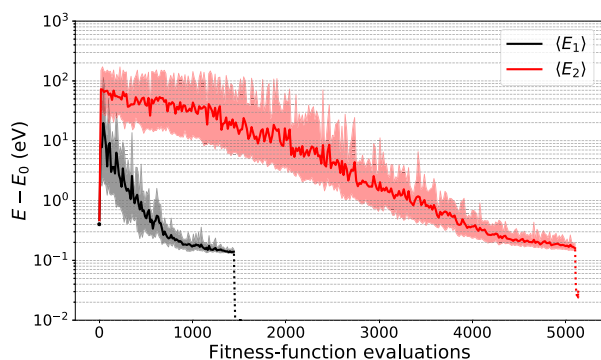ther DFT studies employing local or semi-local functionals[46–49]. In particular, among the various defect configurations that have been observed in these studies, we consider those three with the lowest energy: the Si in an hexagonal interstitial (H), in a tetragonal interstitial (T), and the split $\langle 110 \rangle$ configuration, where two silicon atoms form a dumbbell configurations along the $\langle 110 \rangle$ direction (X). Well-converged simulations tend to agree that the H and X configurations have very similar energies, with the X one most likely being the ground state. On the other hand, the energy of the T configuration is noticeably larger[46,48,49]. Even though simple exchange correlation functionals tends to underestimate the formation energy of the self-interstitial defects, much more accurate approaches tend to agree on the fact that the X and H configurations have very similar energies, noticeably lower than that of the T configuration[50–52]. In this work, we also found that the H and X configurations have very similar energies, but the ground state is represented by H, with the energy of the X configuration being just 23 meV higher. The energy of the T configuration is around 0.37 eV above the ground state. These results are adequate for the purpose of presenting the method proposed in this work, as all the three low-energy minima are present and the multimodal PES offers a realistic test case for our approach.

We choose a T insterstitial as the population founder. While a gradient-based optimization would lead to the T configuration, our EA correctly converges to the H one. In particular, we repeated the simulation three times, each time setting $c_n = 4$ Å. In two calculations $c_a^{(0)}$ was set to 0 Å^2 eV^{-1} and in one to 0.3 Å^2 eV^{-1}. For the former set of calculations, $\sigma^{(0)}$ was set to 0.2 Å, for the latter to 0.08 Å. Finally, c_r was set either 0, or 10 Å. In all cases, the algorithm converged to the ground state. Even tough the choice of the population founder generally falls on a very small set of intuitive structures, in order to check the converge of the algorithm with respect to the choice of such founder, we run six simulations with the same computational parameters except for the population founder. This was chosen by randomly displacing the interstitial Si atom from either the H (three times) or T (the other three times) site. The displacements were sampled each time from a trivariate normal distribution with mean zero and covariance matrix $\sigma_0^2 \mathbf{I}_3$, with σ_0 equal to half of the Si-Si distance in the pristine supercell. Supplementary Fig. 3 shows that in all cases the algorithm converged to the same configuration with a comparable number of fitness-function evaluations. Even in the

case when the population founder is characterized by a complex structure with an energy very close to the global minimum, such as in the case of the X configuration for the Si interstitial, the algorithm is able to find the global minimum using a similar set of hyperparameters values, as illustrated in Supplementary Fig. 4.

As mentioned in section Evolutionary Algorithm, a balance between $\sigma^{(0)}$, $c_a^{(0)}$, and c_r must be kept. As too large initial variances will lead to highly disordered structures characterized by large energies. Hence, if $c_a^{(0)} \neq 0$, the population mean will be shifted to areas of the PES far from any minimum, making the convergence very slow and diminishing the probability to find the global minimum. In general, one should reduce $c_a^{(0)}$ the larger $\sigma^{(0)}$ and c_r are. However, we can use this phenomenon in order to force the algorithm to converge to some local minimum. To give an example, the algorithm converges to the X configuration by setting $\sigma^{(0)}$ to 0.15 Å, $c_a^{(0)}$ to 0.05 Å^2 eV^{-1} and c_r to 25 Å. Additionally, another local minimum with energy of around 0.4 eV above the ground state was found by increasing $\sigma^{(0)}$ to 0.18 Å. This discussion shows how an extensive exploration of the PES can be obtained by increasing systematically the algorithm parameters controlling the variance for atomic displacements, forcing the algorithm to converge not only to the global minimum but to local ones as well. As the evolutionary process is stochastic in nature, results will generally vary from run to run. However, the above analysis describes a trend we observed all the times our method was employed.

Figure 2 shows the average population energy as a function of the number of evaluations of the fitness function. The zero of the scale is set to the energy of the fully relaxed H configuration, which represents the ground state. The bold black line represents a run that converges to this ground state. Note that since a hard cutoff is applied, the energies of the converged structures will be larger than the ground state one. In any case, if the cutoff is large enough, a gradient-based optimization can be used to quickly relax these structures to the proper global minimum, as shown by the dotted lines. One can also note that in the last generations, the average population energy changes very little. In the present work, we consider the algorithm to be converged when the population standard deviation remains below 0.01 eV for at least ten consecutive generations. This in an extremely conservative criterion: for this system, setting a threshold of 0.1 eV for the population standard deviation still yields a converged solution that relaxes to the ground state, but requires about half of the
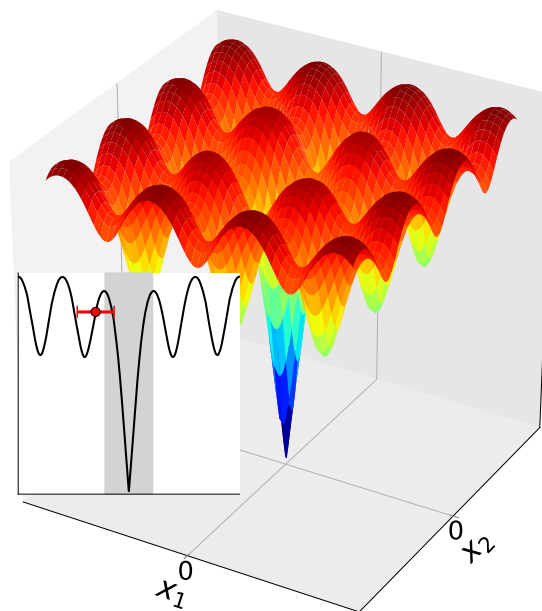
number of fitness-function evaluations, as shown in Supplementary Table 1. Another run of the EA is shown in Fig. 2 by the red line. In this run we used the initial parameters that lead the EA to converge to the X configuration. Both runs start from the same population founder, but, by comparing the two lines, is clear that larger values of $c_a^{(0)}$ and $\sigma^{(0)}$ yield structure with a much higher energy, making the algorithm convergence slower and reducing the probability of finding the global optimum.

Unsupervised Machine Learning Model

While the CMA-ES algorithm has proven to be very reliable on a large set of non-separable multimodal benchmark functions, in some complex cases, particularly large population sizes are required in order to reach the global optimum[41]. Moreover, even when the algorithm converges to the global optimum, it might also sample basins containing other low-energy minima, a phenomenon which is of interest for finding metastable defect configurations. As discussed in the previous section, it is possible to find metastable structures by varying the initial parameters, but it is necessary to rerun the algorithm in order to make it converge to these competing minima. While this offers a simple and systematic method for finding metastable structures, one would wish to limit the number of DFT calculations. A method to exploit all the data generated during an evolutionary run and discover low-energy minima without the need to rerun the algorithm would thus be particularly useful.

