Comprehensive structure exploration and thermodynamics of graphitic heteroatom doped graphene superstructures

Benedict Saunders
Department of Chemistry, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, United Kingdom

Lukas Hörmann* and Reinhard J. Maurer†
Department of Chemistry, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, United Kingdom;
Department of Physics, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, United Kingdom;
and Faculty of Physics, University of Vienna, Kolingasse 14-16, 1090 Vienna, Austria

(Received 19 September 2025; revised 5 October 2026; accepted 8 October 2026; published 2 November 2026)

Graphene has been studied in detail due to its mechanical, electrical, and thermal properties. It is well documented that the introduction of dopants or defects in the lattice can be used to tune material properties for a specific application, such as in electronics, sensors, or catalysis. To design graphene with specific properties, one must achieve control over the composition and concentration of defects. This requires a fundamental understanding of the stability of defects and their interaction in a superstructure. We present a comprehensive defect structure determination approach that enables close to exhaustive enumeration of all relevant defect structures. The approach uses a combination of density functional theory and machine learning to build a transferable energy model for defect formation. We show the capabilities of our approach for free-standing graphene with heteroatom defects, establishing a thermodynamic model to investigate how temperature affects the configuration space of doped graphene. Our analysis yields physical insights into defect interactions. A characteristic peak in the heat capacity indicates a structural transition. We provide a mechanistic explanation for this behavior and demonstrate how the transition temperature shifts with different thermodynamic reservoirs.

DOI: 10.1103/db9g-7gvs

I. INTRODUCTION

Since its isolation in 2004 [1], graphene has garnered significant interest owing to its remarkable electronic [2,3], thermal [4], and mechanical properties [5]. Pristine, undoped, free-standing graphene is a zero-gap semiconductor, evidenced by its characteristic Dirac cones at the K point. A significant amount of research over the past two decades has been dedicated to modifying the structure of graphene to include various dopant atoms or molecules, and even vacancies, with the aim of effectively tuning the chemistry of the material to achieve a particular function [6–8]. Altering the composition or topology of graphene can be used as a means to tune its band structure [9,10]. Doping with heteroatoms has been shown to open the band gap, prompting extensive research on how various dopant configurations affect the band structure [11].

Due to their semiconducting properties, graphene systems doped with heteroatoms have been studied meticulously for their potential applications as photocatalysts [12–14], gas sensors [15,16], as well as hosts for single-atom catalysts (SAC) [17–20]. Nitrogen-doped graphene (NDG), in particular, graphitic nitrogen, where a carbon atom is substituted by a nitrogen atom, is well studied; however, to achieve control over the doping level and dopant dispersity during growth requires understanding of the energetics and thermodynamics of the system in addition to the growth mechanism in the presence of the doping agent. A significant problem that arises when modeling such systems is the sheer number of possible configurations that arise from doping graphene with even only a handful of nitrogen atoms. Various ways of experimentally growing NDG have been explored, including precursor molecules [21], segregation from embedded carbon and nitrogen sources [22], and sputtering [23]. Theory-based studies have investigated the carbon-nitrogen phase diagram, probing the limits of sustainable doping [24], identifying stable NDG structures [25], and demonstrating how the choice of nitrogen and carbon sources can enable the growth of high-concentration NDG [26]. Methods such as cluster expansion, genetic algorithms [24], or particle swarm optimizations [25,26] enabled the exploration of dopant interactions and the material's potential energy surface. However, these approaches are mainly focused on the thermodynamic ground state, seeking the lowest-energy configuration at each composition. The focus on only the lowest-energy configuration is insufficient for investigating the thermodynamic properties of a system, which are determined by an ensemble of possible configurations, and not just the lowest-energy state. Nested sampling offers a promising route to access thermodynamic properties [27,28], but it is less suited to crystalline

*Contact author: lukas.hoermann@univie.ac.at
†Contact author: reinhard.maurer@univie.ac.at

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

materials, where large regions of configuration space are sterically inaccessible. Especially in large-scale defect superstructures, the vast combinatorial complexity remains a major limiting factor.

Focusing on NDG, here we resolve this combinatorial challenge by leveraging a recent machine-learning-assisted structure search method, SAMPLE [29], and hence extract thermodynamic information about the system through the application of statistical mechanics. SAMPLE utilizes Bayesian regression to construct an efficient model for predicting the formation energy as a sum of onsite and pairwise interaction contributions. The final model is highly efficient, enabling brute-force structure search through combinatorial enumeration of millions of defect structures. SAMPLE was originally developed to investigate the configuration space of adsorbed molecules on flat surfaces. To efficiently explore the vast configurational and compositional landscapes of NDG, we have modified the SAMPLE code, extending its capability to predict the formation energies of defective two-dimensional (2D) materials. By considering all possible configurations and their energies as individual states within a statistical mechanical framework, we gain insight into the thermodynamic behavior of the material and enable the calculation of temperature-dependent properties, such as the heat capacity, which in turn can be used to predict the phase stability.

II. METHODS

A. System setup and assumptions

Our thermodynamic analysis is an approximation of a chemical vapor deposition (CVD) experiment. The CVD growth process occurs at elevated temperatures (1000–1300 K) and doped configurations are thermodynamically and kinetically accessible, permitting exploration of the rich configurational space. Once the growth process is complete, the system is slowly cooled to a given temperature, allowing it to settle into a distribution of structures given by our thermodynamic predictions. Our calculated heat capacity, therefore, does not correspond to a low-temperature dynamic transition but rather provides a thermodynamic descriptor of the relative populations of different NDG configurations created upon cooling. NDG can exhibit two primary classes of defects: substitutional defects, where a nitrogen atom replaces a carbon atom, and intrinsic lattice defects, which are structural flaws in the carbon lattice [30–32]. Graphitic nitrogen is experimentally known to be the dominant nitrogen species under high-temperature CVD growth and postannealing conditions, where substitutional nitrogen becomes the most abundant configuration [32–34]. Motivated by this, the main paper focuses on graphitic nitrogen-doped graphene. However, SAMPLE itself is agnostic to the specific defect type. To illustrate this generality, Sec. S3.6 of the Supplemental Material (SM) [35] includes compositional phase diagrams generated with a SAMPLE model for tripyridinic nitrogen defects.

B. Formation energy calculations

First-principles electronic structure calculations are performed using density functional theory (DFT) to generate training sets, validation sets, and reference potentials. All DFT calculations were performed using the all-electron, atom-centered numerical orbital code FHI-aims [36,37] at the PBE [38] level of theory, with the addition of the Tkatchenko-Scheffler Van der Waals correction method [39]. Graphene is described as a free-standing monolayer with a perpendicular vacuum slab of 15.0 Å, which was determined to be sufficient based on convergence testing (see Secs. S1 and S2.2 of the SM [35] for convergence tests). Custom k grids, based on a generalized Monkhorst-Pack scheme [40–42], were used to ensure uniform k-grid sampling across different unit-cell shapes and sizes. Further details are provided in Sec. S1 of the SM [35].

All DFT calculations were conducted using a *tight* default basis set. The energy tolerance was set to 1.0×10^{-6} eV for electronic convergence, and the geometry optimizations were performed using a force tolerance of 0.01 eV Å^{-1}. The angles of the graphene unit cell were fixed during geometry optimizations. All initial graphene geometries were generated from the experimentally determined unit cell [43].

C. SAMPLE for defects in 2D materials

A crucial step for investigating the thermodynamic properties of defective 2D materials is the generation of a comprehensive set of possible defect superstructures. SAMPLE tackles the combinatorial challenge posed by surface defects by restricting the degrees of freedom that define the search space, achieved through a process of coarse graining, which approximates the continuous search space by treating adsorption sites as positions on a grid with discrete translations and rotations of the sites. With such an approximation, the search space of a fixed substrate supercell with a fixed number of adsorbates becomes finite. This is enforced by requiring the substrate supercell and the adsorbate unit cell to be commensurate. With this, the vectors that describe the netplane of the substrate **s** and the adsorbate, or defect in our case, **d** can be mapped onto one another with

$$\begin{pmatrix}\mathbf{d}_1\\ \mathbf{d}_2\end{pmatrix} = \mathbf{M}\cdot\begin{pmatrix}\mathbf{s}_1\\ \mathbf{s}_2\end{pmatrix} \tag{1}$$

where all elements of the 2×2 epitaxial matrix $\mathbf{M}$ must be integers, and $\det(\mathbf{M}) \neq 0$ [44]. By defining a substrate supercell and a concentration of defects, we can enumerate all possible structures for that combination and keep those that are symmetrically unique. Just as we can define an adsorbate geometry commensurate with the substrate supercell, we can define a defect geometry commensurate with the graphene supercell. In the case of a single graphitic nitrogen atom, the unit cell used to describe the defect template is equivalent to that of pristine graphene [Figs. 1(a) and 1(b)]. By approximating surface-adsorbate geometries as large-scale commensurate structures, SAMPLE is well suited to model systems such as doped graphene, due to its highly regular lattice.

SAMPLE stores geometric information in a configuration hash, a short tuple of integer numbers, representing the epitaxial matrix of the supercell, the positions, and types of defects. The positions are given by a unique number scheme of the possible defect sites in the substrate lattice. The defect types are indices pointing to a lookup table. This table stores the

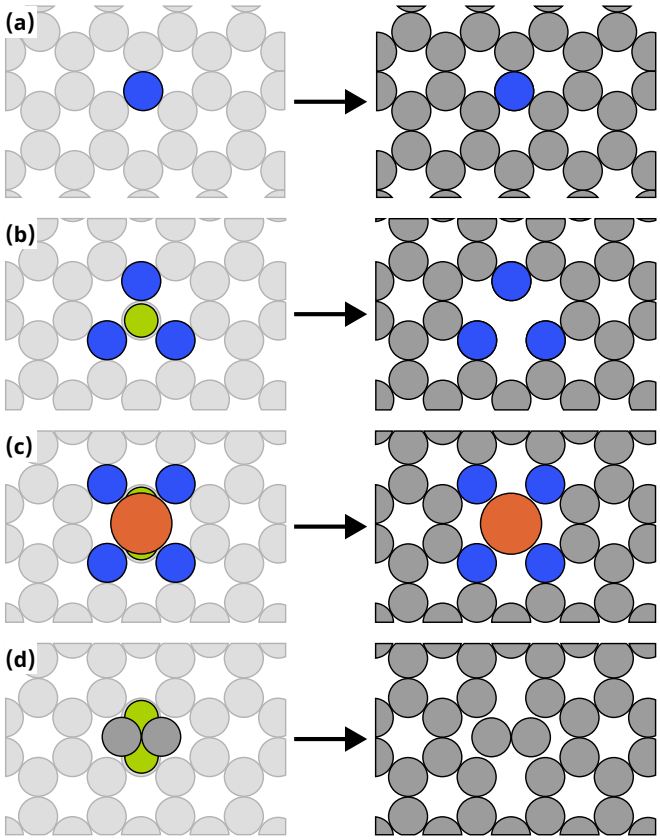


FIG. 1. Various defect templates for graphene (first column) and the respective defective structure created with the template (second column): (a) graphitic nitrogen; (b) tripyridinic nitrogen; (c) a FeN_4 metal complex; (d) a Stone-Wales defect. In (b), (c), and (d) the Rm atom type (yellow green) is added to indicate the removal of an atom to form a vacancy.

optimized geometric structures of the defects (in isolation), called defect templates. The configuration hash has a one-to-one correspondence to the structure and is invariant under symmetry operations applied to the configuration.

To enable the automated generation of 2D-defect superstructures, we made two modifications to the SAMPLE code. The first enables the substitution of atoms in the substrate, such as one would need for a graphitic defect: any atoms in the substrate geometry that overlap with an atom in the defect template, within some tolerance, are automatically overwritten. The second modification is the introduction of a new atomic species that can be included in the defect geometry: removium, designated Rm. As the name suggests, it deletes any atom at the location of the Rm, but otherwise behaves the same as other atoms (e.g., adhering to symmetry operations). As such, Rm can effectively create vacancies in the substrate lattice, but can also be used to delete atoms to create more complex defects too, such as FeN_4 [45], or the Stone-Wales defect [Figs. 1(c) and 1(d)].

When introducing a defect into a lattice (or an adsorbate onto a surface), the SAMPLE approach assumes that the substrate lattice remains unchanged, and all lattice atoms (apart from ones removed or replaced) retain their original positions in the undoped substrate supercell. Typically, the inclusion of defects results in rearrangements in the lattice, but this is neglected by SAMPLE, resulting in an overestimation of formation energies. The lattice distortion that occurs when relaxing a single heteroatom-defect-structure is small (see Sec. III A and Secs. S2.2 and S2.3 of the SM [35]). This is consistent with previous reports: substitutional graphitic nitrogen slightly reduces the lattice constant [46], while recent studies reported sublattice symmetry breaking and C-C bond-length variation with dopant number, although they do not observe large global shear or angular deviations [47,48]. While high defect concentrations can more substantially alter the structure and stability of NDG, these configurations have a negligible impact on the thermodynamic properties, as their Boltzmann factors are small. Therefore, we expect that the assumption of a fixed unit cell does not significantly affect the thermodynamics of graphitic-doped NDG. However, some defects, such as the FeN_4 or tripyridinic substructures shown in Fig. 1, are expected to induce significant lattice distortions. For this reason, we have elected to focus primarily on graphitic nitrogen defects in this study. In Sec. S3 of the SM [35], we also demonstrate the modeling of tripyridinic nitrogen defects with this approach, albeit only for a restricted region of the configurational space.

D. Learning the formation energies

SAMPLE leverages Bayesian regression for accelerated evaluation of formation energies E_i per defect. For consistency with later chapters, we introduce the index i, which denotes the ith configuration. For graphitic nitrogen defects in graphene, a carbon atom is replaced by a nitrogen atom. We assume the system is in contact with reservoirs for carbon and nitrogen, therefore representing the system in a grand canonical ensemble. The energy required to exchange species with the reservoir is given by the respective chemical potentials. While the SAMPLE approach is agnostic to the choice of chemical potential, the choice affects the predicted thermodynamic behavior (see Fig. 4). To demonstrate this, we select different chemical potentials:

Carbon chemical potential (μ_C)

(i) Graphite: The partitioned atomic energy of carbon in bulk graphite is a standard literature reference for carbon systems [26,47,49–51].