To illustrate the ability of the EA to explore basins of the PES different from the one containing the converged solution, we take a multimodal function whose global minimum is difficult to reach, Fig. 3, as an example. Even though the cost function at the global minimum, $\mathbf{x}^\star$, has a considerable lower value than the local minima, convergence to $\mathbf{x}^\star$ is not straightforward. As a concrete example, consider the situation depicted in the inset. The red dot

Fig. 2 Average population energy as a function of the number of fitness-function evaluations. Calculations were done at the LDA level for the Si$_i$ defect. The zero of the energy is set to the ground state configuration for this defect (H). The maximum and minimum energies for each generation are also shown. Black: initial parameters: $c_n = 4$ Å, $c_r = 10$ Å, $\sigma^{(0)} = 0.08$ Å and $c_a^{(0)} = 0.3$ Å^2 eV^{-1}. Red: initial parameters: $c_n = 5$ Å, $c_r = 25$ Å, $\sigma^{(0)} = 0.15$ Å and $c_a^{(0)} = 0.05$ Å^2 eV^{-1}. The dotted lines represent the calculations needed to relax the converged solution to the equilibrium configuration using gradient-based optimization.

Fig. 3 Exploring the minima of an objective function. A reparametrization of the two-dimensional Ackley function, $f(x_1, x_2) = -a \exp\left(-b\sqrt{\frac{1}{2}(x_1^2 + x_2^2)}\right) - \exp\left[\frac{1}{2}(\cos(cx_1) + \cos(cx_2))\right] + e + a$, with $a = 3$, $b = 1$, $c = \pi/3$, chosen in such a way that the basins of the local minima and of the global minimum have a comparable width. In the inset, the function is evaluated on the $x_2 = 0$ plane. The red dot represents the location of $\mathbf{m}^{(0)}$ and the red error bar represents the magnitude of $\sigma^{(0)}$. The gray area represents a region of the search space where a gradient-based optimization approach would lead to the global minimum.

represents the location of the CMA-ES initial mean $\mathbf{m}^{(0)}$ and the red error bar represents the magnitude of $\sigma^{(0)}$. We employ the basic implementation of the CMA-ES and use a population size of 20 individuals. In such conditions, even though the initial global step-size is large enough to allow exploration of the basin of $\mathbf{x}^{\star}$, the algorithm might not converge to this optimum, as the region of the $\mathbf{x}^{\star}$ basin where cost function is lower than other minima is relatively narrow. As a consequence, the update of $\mathbf{m}$ might lead the individuals distribution to shift toward other basins. As a matter of fact, if we define p as the probability that the CMA-ES will converge to $\mathbf{x}^{\star}$ in this problem, and let $\hat{p}$ be the maximum likelihood estimator of p, then, after running the algorithm 3000 times, the calculated value of $\hat{p}$ is found to be ≈ 0.79. This is a rather low probability, considering the fact that the problem is only two-dimensional. For comparison, if one employs the standard parametrization of the Ackley function, keeping the other parameters unaltered, the calculated $\hat{p}$ turns out to be 1. Nevertheless, even when the algorithm does not converge to $\mathbf{x}^{\star}$, it still visits regions of the search space within its basins. For this problem, we found indeed that the estimated expected number of times this happens is around 10. Even though only a very small fraction of individuals visits the basin of the global minimum, in the context of the present work we might expect that structures located in the global minimum basin to be noticeably different from those located in other basins, and thus recognizable by some unsupervised learning method even if there are very few of them.

In particular, analyzing the way the CMA-ES operates in the evolutionary process, one would expect that in the first generations widely different structures are generated. As the evolution proceeds, the algorithm will tend to converge to some basin, and the overall covariance of the individuals distribution will decrease, generating more and more similar structures. Assuming that the structures belonging to a certain basin are characterized by certain common features, employing appropriate structural descriptors these structures will appear as localized clusters. This suggests using a density-based clustering algorithm in order to identify these structures, as described in the example below.

Case study: Pristine silicon with DFT

For presenting the unsupervised model used to find low-energy defect configurations, pristine silicon is an interesting case study. For this system, the global minimum is obviously the defect-free diamond lattice structure; however there are several metastable defect laden structures, including Frenkel pairs, where a Si moves to an interstitial position and leaves behind a vacancy, and the fourfold coordinated defect (FFCD)[47,53], which is possibly the intrinsic defect in silicon with the lowest formation energy. In order to run the proposed EA, we centered the "defect position" to one of the Si atoms in the supercell and used the following computational parameters: $c_n = 4\,\text{Å}$, $c_r = 20\,\text{Å}$, $\sigma^{(0)} = 0.1\,\text{Å}$ and $c_a^{(0)} = 0.05\,\text{Å}^2\,\text{eV}^{-1}$. Running the EA, we found it to converge, as expected, to the diamond structure.

In order to exploit all the data generated during this evolutionary run, we represented all the generated structures using the crystal structural fingerprints proposed by Oganov and Valle[54]. These structural fingerprints give a description of the supercell structure analogous to that of a pair correlation function: a smoothed histogram of the distribution of distances of pairs of atoms is constructed. On the calculated structural fingerprints, we then applied principal component analysis (PCA) and kept only the first two components, which preserve more than 95% of the dataset variance, as shown in the bottom-right inset of Fig. 4. The main panel of Fig. 4 shows the first two principal components of the collected structure. The shape of the scatter plot is very characteristic and we can compare it to that of a comet: a dense region of structures, those obtained in the last generations, form
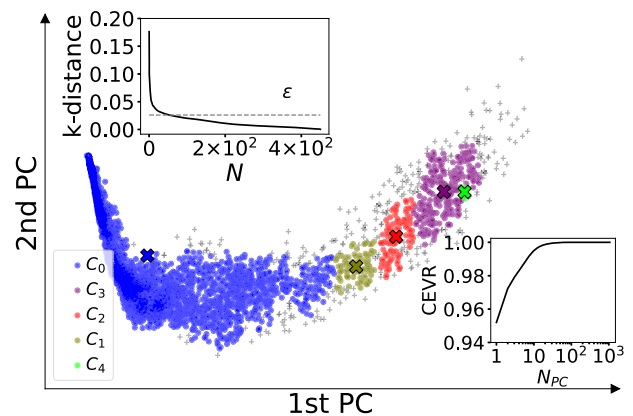

Fig. 4 Representation of the structures visited by the evolutionary algorithm. Results showed for pristine bulk silicon crystal. Scatter plot of the first two principal components of crystal fingerprint descriptors for the structures generated during the EA run (see main text for the details). Cluster centroids are shown with crosses. The inset at the top-left corner shows the k-distance plot used to find the clustering parameter ε. The inset at the bottom-right corner shows the cumulative explained variance ratio (CEVR) as a function of the number of principal components (N_{PC}). Cluster names are given according to their distance from the left-most cluster and they are ordered in the legend according to their size.

the head, while the tail is more rarefied and is composed by the structures generated in the first generations. The first principal component strongly anticorrelates with the generation number so that the further a structure appears to the left of the plot, the more it belongs to later generations. As expected, between the head and the tail one can notice some regions with a noticeably higher density of structures. In order to collect them into distinct clusters, we employed the density-based DBSCAN algorithm[55]. This algorithm uses core points in order to identify clusters. A core point q_C is a point such that a sphere of radius ε centered on it contains at least minPts points. All points with a distance of less than ε from q_C are said to be directly density-reachable from it. A cluster is made of all points for which a chain of sequentially directly density-reachable points connecting to a core point exists. All other points are considered to be noise and represented by crosses in Fig. 4. To obtain this figure, we set minPts to 8 and the value of ε was chosen by analyzing the 8-distance plot shown in the top-left inset of the figure. In particular, the exact value of ε has been chosen in order to obtain distinct and dense clusters. Given the large variance preserved by the projection onto the subspace spanned by the first two principal axes, we can perform the clustering on this two-dimensional space, rather than on the original space. This is particularly advantageous since clustering methods based on neighborhoods, as quantified by the Euclidean metric, would perform very poorly in high-dimensional spaces.

The obtained clusters were labeled according to their Euclidean distance from the converged solution (the pristine Si supercell), with C_0 being the cluster containing the converged structure and C_4 being the cluster whose centroid is the most distant from that of C_0. The Euclidean distance perhaps offers the most natural way of selecting the most interesting clusters, as it produces a measure of how distinct structures differ according to their geometric representation. Additionally, one can also consider distances in distributions, specifically how much the distributions that have generated the individuals in the clusters of interest differ, on average, from the distribution of the converged solution. In order to assess the similarity between two probabilities distributions, we use two quantities: the Kullback–Leibler (KL) divergence and the Bhattacharyya distance. The KL divergence
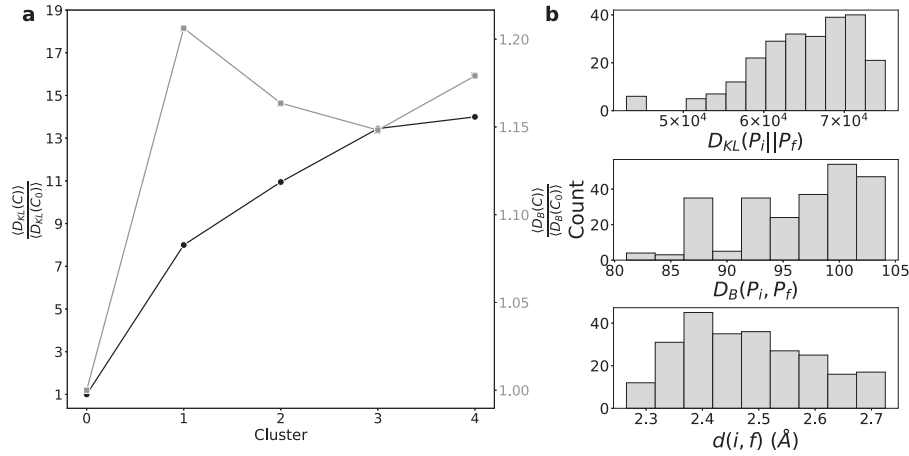
Fig. 5 Similarity measures for clusters of structures. a Cluster relative average KL divergence and relative average Bhattacharyya distance. **b** Distribution of the KL divergence, Bhattacharyya distance, and Euclidean distance for the individuals in cluster C_3.
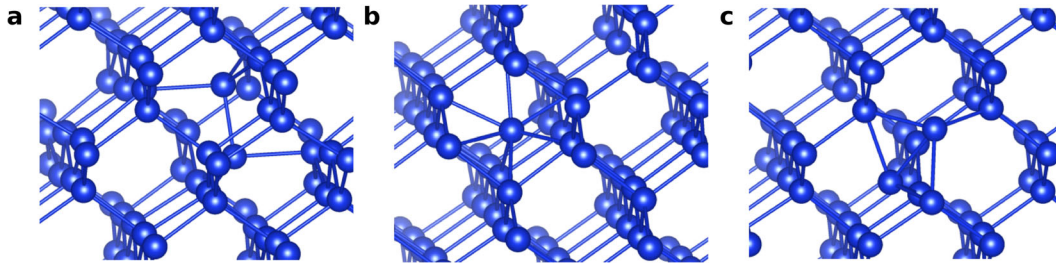
Fig. 6 Defective structures in stoichiometric Si supercells found by the clustering method. a The FFCD. **b** The Frenkel pair where Si_i is in an hexagonal interstice. **c** The Frenkel pair where the interstitial Si is in the split $\langle 110 \rangle$ configuration. Images created with VESTA[73].

quantifies the similarity between probability distributions P and Q according to:

$$D_{KL}(P||Q) = E_{X \sim P}\left[\ln \frac{P(X)}{Q(X)}\right], \tag{9}$$

where X indicates a random variable and the expectation is taken with respect to X which is distributed according to P. The Bhattacharyya distance function takes the form:

$$D_B(P, Q) = -\ln\left(\int \sqrt{P(x)Q(x)}dx\right) \tag{10}$$

Figure 5a shows the clusters relative average KL divergences and Bhattacharyya distances. In particular, the cluster average divergence or distance is calculated according to:

$$\langle D(C) \rangle = \frac{1}{n_c}\sum_i^{n_c} D(P_i, P_f) \tag{11}$$

where P_i represents the distribution of the population where structure i belongs, and P_f represents the distribution of the last population, and n_C the number of structures in that particular cluster.

One can see that in this case the relative KL divergence and Bhattacharyya distance noticeably differs from the value obtained for C_0. While the former increases monotonically with the Euclidean distance from C_0, the latter changes by small amounts. In any case, cluster C_4 seems to be an ideal candidate for discovering new structures as it is the most different in geometric representation, as quantified by the centroid Euclidean distance, and also in distribution, according to the KL divergence, from the converged solution. C_4 contains only three structures, optimizing them to the closest minimum using a gradient-based method results in a structure relaxing back to the pristine system, one

relaxing to the FFCD and one to another defective configuration with a much larger formation energy: around 6.7 eV.

Since C_4 contains very few structures, one would want to consider C_3 as the next cluster of interest. C_3 contains almost 250 structures and it is not convenient to relax all of them. One has to employ a finer selection scheme for shortlisting them. Following the idea of choosing structures according to their Euclidean distance and distribution distance with respect to the converged solution, in Fig. 5b we show the distributions of these quantities for the structures in C_3. As the clustering algorithm is based on the density in the linear space spanned by the first two principal axes, it is not surprising that the variance of the Euclidean distance from the converged solution is relatively small. However, structures in the same cluster can be sampled from rather different distribution; therefore the cluster might show a much larger variance for the KL divergence and Bhattacharyya distance. Hence, we selected only those structures in C_3 whose Euclidean distance from the converged solution falls in the upper quartile range and whose KL divergence and Bhattacharyya distance are above the respective median values. This procedure returns 30 structures, which are then optimized with a gradient-based method. Of these, 22 relax back to the pristine supercell, one relaxes to the FFCD, two structures relax to a defective configuration with a formation energy of around 4.4 eV, which can be described as a Frenkel pair where the interstitial atom is in a H site, one structure relaxes to another Frenkel pair, this time the interstitial is in a X site and has formation energy of around 4.6 eV, and the remaining four structures relax to more disordered configurations with formation energies between 6.4 and 9.5 eV. The three discovered defective structures with the lowest formation energy are shown in Fig. 6. While different filtering criteria than the one suggested here can be used, this analysis shows the efficacy of the presented approach for selecting a small

number of structures which lead to diverse low-energy minima with a modest computational cost.

Supervised machine learning model

The algorithm outlined in section Evolutionary Algorithm is in itself able to find low-energy defect configurations, for all systems considered in this work, at a reasonable computational cost. Nevertheless, the convergence of the self-consistent field cycle might pose a more daunting task in terms of time and computer power for particular systems and/or when very large supercells must be employed. It is therefore desirable to use a ML metamodel able to accurately predict the system electronic energy and thus reduce the number of required first-principles calculations and hence the computational burden. As we mentioned in the introductory section, efficient materials descriptors and models have been shown to be able to approximate DFT energies within few meV per atom[34–38]. Among the various proposed models, we have chosen to employ GP regressors[56], as they have already shown to be very effective in reducing the amount of first-principles calculations required in EA[39,40] and in other optimization problems encountered in materials science, including Bayesian optimization[57–61]. Furthermore, the underlying Bayesian formalism enables the use of an on-the-fly active learning training during the evolutionary process, as will be described later in this section.