(ii) Methane: The partitioned atomic energy of carbon in methane is used to reflect CVD growth of graphene, which typically utilizes small hydrocarbon precursors [52].

Nitrogen chemical potential (μ_N)

(i) Dinitrogen: The partitioned atomic energy of nitrogen in molecular nitrogen is the conventional reference [47,51], though the extreme stability of the N≡N triple bond yields a very low μ_N, which often misrepresents reactive environments.

(ii) Ammonia, pyridine: The partitioned atomic energy of nitrogen in ammonia or pyridine is representative of the reactive nitrogen-containing molecules commonly supplied in CVD experiments [52].

We define the partitioned atomic energy as the effective DFT energy of the target species within its reservoir. For single-element reservoirs (dinitrogen, graphite) this is simply the energy per atom. For molecular reservoirs it is obtained

by partitioning the total DFT energy of the molecule among its constituent atoms, using reference energies for the remaining species and assuming additive bond contributions [53] (see Sec. S2.1 of the SM [35]).

In all cases where the chemical potentials used are not explicitly stated, dinitrogen-graphite is used. The exact values of the chemical potentials, μ_C and μ_N, used in Eq. (2) are given in Table S1 [35]. The formation energy can be written as follows:

$$E_i = \frac{1}{N_i}(E_{SYS} - E_{GR} + N_C\mu_C - N_N\mu_N). \tag{2}$$

Here E_{SYS} and E_{GR} are the energies of the defective and pristine graphene sheet, respectively. N_i is the number of defects within the supercell adopted in the computation, while N_N and N_C are the numbers of nitrogen atoms and carbon vacancies associated with these defects. To learn and predict the formation energy per defect E_i, SAMPLE uses a cluster-expansion energy model with one-body energy interactions g and two-body interaction energies between pairs of defects p:

$$E_i = \sum_g n_g U_g + \sum_p n_p V_p. \tag{3}$$

Here n_g and n_p specify the occurrences of each of the different one- and two-body interactions within a given configuration. The values in the vector $\mathbf{U}$ correspond to the dopant-lattice interaction in isolation, analogous to the adsorption energy of a single molecule on a surface, and $\mathbf{V}$ discretizes the interaction energies between pairs of defects at different distances. Once trained, these vectors can be considered as the formation energies of isolated defects and the energies of pairs of defects on the coarse-grained lattice, respectively. Concatenating $\mathbf{U}$ and $\mathbf{V}$ into a single interaction vector $\boldsymbol{\omega}$, together with the model vector $\mathbf{n}$, which lists the occurrences of individual interactions in a configuration, the energy per dopant of a configuration can be evaluated using linear regression:

$$E_i = \mathbf{n}_i \cdot \boldsymbol{\omega}. \tag{4}$$

As $\mathbf{U}$ and $\mathbf{V}$, hence $\boldsymbol{\omega}$, contain the energy contributions of every possible combination of one- and two-body terms considered by a given SAMPLE model, only $\mathbf{n}$ changes with configuration. As such, all of the configurations to be considered can be agglomerated into a single model matrix $\mathbf{X}$, such that the energies of a configuration set can be simply expressed as

$$\mathbf{E} = \mathbf{X} \cdot \boldsymbol{\omega}. \tag{5}$$

SAMPLE is trained using Bayesian regression to fit the vector $\boldsymbol{\omega}$. The reason for using this method is twofold: First, fitting $\boldsymbol{\omega}$ is an underdetermined problem, and there are far more interactions in the configuration space than one could reasonably have in the training data. This means that we can fit the model on only a few hundred DFT data points. Second,

$$p(\boldsymbol{\omega}|\boldsymbol{E}_{DFT}) = \frac{p(\boldsymbol{E}_{DFT}|\boldsymbol{\omega})p(\boldsymbol{\omega})}{p(\boldsymbol{E}_{DFT})}. \tag{6}$$

Bayes' theorem, given in Eq. (6), allows SAMPLE to ascertain the expectation values and respective probability distributions of the one- and two-body vectors. In SAMPLE's case, the denominator $p(\boldsymbol{E}_{DFT})$, termed the *marginal probability*, can be neglected since it is simply a normalization constant. The *prior*, $p(\boldsymbol{\omega})$, and the *likelihood*, $p(\boldsymbol{E}_{DFT}|\boldsymbol{\omega})$, are both Gaussian distributions, and therefore their product equals the *posterior*, $p(\boldsymbol{\omega}|\boldsymbol{E}_{DFT})$. The likelihood determines the probability that a given interaction vector $\boldsymbol{\omega}$ will reproduce the energies $\boldsymbol{E}_{DFT}$ calculated with DFT for a given set of configurations (i.e., the training data) within a predefined tolerance. This training set consists of a few hundred DFT calculations of defect superstructures that are sampled from the comprehensive configuration space discussed in the previous section, through D-optimal selection [54]. The prior, on the other hand, provides the physical knowledge about the system, and as such, we can bias the likelihood such that the model may give more accurate predictions, typically using data from DFT calculations. A more thorough explanation of the SAMPLE method can be found in the original paper [29].

E. Thermodynamics with formation energies

The ability of SAMPLE to comprehensively map the configuration space enables brute-force calculation of the partition function and all thermodynamic properties that can be derived from it. For each size of the simulation cell (see Sec. S3.1 of the SM [35]), the composition space can be treated as a grand canonical ensemble for fixed values of μ_C and μ_N. In our thermodynamic formulation, each microstate corresponds to a specific NDG configuration. To make systems with different coverages comparable, we determine the energy per simulation cell:

$$\epsilon_i = E_i N_i. \tag{7}$$

Using the formation energy of the simulation cell, we construct the grand canonical partition function for each cell size. Hereby, we make use of the fact that the formation energy already includes the chemical potentials that govern particle exchange between the system and reservoirs [see Eq. (2)]. This approach is appropriate because the simulation cell defines the fundamental periodic unit that accounts for the total energetic cost of a specific configuration, including many-body interactions and surface strain. By treating the entire cell as a single state in the ensemble, we ensure that the competition between different phases is governed by the total free energy of the interface, rather than being biased by the varying number of constituent atoms in different supercell geometries:

$$\mathcal{Z} = \sum_i \exp(-\beta\epsilon_i). \tag{8}$$

In this equation, $\beta = \frac{1}{k_B T}$, k_B is the Boltzmann constant, and T is the temperature of the system. Boltzmann averages of observables can be calculated directly. We can calculate the Boltzmann probability corresponding to each microstate i:

$$P_i = \frac{\exp(-\beta\epsilon_i)}{\mathcal{Z}} \tag{9}$$

which we can then use to calculate expected formation energy and particle concentration (normed to intensive energy per

lattice site) as

$$\langle E_{\text{form}} \rangle = \frac{1}{N_{\text{tot}}} \sum_i P_i \epsilon_i. \tag{10}$$

However, when calculating the heat capacity in the grand canonical ensemble, fluctuations of defect concentration $\theta_i = \frac{N_i}{N_{\text{tot}}}$ need to be considered, where N_i is the number of defects and N_{tot} is the total number of atoms in the configuration. This is given by the variance $\text{Var}(\theta) = \langle \theta^2 \rangle - \langle \theta \rangle^2$ of the concentration and the covariance between the energy and concentration $\text{Cov}(\epsilon, \theta) = \langle \theta \epsilon \rangle - \langle \theta \rangle \langle \epsilon \rangle$ (see Sec. S6 of the SM [35] for a detailed derivation). The heat capacity of a grand canonical ensemble can hence be expressed as follows:

$$C_{\text{V}} = \frac{1}{k_B T^2} \left[\text{Var}(\epsilon) - \frac{\text{Cov}(\epsilon, \theta)^2}{\text{Var}(\theta)} \right] \tag{11}$$

and hence the specific heat capacity is calculated trivially as

$$c_{\text{V}} = \frac{C_{\text{V}}}{N_{\text{tot}}}. \tag{12}$$

Applying these expressions, we can explore the thermodynamic behavior of the system for a given surface area under fluctuating numbers of defects and hence gain further insight into the thermodynamic stability and the response to temperature changes.

III. RESULTS

The results of this study are divided into three parts. First, we report an investigation into the interactions of defects in graphene, as predicted by DFT. The SAMPLE model is fitted based on these data and validated against the DFT reference data. Second, we apply our SAMPLE model to explore the vast configuration space of doped graphene, a task that would be computationally prohibitive based solely on DFT. Finally, the derivation of thermodynamic properties of graphene doped with graphitic-nitrogen environments provides insights into the stability and phase behavior of the material under varying conditions.

A. Defect-defect interactions

When investigating the configuration space of a doped material, two main energy contributions must be considered: (i) the intrinsic formation energy of each defect in the host lattice, and (ii) the interaction energy between defects, which becomes particularly significant when defects are in close proximity. Both contributions play a crucial role in determining the overall formation energy and thus the underlying chemistry and physics of the system.

DFT calculations. The intrinsic formation energy $E_{\text{intrinsic}}$ is the energy required to exchange carbon for nitrogen in the graphene plane. We determine it using Eq. (2) for a unit cell in which only a single defect is placed. The unit cell is chosen large enough that the defects no longer interact with each other (see Sec. S2.2 of the SM [35]). For the interaction energy, we can apply

$$E_{\text{interaction}} = E_{\text{pair}} - 2E_{\text{intrinsic}}, \tag{13}$$

where E_{pair} is the formation energy [calculated with Eq. (2)] of a structure containing a pair of defects.

We find that the formation energy of an isolated graphitic nitrogen site in graphene in the dilute limit is 0.80 eV, calculated using a 12×12 supercell to ensure a safe margin beyond convergence (see Sec. S2.4 [35], and Refs. [26,47,50,51,55–59] therein). We also probed the effect of geometry optimizations on the formation energy by placing a single nitrogen defect into supercells of varying sizes. For all supercells larger than the primitive graphene unit cell, the total energy of the doped systems decreases between 0.03 and 0.04 eV. Therefore, the lattice relaxation does not significantly affect the formation energies of isolated graphitic nitrogen defects and the interaction energy with periodic replicas of the defects. The dilute limit is reached with a 11×11 supercell, where the smallest distance between two nitrogen defects is 14.8 Å (see Fig. S2 of the SM [35]).

With a term for the isolated formation energy, we can now calculate various interaction energies of graphitic nitrogen sites with DFT. Figure 2(a) shows the interaction energy of multiple pairs of defects as a function of separation distance, covering 17 of the symmetrically unique interactions between 1.4 and 10.7 Å in a 16×16 graphene supercell. Although there is a clear overall trend of decreasing interaction energy with increasing distance, the trend itself is very much nontrivial. Unsurprisingly, the largest interaction energy is exhibited by the configuration in which the two defects are on neighboring lattice sites, 1.20 eV.

The interaction profile between two dopant pairs in the lattice is nontrivial. To show this, we divide pair configurations into three categories: *Z* configurations follow a zigzag path, *H* configurations effectively hop over consecutive rings, and *O* configurations are some combination of the two [see Fig. 2(a)]. The *Z* configurations exhibit an interaction energy profile that closely follows an exponential decay with increasing distance. In contrast, the *H* configurations include some of the most stable defect arrangements, particularly *H1*, which represents a significant local energy minimum after hopping over a single hexagonal ring, and *H3*, which corresponds to a second local minimum resembling a 4,4′-bipyridine-like system. The *O* configurations, which combine elements of both *Z* and *H* pathways, also show unexpectedly stable arrangements, particularly *O1* and *O4*.

Additionally, Fig. 2(a) shows the gain in interaction energy upon geometry optimization. This energy gain decreases with increasing pair distance and becomes negligible at larger separations that characterize energetically favorable structures. As detailed in Secs. S2.3 and S4.1 of the SM [35] (see also Ref. [60] therein), geometry relaxations can be neglected given our objective of investigating the qualitative thermodynamic behavior of NDG and given the fact that structure relaxation most strongly affects close distances with high energies that are less likely to contribute to low-energy configurations.

SAMPLE. Based on DFT reference data, a SAMPLE model was constructed to determine the formation energies of graphene supercells doped with graphitic nitrogen defects. Against the validation set, the model exhibits a root-mean-square error of 0.022 eV per defect for the formation energy prediction. Details regarding specific hyperparameters in the energy model can be found in Sec. S3.4 of the SM [35]

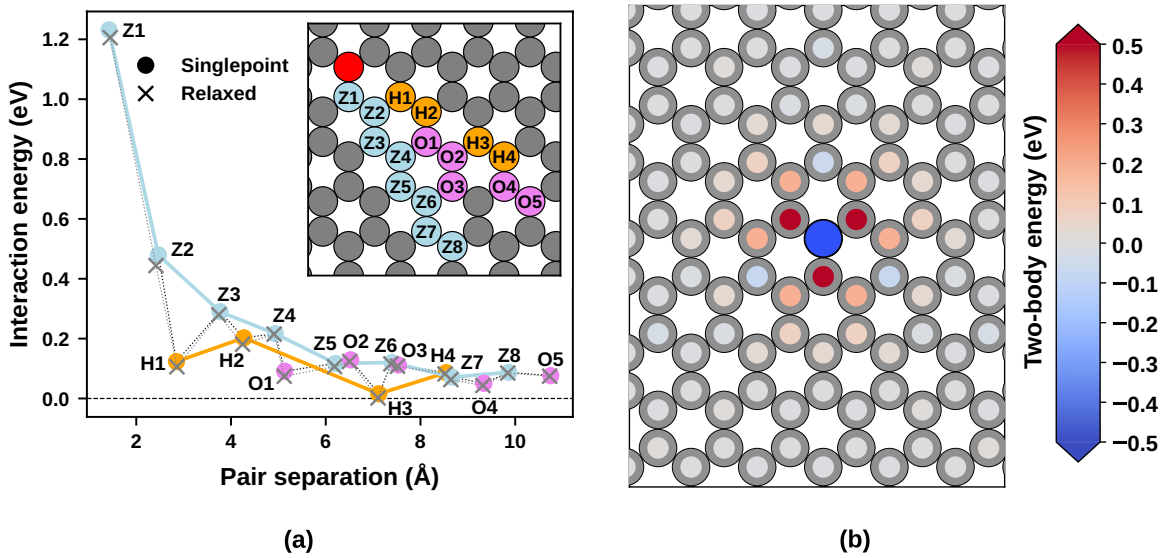


FIG. 2. Nitrogen defect interaction energy with DFT (a) and two-body energies from SAMPLE model (b). The DFT interaction energies are split into three groups according to their relative position in the lattice: Z follows a zigzag pattern, H "hops" over rings, and O configurations are combinations of both moves. Some of the learned SAMPLE two-body energies in (b) exceed 0.50 eV; however, the color scale has been shifted to emphasize that SAMPLE was able to infer the nontrivial interactions as observed with DFT.