A GP is a collection of random variables such that any finite subset thereof is a jointly-Gaussian random vector. In our case, we are interested in modeling the cost function f as a Gaussian random variable. The value of f depends on the actual crystal structure, which is described by appropriate descriptors and forms the input $\mathbf{x}$ of the model. The joint distribution of f in uniquely defined by a mean function $m(\mathbf{x})$ and a kernel function $k(\mathbf{x}, \mathbf{x}')$. In particular, we can conceptually divide the input vectors into a set which is used for training, for which we know the value of the cost function, and a set representing previously unseen structures. To distinguish between quantities already observed and quantities that concern the prediction of the cost function, we add a superscript * to the latter. According to this notation, let $\mathbf{X}^*$ be a $n_* \times d$ matrix where each row represents an input vector $\mathbf{x}_i^*$, $i = 1, \dots, n_*$ at which one would want to predict the value of f, then the joint prior distribution of the random vector $\mathbf{f}^*$, where $f_i^* = f(\mathbf{x}_i^*)$, $i = 1, \dots, n_*$, is, according to the model:

$$\mathbf{f}^* | \mathbf{X}^*, m, k \sim \mathrm{N}(\mathbf{m}^*, \mathrm{K}(\mathbf{X}^*, \mathbf{X}^*)), \tag{12}$$

where the components of $\mathbf{m}^*$ are given by $m_i^* = m(\mathbf{x}_i^*)$ and the elements of the covariance matrix K are given by: $K_{ij} = k(\mathbf{x}_i^*, \mathbf{x}_j^*)$. Note that the mean function m and the kernel function k are always the same whether we are considering observed points or not. Conditioning the prior distribution of $\mathbf{f}^*$, equation (12), on the values of f observed on a set of input vectors $\mathbf{X} \in \mathrm{IR}^{n \times d}$, yields the joint posterior distribution:

$$\mathbf{f}^* | \mathbf{X}^*, \mathbf{X}, \mathbf{f}, m, k \sim \mathrm{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*), \tag{13}$$

where the posterior mean $\boldsymbol{\mu}^*$ and covariance matrix $\boldsymbol{\Sigma}^*$ are given, respectively, by:

$$\boldsymbol{\mu}^* = \mathbf{m}^* + K(\mathbf{X}^*, \mathbf{X}) K^{-1}(\mathbf{X}, \mathbf{X})(\mathbf{f} - \mathbf{m}) \tag{14}$$

$$\boldsymbol{\Sigma}^* = K(\mathbf{X}^*, \mathbf{X}^*) - K(\mathbf{X}^*, \mathbf{X}) K^{-1}(\mathbf{X}, \mathbf{X}) K(\mathbf{X}, \mathbf{X}^*). \tag{15}$$

Once the posterior of $\mathbf{f}^*$ is obtained, one can make predictions at inputs $\mathbf{X}^*$ by using the posterior mean $\boldsymbol{\mu}^*$: the estimator which minimizes the mean squared error of the prediction. The most critical aspect affecting the posterior distribution is the choice of the kernel function k. Several kernels have been proposed in the literature. In the present work, we employ a linear combination of

two squared exponential kernels:

$$k(\mathbf{x}, \mathbf{x}') = C\left(x \exp -\frac{||\mathbf{x} - \mathbf{x}'||_2^2}{2l_1^2} + (1-x) \exp -\frac{||\mathbf{x} - \mathbf{x}'||_2^2}{2l_2^2} \right). \tag{16}$$

The length scale hyperparameters l_1 and l_2 have been constrained on the range $[10^2, 10^6]$ and $[10^{-2}, 10^2]$, respectively, in order to add flexibility to the prediction distribution. A similar kernel has been used in reference[40] in a GA context and has shown good performances. The values of the hyperparameters C, l_1, and l_2 are obtained during the training phase by maximizing the log marginal likelihood, as described in reference[56]; while the value of x is chosen by the user in [0, 1]. Regarding the mean function, we took it as the sample mean of the dataset.

An advantage of having the posterior covariance matrix, is that for each prediction we have an estimation of its uncertainty. We use this quantity in order to decide which structures generated by the EA should have their fitness evaluated directly by a DFT calculation, or approximated by the GP metamodel. In particular, at each generation g, the GP metamodel is trained using the individuals in the previous generations whose energy has been calculated directly by DFT. The model then calculates the prediction standard error for each individual in generation g. If this quantity is larger than a user-define threshold, then the DFT energy is calculated for that individual, otherwise the GP-predicted energy is used. This is an effective way to train the metamodel on-the-fly allowing for a certain degree of active learning: the metamodel can automatically choose the proper training data needed to maintain a predefined accuracy for the approximated PES. It is important to note that it is not necessary for the latter to be an exact reproduction of the true PES: the approximation should be good enough to represents the main basins of the PES, but these need not to be an exact replication of the real ones. The main purpose of the metamodel is to offer a surrogate PES which can be used for sampling the real one at a much lower computational cost.

With the metamodel in use, the number of fitness-function evaluations can be considerably reduced, as shown in Fig. 7. In particular, we considered the same Si_i defect and computational parameters as in section "Case Study: the Si_i in Si with DFT".
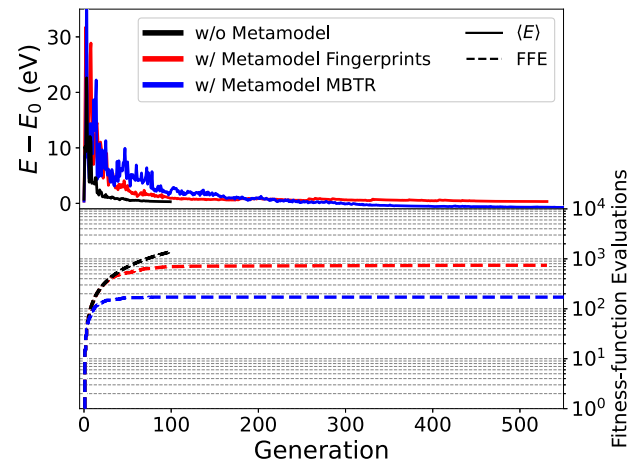
Fig. 7 Reducing computational costs using the ML metamodel. Comparison between two types of runs of the EA for the Si_i defect. One type employs only DFT calculations (black lines) and one is performed with the use of the machine learning metamodel, emyploying either the crystal fingerprints descriptors (red lines) or the MBTR ones (blue lines). Dashed lines represent the number of required fitness-function evaluations. The bold lines, the population average energy referenced to the ground state energy of Si_i obtained in LDA calculations.

Two types of runs were performed: for one type we employed the metamodel during the evolutionary process while the other run we did not. In both cases, the converged solution relaxed to the ground state. The metamodel was trained employing both the crystal fingerprint descriptors proposed by Valle and Oganov[54], which we already introduced above, and the more recently proposed descriptors: the many-body tensor representations (MBTR)[62]. The latter can be considered a generalization of the former, including more general relationships between atoms in the system and taking many-body terms into account. In particular, we employed only two- and three-body terms and we calculated these descriptors using the `DScribe` library[63].

The minimum standard error allowed for the metamodel was set to 10 meV per atom. From Fig. 7, one can notice that when the metamodel is employed, no explicit energy calculations need to be performed after a certain amount of generations. Overall, even tough the metamodel run needs a larger number of generations in order to converge (partly because no gradient term is used in updating the population mean), it requires a considerable less number of fitness-function evaluations. In particular, this number is reduced by more than a factor of two employing the crystal fingerprints descriptors and by more than a factor of eight employing the MBTR ones. As the development of efficient materials descriptors is a very active field of research in materials science, the number of DFT calculations might be reduced even more with the most recent descriptors.

Case study: The neutral $\square_O$ in TiO$_2$ anatase

We finally applied the method proposed in this study to a complex defect in a transition-metal oxide. The uncharged oxygen vacancy, $\square_O$, is perhaps the native defect of anatase TiO$_2$ that has received the most attention both from experimental and computational studies. Several first-principles calculations have been performed using a wide range of functionals and different structures have been reported. Simply removing an oxygen atom from the host material and relaxing the atomic positions typically leads to the "simple" vacancy structure, which is characterized by the localization of the two extra electrons on the vacancy site and has the point group symmetry $mm2$, Fig. 8a, b. However, another configuration with considerable lower energy can be obtained. This is called the "split" vacancy configuration and is distinguished by a broken symmetry (point group $\bar{2}$) and the localization of an extra electron on a neighboring Ti atom[64–67], Fig. 8c, d. Due to the inability of standard DFT to treat strongly localized electrons, the
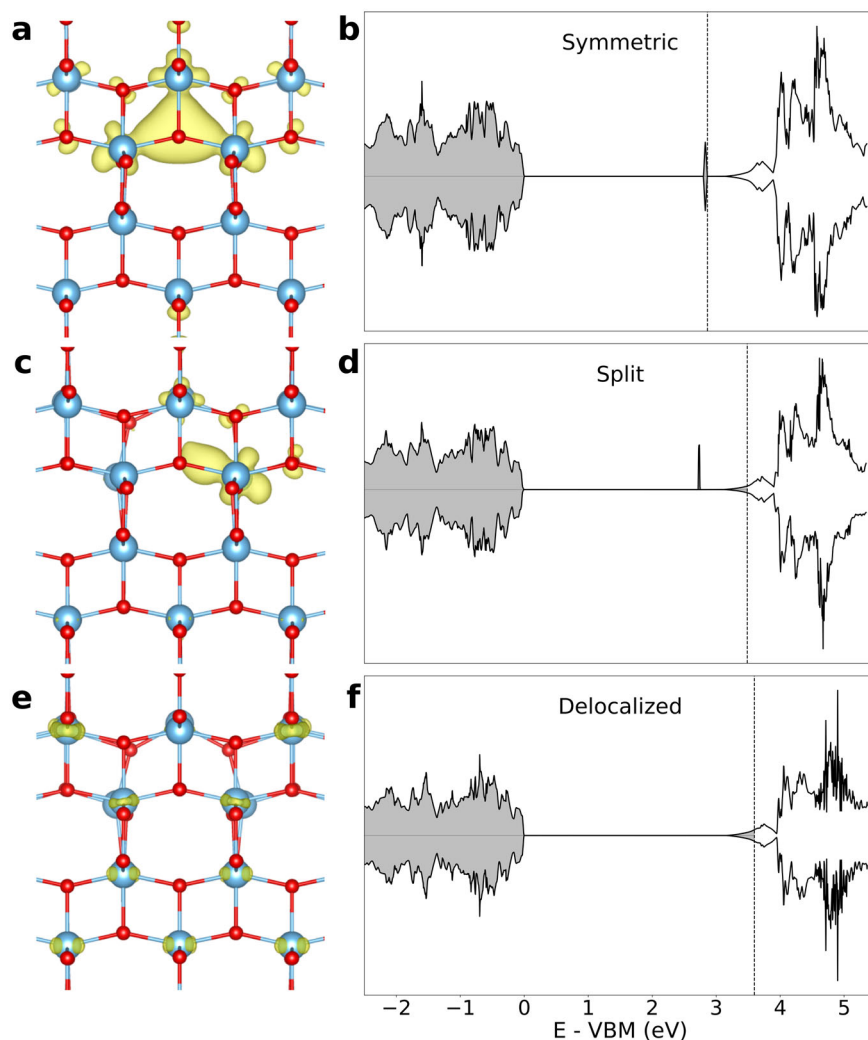

Fig. 8 **Electronic structure of the possible (meta)stable $\square_O$ configurations in TiO$_2$ anatase.** Isosurfaces for: **a** the simple vacancy, **c** the split vacancy, and **e** the delocalized vacancy plotted for the valence electron density projected onto the defect-induced level. Isosurfaces levels are shown at 0.005 e Å^{-3}. Electronic density of states for these configurations, calculated at the HSE15 level using a $2 \times 2 \times 2$ k-point mesh: **b** the simple vacancy, **d** the split vacancy, and **f** the delocalized vacancy. The dotted lines represent the Fermi level. The gray areas emphasize occupied states. The energy is referenced with respect to the valence band maximum (VBM).

split configuration is most prominent using either the DFT+U approach or hybrid functionals. We have recently shown that DFT+U, which is computationally much less demanding, is able to yield geometric structures in good agreement with those obtained with hybrid functionals[67]. The agreement of the defect formation energy is considerably worse: using DFT+U, we found that the simple vacancy configuration has energy 0.77 eV larger than the split vacancy, while using HSE, this difference is of about 0.38 eV[67]. In both cases these minima are however clearly distinguishable. Based on these observations, we will explore the low-energy configurations of this defect at the DFT+U level, using the EA and the density-based clustering method in order to explore the PES. A final set of selected structures is then studied also with a hybrid functional.

The results we present below are obtained by setting the following values for the EA hyperparameters: $c_n = 6.5$ Å, $c_r = 20$ Å, $\sigma^{(0)} = 0.12$ Å and $c_a^{(0)} = 0.3$ Å^2 eV^{-1}. We have run the EA multiple times with different hyperparameters values (c_n from 4 Å to 6.5 Å, c_r from 3 Å to 50 Å, $\sigma^{(0)}$ from 0.10 Å to 0.2 Å and $c_a^{(0)}$ from 0 Å^2 eV^{-1} to 0.3 Å^2 eV^{-1}). In all cases the solution of the EA converged to the same defect configuration.

Employment of the metamodel does also lead to the same solution, if similar values for the hyperparameters are kept and the threshold value for the prediction standard error is set below 10 meV per atom. As mentioned in the previous section, the main advantage of the metamodel consists in reducing the number of DFT calculations. As shows in Fig. 9, even with a complicated PES as the one produced by the DFT+U method, the crystal fingerprint descriptors of Valle and Oganov allow to obtain accurate predictions of the DFT energy after few tens of generations. In particular, the figure illustrates the validation root mean square error (RMSE) and prediction standard errors as a function of the generation. The results at generation g are obtained by training a GP regressor on the previous generations. The energy values at generation g are then used as the validation set where the RMSE is calculated. One can see that a prediction RMSE and standard error within 10 meV per atom are readily achieved after a couple tens of generations. As a matter of fact, it is in the nature of the EA that the more the evolution proceeds, the more similar structures are generated. Hence, the prediction RMSE will ultimately decrease, as can be seen from equation (14), which shows that the prediction of a GP is a linear smoothing of the training set targets.