(see also Ref. [61] therein). The two-body energies learned by SAMPLE are an approximation of the defect-defect interaction energies, and are illustrated as a heatmap in Fig. 2(b). The two-body energies follow the trends observed in the DFT data in Fig. 2(a), but are slightly offset from the DFT interaction energies. This is because the two-body energies and the DFT interaction energies have different references. As this offset is added to the one-body energies, this offset does not affect the total energies predicted with Eq. (3). Further details regarding the model's accuracy are provided in Sec. S3.5 of the SM [35].

B. Configuration space exploration

To explore the configuration space of NDG, we generated a total of 162 239 485 configurations with a cell size of up to 100 graphene lattice sites. The nitrogen concentration ranges from 1% to 100%, with sampling biased toward low-nitrogen-content configurations (see Sec. S3.1 of the SM [35]), as these are energetically more favorable and are expected to dominate the thermodynamic properties (see Sec. III C). Symmetrically equivalent structures were screened out using a method described in Sec. S3.2 of the SM [35] (see also Refs. [62–64] therein), reducing the dataset to 105 614 115. Applying the SAMPLE model that we constructed above, we can illustrate the entire configuration space of graphitic nitrogen dopants by considering the formation energy as a function of nitrogen content, as in Fig. 3, where we use dinitrogen and graphite as the chemical potentials for nitrogen and carbon, respectively. We can describe the topology of the compositional phase diagram with a convex hull and, as such, the vertical distance from a given point to the hull is used as a measure of relative stability within the configuration space.

Overall, configurations with the fewest and most sparsely distributed dopants exhibit the lowest formation energies. However, the wide range of formation energies at each value of θ highlights that the specific atomic configuration, not just the defect concentration, plays a crucial role in determining the energetic favorability. For example, in Fig. 3 at $\theta = 0.5$, the most stable configuration exhibits an alternating arrangement, where the nitrogen and carbon atoms occupy graphene's inherent A and B sublattice sites, respectively, whereas the least stable configuration shows segregated nitrogen and carbon islands within the plane. It follows that the least stable configuration overall has 100% nitrogen content.

The formation energy [see Eq. (2)] is a function of the chemical potentials for carbon and nitrogen, which depend on the experimental setup. While relative energies at a given composition are independent of the chemical potentials, the energy ranking of structures with different compositions may change. To gauge this effect, we present the potential phase diagrams based on the chemical potential combinations: pyridine-graphite, ammonia-graphite, ammonia-methane, in Sec. S5.1 of the SM [35]. Under these changes to the chemical potentials, the phase diagrams remain qualitatively similar, while the spread in formation energies and the slope of the lower convex hull change.

Finally, the gaps observed throughout the phase diagram in Fig. 3 arise from the fact that we limit the configurations we consider to a maximum supercell size. Filling these gaps would require supercells with more than 100 graphene lattice sites, which would lead to a greater combinatorial explosion, becoming computationally intractable. Nevertheless, SAMPLE has generated a comprehensive set of configurations for the given constraints, which are used to evaluate the energy phase diagram of NDG deterministically.

C. Thermodynamics

Equipped with a set of configurations and their corresponding formation energies, we can apply statistical mechanics to predict the thermodynamic properties of the system. Treating the entire configuration space as a grand canonical ensemble,

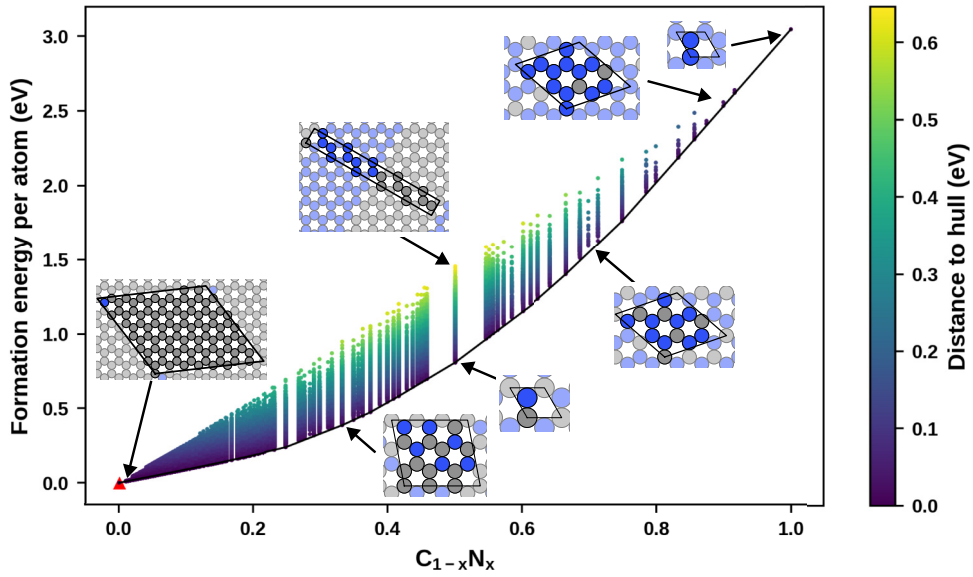


FIG. 3. The compositional phase diagram of graphene doped with graphitic nitrogen sites. Pristine graphene is represented as a red triangle with a formation energy of 0.00 eV. Also shown is a selection of structures to demonstrate the variety of supercell shapes and sizes considered by SAMPLE.

we can construct the partition function explicitly by considering each configuration and its respective formation energy as a microstate of the ensemble [see Eq. (7)]. To ensure valid sampling, our analysis uses only supercells with a fixed size of 100 graphene lattice sites. At this cell size, the thermodynamic properties are well converged, as shown in Fig. S11 [35].

Figure 4 shows the calculated thermodynamic quantities for NDG, illustrating how fluctuations in the concentration θ cause more and more states to become accessible with increasing temperature. This leads to a sharp increase in the internal energy and entropy. The temperature of this increase strongly depends on the choice of chemical potential. Irrespective of the chemical potential, the heat capacity exhibits an anomalous peak, which is typically indicative of a transition within the system. The position of this peak shifts with the chemical potential: For the pyridine-graphite chemical potentials, it lies at approximately 700 K, for ammonia-methane at 4700 K, for dinitrogen-graphite at 5300 K, and for ammonia-graphite the peak is above 7500 K. Temperatures of additional characteristic features in Fig. 4 are shown in Table S9 [35]. We note that temperatures of 4000 K and more are above the sublimation point of graphite, where the model of NDG is no longer valid.

By taking the first derivative, we further show that this peak is continuous, with a clear minimum inflection point, indicating that the transition is complete at that temperature. Beyond this inflection point temperature, the derivative of the heat capacity becomes asymptotic to zero, indicating that the heat capacity itself has also become asymptotic, but to a nonzero value as a consequence of the transition.

Owing to the finite size of the systems we have investigated, we cannot achieve a defect concentration below 1%. In fact, there exist an infinite number of possible configurations, given an arbitrary supercell size. Sampling the configuration space uniformly with a maximum supercell size is expected to yield a representative, finite subset of NDG configurations. However, the sheer size of the configuration space forces us to focus on low-concentration configurations. Due to the mostly repulsive defect interactions, low-concentration configurations dominate the thermodynamic properties in the present temperature range of 0 to 7500 K, as shown in Table S7 [35]. We expect that the addition of high-concentration configurations does not significantly affect the thermodynamic properties at lower temperatures: a more complete picture of the phase space would essentially push the critical point to a slightly higher temperature, while broadening the peak. A thorough discussion about sampling quality and the convergence of the thermodynamic quantities can be found in Secs. S4.2 and S5.1 in the SM [35].

IV. DISCUSSION

Defect-defect interactions. Investigating the pair interaction energies with DFT yields similar results to those shown by previous studies [25,26,65]. Overall, nitrogen-nitrogen pair interactions are repulsive. This trend is likely a result of repulsive Coulombic electrostatic interaction, as previously argued [26]. However, we find a minimum at the H3 position, as reported in Sec. III A, corresponding to a configuration that does not appear to have been discussed in previous works. Nevertheless, the trends and anisotropies demonstrated by the defects' interaction profile are consistent with the literature; our results show that these profiles are more than just a function of distance and sublattice site, but very much dependent on the nature of the path taken between the defect sites. For example, the presence of a "ring crossing" between the two defects contributes to the stability of the configuration, while increasing separation via a zigzag pathway just leads to a monotonic decay of the repulsion. Bu *et al*. have previously argued that this is a result of the anisotropic distribution of the electron charges [26].

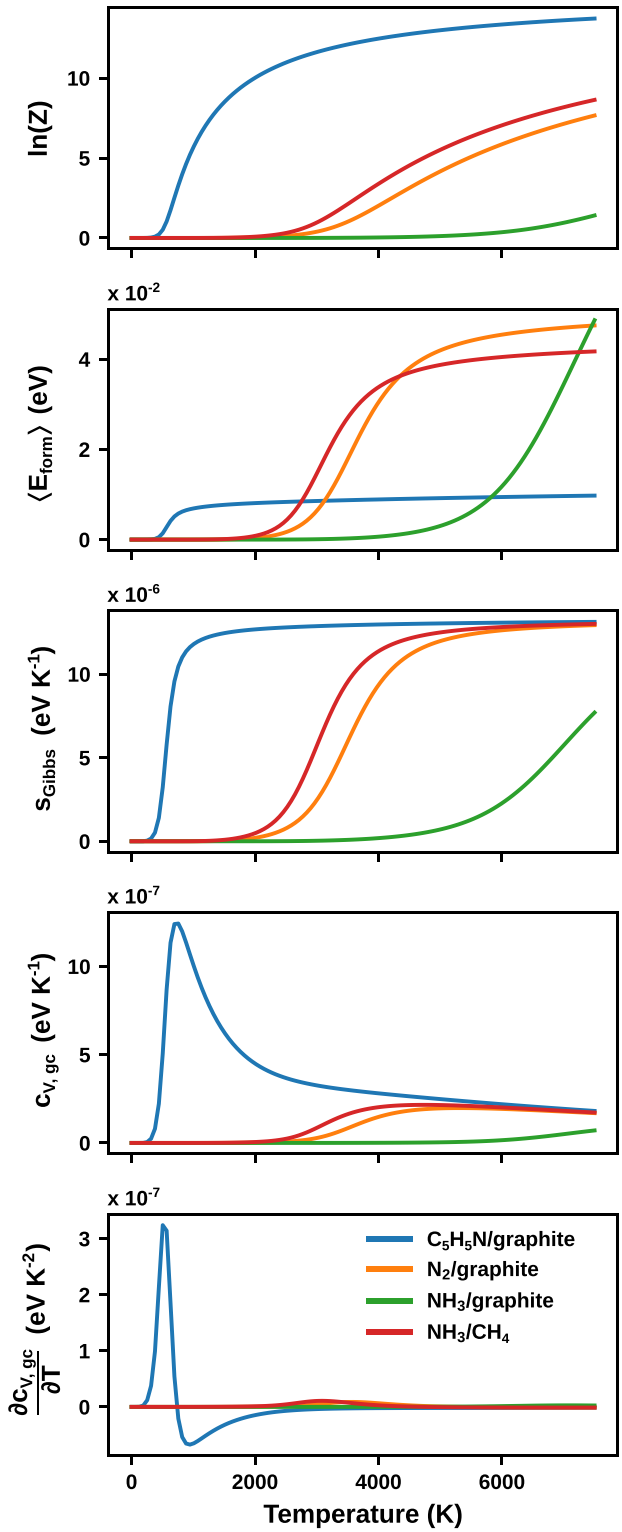


FIG. 4. Predicted thermodynamic properties as a function of temperature, using all configurations with 100 atoms as a grand canonical ensemble. From top to bottom: the logarithm of the partition function, expected energy, entropy, heat capacity, and first derivative of the heat capacity calculated numerically.

Compositional phase diagram. We find that the compositional phase diagram is qualitatively robust against changes in chemical potential. Using the common literature references (dinitrogen and graphite), the topology of the phase diagram for graphitic-N defects in graphene matches that observed in previous works by Shi *et al.* [24] and Bu *et al.* [26], the difference in this study being that we can guarantee that the configurations that lie on the convex hull are the most stable for a given dopant concentration. By investigating a configuration space with much larger graphene supercells than previous works, we also demonstrate that regions of the phase diagram, which were previously shown to exhibit local concavities, are an artifact of the maximum size of system employed.

The formation energy exhibits an overall exponential increase with increasing defect concentration, consistent with trends reported in the literature. In particular, the sharp change in the convex hull gradient at $\theta = 0.25$ corresponds to the stoichiometry C_3N and agrees with previous studies. No other comparably pronounced features are observed along the remainder of the convex hull. This is consistent with the hypothesis of Shi *et al.* [24], who proposed a "stability switchover" between $\theta = 0.33$ and 0.375, confirmed using phonon calculations in their study. However, the apparent range of this switchover may be an artifact of the small system sizes considered, while the sharp rise in formation energy at $\theta = 0.25$ could mark the true onset of this effect.

Thermodynamic properties and phase transitions. We observe a transition in the heat capacity, evidenced by a peak. The position of this peak is strongly dependent on the choice of chemical potential. Using dinitrogen-graphite as chemical potentials, as commonly done in literature, we find this peak at a temperature of approximately 5300 K, where basic assumptions of our model break down. This high temperature is the result of the strong binding energy of nitrogen molecules, which makes the incorporation of nitrogen into graphene energetically unfavorable. However, in CVD experiments, nitrogen is commonly supplied to the system in the form of nitrogen-containing molecules that give rise to a significantly higher (less negative) effective chemical potential of nitrogen. Using pyridine-graphite as a source, for instance, shifts the peak to approximately 700 K, which is in the range of common deposition temperatures.