With all the parameters described in the previous section, we found the EA to converge always to a configuration which is
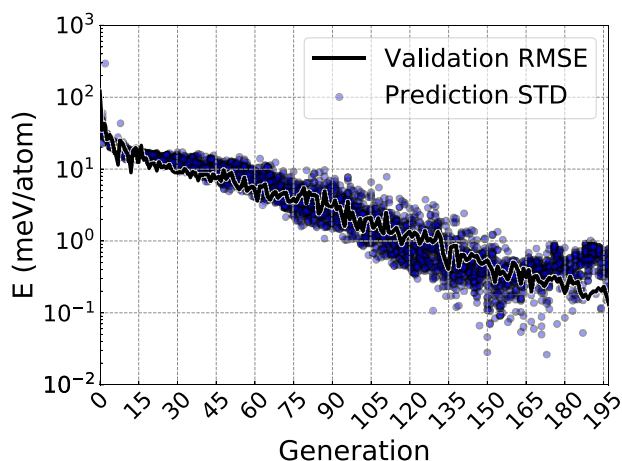

Fig. 9 Performance of the crystal fingerprint descriptors during the evolutionary process. Results shown for $\square_O$ in anatase TiO$_2$. Bold black lines represent the calculated root mean square errors. Blue dots represents the values of the prediction standard error. Each generation corresponds to 20 fitness-function evaluations.

neither the simple nor the split vacancy one. This configuration resembles somewhat the simple vacancy, but has lower symmetry and considerable lower energy: its formation energy is only around 30 meV larger than that of the split vacancy. We named this configuration as the "pseudo-symmetric" vacancy.

In order to look at other explored regions of the PES, we again represented the structures visited by the EA using the crystal fingerprints and plot the two principle components as a scatter plot in Fig. 10a. We obtain an analogous cloud as that of Fig. 4. Also in this case the first principal components strongly anticorrelates with the generation number so that the later generations are found to the left. We employed the clustering approach described in section Unsupervised Machine Learning Model collecting all the structures visited by a run of the EA. We set the DBSCAN parameter minPts to 8 and found the value of ε by analyzing the 8-distance plot shown in the inset of Fig. 10a. Several clusters were found, with C_0 being the largest one since it contains individuals generated in the last generations, where the structures are similar to one another. Nine other clusters were found. As in section "Case Study: Pristine silicon with DFT", their names reflects the magnitude of the Euclidean distance from C_0. The behavior of relative average KL divergence and Bhattacharyya distance per cluster is similar to the one described by Fig. 5a, but this time both quantities increase monotonically with the cluster's centroid distance from C_0. In particular, both quantities also show a steep increase in magnitude already from C_1, illustrating that the algorithm is indeed able to separate the converged structure from others. According to these observations, we took the structures in C_9, C_6, which are relatively small clusters with large values of the relative $\langle D(C) \rangle$ and we also included the structures in C_5, since the cluster is very small. The overall number of structures in the three clusters is the following: ten for C_9, seven for C_6, and six for C_5.

We relaxed all 23 structures from the three clusters using a gradient-based optimization algorithm. Multiple new minima were found, including a configuration with energy lower than that of the split vacancy by around 50 meV, but with a closely related structure. As a matter of fact, most of the new local minima show a structure either similar to that of the pseudo-symmetric vacancy or to the split vacancy, but have distinguishable different levels of lattice distortions. Adding to these structures those of the known symmetric and split vacancy configurations, we obtain a total of 25 structures, whose distribution in formation energy, with respect to that of the split vacancy, is shown in Fig. 10b. The presence of several low-energy minima indicates that the PES described by the DFT+U method is artificially complex. As the next step, we grouped all these 25 structures into nine families, according to the values of their first two principal components; specifically, the family with index 5, includes the pseudo-symmetric configuration, that with index 2 consists of the symmetric vacancy and that with index 4 includes the split vacancy. Other indices are assigned to families that do not comprise these configurations. From Fig. 10b, one can see that in general each family comprises structures with very similar formation energies. However, some exceptions appear: structures belonging to the same family might have a notably different energy and structures with very similar energy might belong to different families. We then selected a representative for each of the nine families and we relaxed its structure employing the HSE15 functional using just the Γ point. All except two (the representatives of family 3 and 1, both belonging to C_9) relaxed to the same split-vacancy configuration reported in ref. [67] (including the various pseudo-symmetric vacancy configurations), showing that many of the minima appearing on the DFT+U PES are indeed spurious and disappear when a more accurate functional is used. The remaining two structures relaxed to a configuration characterized by large displacements of the two oxygen atoms close to $\square_O$, which is neither the simple nor the split vacancy. We name this structure the "delocalized vacancy" configuration, as the extra electrons are delocalized in conduction
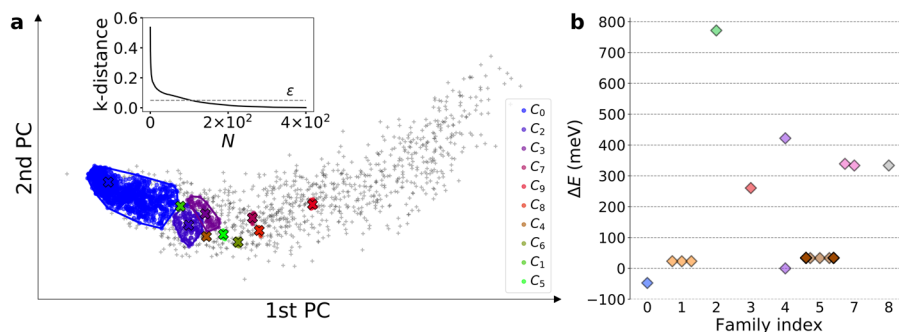
Fig. 10 Structures and local minima found at the DFT+U level. a The clusters found among the visited structures of an EA run for the $\square_O$ in anatase TiO_2. Cluster centroids are shown with crosses. The inset shows the k-distance plot used to find the clustering parameter ε. Cluster names are ordered in the legend according to their size. **b** Distribution of the minima found at the DFT+U level from the selected clusters.

band states, as shown in Fig. 8e, f. We then reoptimize both the split-vacancy and the delocalized vacancy structures employing a larger $2 \times 2 \times 2$ mesh for reciprocal-space integration, and found that both structures remain (meta-)stable with the energy of the delocalized vacancy only ≈ 59 meV higher than that of the split vacancy.

In any case, the existence of the delocalized vacancy configuration completes the atomistic description of the neutral $\square_O$ in TiO_2 anatase. In fact, considering the electronic structure of this defect, three different scenarios are possible: in the simple vacancy the two extra electrons are localized in the vacancy site, in the split vacancy one electron is excited to a delocalized conduction band state, while the other is strongly localized on a neighboring Ti atom, reducing it to the $+3$ charge state and inducing strong lattice distortions. Finally, in the delocalized vacancy, both electrons are delocalized in the conduction band.

DISCUSSION

In this work we have proposed an approach based on a variation of the CMA-ES algorithm optimized for finding low-energy structures of point defects in solids. We have shown that the algorithm is generally able to find the ground state defect configuration for a set of defects showing different electronic and structural properties. Using this algorithm, we were able to find known low-energy defective structures and additional ones as well. The computational costs required by the algorithm are affordable even at the first-principles level, and can be further reduced by employing a GP regressor which is trained on-the-fly during the evolutionary process and is devised to guarantee a minimum accuracy of the approximated PES.

The application of our proposed approach also demonstrates how the DFT+U leads to an overly-complex PES, with several spurious local minima, for the neutral $\square_O$ in TiO_2 anatase, making the convergence of the EA to the global minimum difficult. However, the application of the unsupervised clustering post-processing step allows for an efficient selection of interesting structures, which can then be short-listed an studied with more accurate functionals. Employing this procedure, we were able to discover a low-energy configuration for the oxygen vacancy in TiO_2 anatase, characterized by the presence of two electrons delocalized in the conduction band, which completes the atomistic description of this well-studied defect.