This peak can be interpreted as a Schottky anomaly: At low temperatures, only a limited number of states are thermodynamically accessible. As the temperature surpasses a critical point, a significantly larger number of states become accessible, resulting in a peak in the heat capacity. As Fig. 4 demonstrates, there is a sudden increase in the system's entropy, which corresponds to the rise in the number of accessible configurations. Among the 105 million configurations explored, the compositional phase diagram of NDG exhibits two distinct regimes as dopant concentration increases: clustered, islandlike defect superstructures and uniformly dispersed defects across the 2D plane. The most prominent example is illustrated in Fig. 3 at 50% concentration; the configuration with the largest formation energy is composed of large strips of hexagonal 2D nitrogen and graphene, and the lowest-energy configuration has C and N atoms at the *A* and *B* sublattice sites of graphene. It is impossible for a nitrogen-nitrogen neighbor defect pair not to be present in any given configuration with a defect concentration greater than 50%. Previous work [66] has shown

that the presence of these nitrogen-nitrogen neighbor pair motifs is detrimental to the stability of the doped graphene structure. Naturally, clustered configurations are composed of large numbers of nitrogen-nitrogen neighbor pairs, which greatly increases their formation energy. Due to symmetry, low-energy configurations with maximally spaced nitrogen atoms are few in number compared to configurations containing clusters of nitrogen. Before the critical temperature is reached, only uniformly dispersed defect superstructures contribute to the partition function. Above the critical temperature, clusters of defect superstructures become available.

Similar peaks in heat capacities have been reported for various systems, such as order-disorder phase transitions in molecular solids [67,68], charge disorder in bulk materials [69], binary clusters described with a Lennard-Jones potential [28], and for the Ising model [70]. Di Cairano *et al.* [71] interpreted this peak as a third-order phase transition. Using the sine-Gordon model in a microcanonical ensemble on a square two-dimensional lattice, the authors found a "roughness"-driven phase transition where the heat capacity displays a peak but remains a continuous function. In the sine-Gordon model, each site can assume arbitrary field values, qualitatively resembling our system, where nitrogen and carbon represent different field values when only one unit cell is considered. Both models feature one- and two-body interactions. However, the sine-Gordon model includes only next-nearest-neighbor interactions, while in the nitrogen-defect model, the most repulsive interactions occur at the closest range and decay rapidly (see Fig. 2).

To study the effect of the energy distribution in the microstates, and how it may affect the energy fluctuations, and hence the heat capacity, we employ a toy model that allows us to arbitrarily generate energy spectra (see Sec. S5.3 of the SM [35]). We find that systems with degenerate ground states and final states result in a broader peak, reaching into very high temperatures. Conversely, for energy spectra with a single ground state, final state, but degenerate intermediate states, the heat capacity is much narrower. Such energy spectra are a result of the combinatorics governing the energetic favorability of structures: Only very few configurations of a system's degrees of freedom will lead to very low or high energies, while a large number of configurations have moderate energies. This was observed in several studies on surface structure prediction [29,72–74]. This transition represents a combinatorially driven order-disorder phenomenon, where rising temperature enhances the accessibility of clustered nitrogen-nitrogen defect interactions. Crucially, it is not the energetic contributions of nitrogen pairs alone that drive the transition, since such motifs exist even in dispersed configurations, but the point at which clustered configurations start to dominate the partition function due to their abundance compared to dispersed configurations.

V. CONCLUSION

We have presented a modification of the SAMPLE approach capable of generating defective two-dimensional structures. Using this approach, we have generated more than 1×10^8 configurations of NDG, enabling us to explore the compositional phase diagram and the thermodynamic properties. In the phase diagram, we observe a sharp change in the convex hull gradient at $\theta = 0.25$, corresponding to the stoichiometry C_3N, which is in agreement with literature. The internal energy and entropy exhibit a sigmoidal behavior when plotted against the temperature. The heat capacity shows a characteristic peak, which is the result of a combinatorially driven order-disorder transition. The positions of the peaks strongly depend on the chemical potentials of nitrogen and carbon, which affect the thermodynamic favorability of incorporating nitrogen into graphene. Future work could focus on extending this approach to other defective 2D materials, such as the full boron-nitrogen-carbon phase diagram and defects such as the Stone-Wales defect, to explore their configurational landscapes and phase behavior under varying thermodynamic conditions. More efficient data processing will enable the exploration of even larger configuration spaces.

ACKNOWLEDGMENTS

The authors acknowledge financial support from the UKRI Future Leaders Fellowship program (Grant No. MR/X023109/1), an AIchemy Pump Prime grant (Grant No. 122408), funded through the EPSRC (Grant No. EP/Y028759/1), and UKRI Horizon grants (ERC StG, Grant No. EP/X014088/1 and MSCA-PF, Grant No. EP/Y024923/1). We gratefully acknowledge O. T. Hofmann and C. Wachter for support with the SAMPLE code and L. Bartók-Pártay and T. Chatbipho for many fruitful discussions. High-performance computing resources were provided via the EuroHPC supercomputer LUMI, hosted by CSC (Finland) and the LUMI consortium through a EuroHPC Regular Access call (Grant No. EHPC-REG-2023R03-123). Additional computation resources were provided by the Scientific Computing Research Technology Platform of the University of Warwick, the EPSRC-funded Materials Chemistry Consortium (Grants No. EP/R029431/1 and No. EP/X035859/1), and the UK Car-Parrinello consortium (Grant No. EP/X035891/1) for the ARCHER2 UK National Supercomputing Service, the EPSRC-funded HPC Midlands+ computing centre for access to Sulis (Grant No. EP/P020232/1).

DATA AVAILABILITY

The data supporting this research are publicly available on the NOMAD repository [75]. The SAMPLE code is publicly available [76].

[1] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, and A. A. Firsov, Electric field effect in atomically thin carbon films, Science **306**, 666 (2004).

[2] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, M. I. Katsnelson, I. V. Grigorieva, S. V. Dubonos, and A. A. Firsov, Two-dimensional gas of massless Dirac fermions in graphene, Nature (London) **438**, 197 (2005).