METHODS
DFT calculations

All first-principles calculations were performed employing the VASP code[68], using the projector augmented-wave (PAW) method[69]. For the systems involving bulk silicon the local-density approximation (LDA) for the exchange-correlation functional[10,70] was used. We employed a $2 \times 2 \times 2$ expansion of the conventional cubic cell, obtaining a supercell containing 64 atoms. The basis set included all plane waves with a kinetic energy within 307 eV and the $3s$ and $3p$ electrons were considered as valence ones. Reciprocal-space integrals were approximated by using the Γ-point only.

For TiO_2 anatase, the vast majority of calculations was performed employing the DFT+U method[71], where a U of 5.8 eV was applied to the d electrons of the Ti atoms, keeping all other computational parameters as those described in reference[67], with the exception that during the EA run, only the Γ-point was used. After the EA has converged a $2 \times 2 \times 2$ k-point mesh was used in order to relax the obtained solution to the closest minimum. Finally, for selected structures, we performed also calculations employing the Heyd–Scuseria–Ernzerhof screened hybrid functional (HSE)[72]. The amount of Hartree–Fock exchange in the HSE functional was set to 15% as described in ref. [67] (HSE15). In all cases a supercell with 108 atoms was used.

For all considered systems, relaxation of the atomic coordinates by means of gradient-based methods was done without any constraint on the atomic positions until the remaining forces on each atom were less than 0.01 eV Å^{-1}.

ML models

The Gaussian-process regressor was implemented from the corresponding classes defined in the `gpytorch` library. The `DBSCAN` clustering and PCA were performed employing the implementation available in the `scikit-learn` library.

DATA AVAILABILITY

A dataset containing the structures described in this work, including those generated during the evolutionary processes, is available on Zenodo at reference[74].

CODE AVAILABILITY

The code employed in this study, named `Clinamen`, is available on the GitLab repository: https://gitlab.com/Marrigoni/clinamen and its documentation is hosted on Read the Docs at: https://clinamen.readthedocs.io/

REFERENCES

1. Pizzini, S. *Physical Chemistry of Semiconductor Materials and Processes* (John Wiley & Sons, Ltd, 2015).
2. Queisser, H. J. & Haller, E. E. Defects in semiconductors: some fatal, some vital. *Science* **281**, 945–950 (1998).
3. McCluskey, M. D. & Haller, E. E. *Dopants and Defects in Semiconductors* (CRC Press, 2012).
4. Laks, D. B., Van de Walle, C. G., Neumark, G. F. & Pantelides, S. T. Role of native defects in wide-band-gap semiconductors. *Phys. Rev. Lett.* **66**, 648–651 (1991).
5. Zunger, A. Practical doping principles. *Appl. Phys. Lett.* **83**, 57–59 (2003).
6. Maier, J. Nanoionics: ion transport and electrochemical storage in confined systems. *Nat. Mater* **4**, 805–815 (2005).

7. del Alamo, J. A. Nanometre-scale electronics with iii-v compound semiconductors. *Nature* **479**, 317–323 (2011).

8. Yu, X., Marks, T. J. & Facchetti, A. Metal oxides for optoelectronic applications. *Nat. Mater* **15**, 383–396 (2016).

9. Walsh, A. & Zunger, A. Instilling defect tolerance in new compounds. *Nat. Mater* **16**, 964–967 (2017).

10. Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965).

11. Van de Walle, C. G. & Neugebauer, J. First-principles calculations for defects and impurities: applications to iii-nitrides. *J. Appl. Phys.* **95**, 3851–3879 (2004).

12. Freysoldt, C. et al. First-principles calculations for point defects in solids. *Rev. Mod. Phys.* **86**, 253–305 (2014).

13. Drabold, D. A. & Estreicher, S. *Theory of Defects in Semiconductors* (Springer, 2007).

14. Lany, S. & Zunger, A. Anion vacancies as a source of persistent photoconductivity in ii-vi and chalcopyrite semiconductors. *Phys. Rev. B* **72**, 035215 (2005).

15. Kundu, A. et al. Effect of local chemistry and structure on thermal transport in doped GaAs. *Phys. Rev. Mater.* **3**, 094602 (2019).

16. Eiben, A. E. & Smith, J. E. *Introduction to Evolutionary Computing* (Springer, 2015).

17. Bozorg-Haddad, O., Solgi, M. & Loáiciga, H. A. *Meta-Heuristic and Evolutionary Algorithms for Engineering Optimization* (John Wiley & Sons, Ltd, 2017).

18. Deaven, D. M. & Ho, K. M. Molecular geometry optimization with a genetic algorithm. *Phys. Rev. Lett.* **75**, 288–291 (1995).

19. Jóhannesson, G. H. et al. Combined electronic structure and evolutionary search approach to materials design. *Phys. Rev. Lett.* **88**, 255506 (2002).

20. Oganov, A. R. & Glass, C. W. Crystal structure prediction using ab initio evolutionary techniques: principles and applications. *J. Chem. Phys.* **124**, 244704 (2006).

21. Martinez, U., Vilhelmsen, L. B., Kristoffersen, H. H., Stausholm-Møller, J. & Hammer, B. Steps on rutile TiO$_2$(110): active sites for water and methanol dissociation. *Phys. Rev. B* **84**, 205434 (2011).

22. Vilhelmsen, L. B. & Hammer, B. Systematic study of Au$_6$ to Au$_{12}$ gold clusters on mgo (100) *f* centers using density-functional theory. *Phys. Rev. Lett.* **108**, 126101 (2012).

23. Lysgaard, S., Landis, D. D., Bligaard, T. & Vegge, T. Genetic algorithm procreation operators for alloy nanoparticle catalysts. *Top. Catal.* **57**, 33–39 (2014).

24. Merte, L. R. et al. Structure of the SnO$_2$ (110) − 4 × 1 surface. *Phys. Rev. Lett.* **119**, 096102 (2017).

25. Oganov, A. R., Pickard, C. J., Zhu, Q. & Needs, R. J. Structure prediction drives materials discovery. *Nat. Rev. Mater.* **4**, 331–348 (2019).

26. Kaczmarowski, A., Yang, S., Szlufarska, I. & Morgan, D. Genetic algorithm optimization of defect clusters in crystalline materials. *Comput. Mater. Sci.* **98**, 234–244 (2015).

27. Patra, T. K. et al. Defect dynamics in 2-D MoS$_2$ probed by using machine learning, atomistic simulations, and high-resolution microscopy. *ACS Nano* **12**, 8006–8016 (2018).

28. Atilgan, E. & Hu, J. First-principle-based computational doping of SrTiO$_3$ using combinatorial genetic algorithms. *Bull. Mater. Sci.* **41**, 1–9 (2018).

29. Cheng, Y., Zhu, L., Zhou, J. & Sun, Z. pygace: combining the genetic algorithm and cluster expansion methods to predict the ground-state structure of systems containing point defects. *Comput. Mater. Sci.* **174**, 109482 (2020).

30. Morris, A. J., Pickard, C. J. & Needs, R. J. Hydrogen/silicon complexes in silicon from computational searches. *Phys. Rev. B* **78**, 184102 (2008).

31. Pickard, C. J. & Needs, R. J. Ab initiorandom structure searching. *J. Phys. Condens. Matter* **23**, 053201 (2011).

32. Mulroue, J., Morris, A. J. & Duffy, D. M. Ab initio study of intrinsic defects in zirconolite. *Phys. Rev. B* **84**, 094118 (2011).

33. Hansen, N. & Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *Evol. Comput.* **9**, 159–195 (2001).

34. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).

35. Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **104**, 136403 (2010).

36. Rupp, M., Tkatchenko, A., Müller, K.-R. & von Lilienfeld, O. A. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **108**, 058301 (2012).

37. Li, Z., Kermode, J. R. & De Vita, A. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Phys. Rev. Lett.* **114**, 096405 (2015).

38. Bartók, A. P., Kermode, J., Bernstein, N. & Csányi, G. Machine learning a general-purpose interatomic potential for silicon. *Phys. Rev. X* **8**, 041048 (2018).

39. Jennings, P. C., Lysgaard, S., Hummelshøj, J. S., Vegge, T. & Bligaard, T. Genetic algorithms for computational materials discovery accelerated by machine learning. *Npj Comput. Mater.* **5**, 46 (2019).

40. Bisbo, M. K. & Hammer, B. Efficient global structure optimization with a machine-learned surrogate model. *Phys. Rev. Lett.* **124**, 086102 (2020).

41. Hansen, N. & Kern, S. in *Parallel Problem Solving from Nature - PPSN VIII* (eds Yao, X. et al.) (Chapter 29, Springer, 2004).

42. Hansen, N. *The CMA Evolution Strategy: A Comparing Review* (Springer, 2006).

43. Hansen, N. The cma evolution strategy: a tutorial. Preprint at arXiv:1604.00772 (2016).

44. References to cma-es applications. http://www.cmap.polytechnique.fr/nikolaus.hansen/cmaapplications.pdf (2009).

45. Pun, G. P. P. & Mishin, Y. Optimized interatomic potential for silicon and its application to thermal stability of silicene. *Phys. Rev. B* **95**, 224103 (2017).

46. Leung, W.-K., Needs, R. J., Rajagopal, G., Itoh, S. & Ihara, S. Calculations of silicon self-interstitial defects. *Phys. Rev. Lett.* **83**, 2351–2354 (1999).

47. Goedecker, S., Deutsch, T. & Billard, L. A fourfold coordinated point defect in silicon. *Phys. Rev. Lett.* **88**, 235501 (2002).

48. Mattsson, A. E., Wixom, R. R. & Armiento, R. Electronic surface error in the Si interstitial formation energy. *Phys. Rev. B* **77**, 155211 (2008).

49. Ganchenkova, M. G. et al. Influence of the ab-initio calculation parameters on prediction of energy of point defects in silicon. *Mod. Electron. Mater.* **1**, 103–108 (2015).

50. Rinke, P., Janotti, A., Scheffler, M. & Van de Walle, C. G. Defect formation energies without the band-gap problem: combining density-functional theory and the *GW* approach for the silicon self-interstitial. *Phys. Rev. Lett.* **102**, 026402 (2009).

51. Bruneval, F. Range-separated approach to the RPA correlation applied to the van der waals bond and to diffusion of defects. *Phys. Rev. Lett.* **108**, 256403 (2012).

52. Gao, W. & Tkatchenko, A. Electronic structure and Van der Waals interactions in the stability and mobility of point defects in semiconductors. *Phys. Rev. Lett.* **111**, 045501 (2013).

53. Cargnoni, F., Gatti, C. & Colombo, L. Formation and annihilation of a bond defect in silicon: an ab initio quantum-mechanical characterization. *Phys. Rev. B* **57**, 170–177 (1998).

54. Valle, M. & Oganov, A. R. Crystal fingerprint space – a novel paradigm for studying crystal-structure sets. *Acta Cryst. A* **66**, 507–517 (2010).

55. Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proc. Second International Conference on Knowledge Discovery and Data Mining,*(eds Simoudis, E., Han, J. & Fayyad, U) KDD'96. 226–231 (AAAI Press, 1996).

56. Rasmussen, C. & Williams, C. *Gaussian Processes for Machine Learning* (MIT Press, 2006).

57. Denzel, A. & Kästner, J. Gaussian process regression for geometry optimization. *J. Chem. Phys.* **148**, 094114 (2018).

58. Koistinen, O.-P., Ásgeirsson, V., Vehtari, A. & Jónsson, H. Nudged elastic band calculations accelerated with Gaussian process regression based on inverse interatomic distances. *J. Chem. Theory Comput.* **15**, 6738–6751 (2019).

59. Todorović, M., Gutmann, M. U., Corander, J. & Rinke, P. Bayesian inference of atomistic structure in functional materials. *Npj Comput. Mater.* **5**, 35 (2019).

60. Garijo del Río, E., Mortensen, J. J. & Jacobsen, K. W. Local Bayesian optimizer for atomic structures. *Phys. Rev. B* **100**, 104103 (2019).

61. Järvi, J., Rinke, P. & Todorović, M. Detecting stable adsorbates of (1s)-camphor on cu (111) with Bayesian optimization. *Beilstein J. Nanotechnol.* **11**, 1577–1589 (2020).

62. Huo, H. & Rupp, M. Unified representation of molecules and crystals for machine learning. Preprint at arXiv:1704.06439 (2017).

63. Himanen, L. et al. Dscribe: library of descriptors for machine learning in materials science. *Comput. Phys. Commun.* **247**, 106949 (2020).

64. Finazzi, E., Di Valentin, C., Pacchioni, G. & Selloni, A. Excess electron states in reduced bulk anatase TiO$_2$: comparison of standard GGA, GGA+U, and hybrid DFT calculations. *J. Chem. Phys.* **129**, 154113 (2008).

65. Mattioli, G., Filippone, F., Alippi, P. & Amore Bonapasta, A. Ab initio study of the electronic states induced by oxygen vacancies in rutile and anatase TiO$_2$. *Phys. Rev. B* **78**, 241201 (2008).

66. Morgan, B. J. & Watson, G. W. Intrinsic n-type defect formation in TiO$_2$: a comparison of rutile and anatase from GGA+U calculations. *J. Phys. Chem. C* **114**, 2321–2328 (2010).

67. Arrigoni, M. & Madsen, G. K. H. A comparative first-principles investigation on the defect chemistry of tio2 anatase. *J. Chem. Phys.* **152**, 044110 (2020).

68. Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169–11186 (1996).

69. Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B* **50**, 17953–17979 (1994).

70. Ceperley, D. M. & Alder, B. J. Ground state of the electron gas by a stochastic method. *Phys. Rev. Lett.* **45**, 566–569 (1980).

71. Anisimov, V. I., Solovyev, I. V., Korotin, M. A., Czyżyk, M. T. & Sawatzky, G. A. Density-functional theory and NiO photoemission spectra. *Phys. Rev. B* **48**, 16929–16934 (1993).

72. Heyd, J., Scuseria, G. E. & Ernzerhof, M. Hybrid functionals based on a screened coulomb potential. *J. Chem. Phys.* **118**, 8207–8215 (2003).

73. Momma, K. & Izumi, F. *VESTA3* for three-dimensional visualization of crystal, volumetric and morphology data. *J. Appl. Crystallogr.* **44**, 1272–1276 (2011).

74. Arrigoni, M. & Madsen, G. K. H. Evolutionary computing and machine learning for the discovering of low-energy defect configurations. *Zenodo* https://doi.org/10.5281/zenodo.4265094 (2020).

ACKNOWLEDGEMENTS

AUTHOR CONTRIBUTIONS

M.A. conceived the initial study and developed the methodology, wrote the code and performed the calculations. G.K.H.M. supervised the project, contributed in finalizing the methodological approach and provided the funds. Both authors contributed equally in analyzing the results and in writing and revising the paper.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41524-021-00537-1.

Correspondence and requests for materials should be addressed to M.A.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.