[3] A. H. Castro Neto, F. Guinea, N. M. R. Peres, K. S. Novoselov, and A. K. Geim, The electronic properties of graphene, Rev. Mod. Phys. **81**, 109 (2009).
[4] A. A. Balandin, Thermal properties of graphene and nanostructured carbon materials, Nat. Mater. **10**, 569 (2011).
[5] G. Tsoukleri, J. Parthenios, K. Papagelis, R. Jalil, A. C. Ferrari, A. K. Geim, K. S. Novoselov, and C. Galiotis, Subjecting a graphene monolayer to tension and compression, Small **5**, 2397 (2009).
[6] V. Mazánek, J. Luxa, S. Matějková, J. Kučera, D. Sedmidubský, M. Pumera, and Z. Sofer, Ultrapure graphene is a poor electrocatalyst: Definitive proof of the key role of metallic impurities in graphene-based electrocatalysis, ACS Nano **13**, 1574 (2019).
[7] L. Wang, Z. Sofer, and M. Pumera, Will any crap we put into graphene increase its electrocatalytic effect? ACS Nano **14**, 21 (2020).
[8] P. Han, X. Yu, D. Yuan, M. Kuang, Y. Wang, A. M. Al-Enizi, and G. Zheng, Defective graphene for electrocatalytic CO2 reduction, J. Colloid Interface Sci. **534**, 332 (2019).
[9] D. J. Rizzo, G. Veber, T. Cao, C. Bronner, T. Chen, F. Zhao, H. Rodriguez, S. G. Louie, M. F. Crommie, and F. R. Fischer, Topological band engineering of graphene nanoribbons, Nature (London) **560**, 204 (2018).
[10] M. Yankowitz, J. Jung, E. Laksono, N. Leconte, B. L. Chittari, K. Watanabe, T. Taniguchi, S. Adam, D. Graf, and C. R. Dean, Dynamic band-structure tuning of graphene moiré superlattices with pressure, Nature (London) **557**, 404 (2018).
[11] P. Rani and V. Jindal, Designing band gap of graphene by B and N dopant atoms, RSC Adv. **3**, 802 (2013).
[12] Z. Zeng, F.-X. Xiao, X. Gui, R. Wang, B. Liu, and T. T. Yang Tan, Layer-by-layer assembly of nitrogen-doped graphene quantum dots monolayer decorated one-dimensional semiconductor nanoarchitectures for solar-driven water splitting, J. Mater. Chem. A **4**, 16383 (2016).
[13] C. Bie, H. Yu, B. Cheng, W. Ho, J. Fan, and J. Yu, Design, fabrication, and mechanism of nitrogen-doped graphene-based photocatalyst, Adv. Mater. **33**, 2003521 (2021).
[14] H. Chen, Z. Mo, Z. Wang, P. Yan, P. Sun, G. Wu, J. Zhang, X. Zhu, L. Wang, and H. Xu, Implanting nitrogen-doped graphene quantum dots on porous ultrathin carbon nitride for efficient metal-free photocatalytic hydrogen evolution, J. Environ. Chem. Eng. **11**, 109801 (2023).
[15] A. Mirzaei, S. P. Bharath, J.-Y. Kim, K. K. Pawar, H. W. Kim, and S. S. Kim, N-doped graphene and its derivatives as resistive gas sensors: An overview, Chemosensors **11**, 334 (2023).
[16] U. Nath and M. Sarma, Pyridinic dominance N-doped graphene: A potential material for SO_2 gas detection, J. Phys. Chem. A **127**, 1112 (2023).
[17] T. He, A. R. Puente Santiago, and A. Du, Atomically embedded asymmetrical dual-metal dimers on N-doped graphene for ultra-efficient nitrogen reduction reaction, J. Catal. **388**, 77 (2020).
[18] A. Baby, L. Trovato, and C. Di Valentin, Single atom catalysts (SAC) trapped in defective and nitrogen-doped graphene supported on metal substrates, Carbon **174**, 772 (2021).
[19] Y. Yang, J. Li, C. Zhang, Z. Yang, P. Sun, S. Liu, and Q. Cao, Theoretical insights into nitrogen-doped graphene-supported Fe, Co, and Ni as single-atom catalysts for CO_2 reduction reaction, J. Phys. Chem. C **126**, 4338 (2022).
[20] L. A. Alvarado-Leal, J. I. Paez-Ornelas, H. N. Fernández-Escamilla, E. G. Perez-Tijerina, J. Guerrero-Sánchez, J. M. Romo-Herrera, and N. Takeuchi, Influence of the atomic local environment on the ORR activity of single-atom catalysts in N-doped graphene, ACS Appl. Energy Mater. **7**, 4794 (2024).
[21] D. Usachov, O. Vilkov, A. Grüneis, D. Haberer, A. Fedorov, V. K. Adamchuk, A. B. Preobrajenski, P. Dudin, A. Barinov, M. Oehzelt, C. Laubschat, and D. V. Vyalikh, Nitrogen-doped graphene: Efficient growth, structure, and electronic properties, Nano Lett. **11**, 5401 (2011).
[22] C. Zhang, L. Fu, N. Liu, M. Liu, Y. Wang, and Z. Liu, Synthesis of nitrogen-doped graphene using embedded carbon and nitrogen sources, Adv. Mater. **23**, 1020 (2011).
[23] W. Zhao, O. Höfert, K. Gotterbarm, J. F. Zhu, C. Papp, and H.-P. Steinrück, Production of nitrogen-doped graphene by low-energy nitrogen implantation, J. Phys. Chem. C **116**, 5062 (2012).
[24] Z. Shi, A. Kutana, and B. I. Yakobson, How much N-doping can graphene sustain? J. Phys. Chem. Lett. **6**, 106 (2015).
[25] H. J. Xiang, B. Huang, Z. Y. Li, S.-H. Wei, J. L. Yang, and X. G. Gong, Ordered semiconducting nitrogen-graphene alloys, Phys. Rev. X **2**, 011003 (2012).
[26] S. Bu, N. Yao, M. A. Hunter, D. J. Searles, and Q. Yuan, Design of two-dimensional carbon-nitride structures by tuning the nitrogen concentration, npj Comput. Mater. **6**, 128 (2020).
[27] L. B. Pártay, A. P. Bartók, and G. Csányi, Nested sampling for materials: The case of hard spheres, Phys. Rev. E **89**, 022302 (2014).
[28] R. Yang, J. Chen, D. Thibodeaux, and R. B. Wexler, Freebird.jl: An extensible toolbox for simulating interfacial phase equilibria, J. Chem. Theory Comput. **21**, 10765 (2025).
[29] L. Hörmann, A. Jeindl, A. T. Egger, M. Scherbela, and O. T. Hofmann, SAMPLE: Surface structure search enabled by coarse graining and statistical learning, Comput. Phys. Commun. **244**, 143 (2019).
[30] Y. Fujimoto and S. Saito, Formation, stabilities, and electronic properties of nitrogen defects in graphene, Phys. Rev. B **84**, 245446 (2011).
[31] A. Kumar, A. Ganguly, and P. Papakonstantinou, Thermal stability study of nitrogen functionalities in a graphene network, J. Phys.: Condens. Matter **24**, 235503 (2012).
[32] E. Doustkhah, A. Kotb, S. Tafazoli, T. Balkan, S. Kaya, D. A. Hanaor, and M. H. N. Assadi, Templated synthesis of exfoliated porous carbon with dominant graphitic nitrogen, ACS Mater. Au **3**, 231 (2023).
[33] R. Lv, Q. Li, A. R. Botello-Méndez, T. Hayashi, B. Wang, A. Berkdemir, Q. Hao, A. L. Elías, R. Cruz-Silva, H. R. Gutiérrez, Y. A. Kim, H. Muramatsu, J. Zhu, M. Endo, H. Terrones, J.-C. Charlier, M. Pan, and M. Terrones, Nitrogen-doped graphene: Beyond single substitution and enhanced molecular sensing, Sci. Rep. **2**, 586 (2012).
[34] R. Trusovas, K. Ratautas, G. Račiukaitis, J. Barkauskas, I. Stankevičienė, G. Niaura, and R. Mažeikienė, Reduction of graphite oxide to graphene with laser irradiation, Carbon **52**, 574 (2013).
[35] See Supplemental Material at http://link.aps.org/supplemental/10.1103/db9g-7gvs for a detailed description of the computational methods, convergence tests, formation energies of defects, extended derivations, and supplementary figures.
[36] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, *Ab initio* molecular simulations with

numeric atom-centered orbitals, Comput. Phys. Commun. **180**, 2175 (2009).
[37] V. Havu, V. Blum, P. Havu, and M. Scheffler, Efficient $O(N)$ integration for all-electron electronic structure calculation using numeric basis functions, J. Comput. Phys. **228**, 8367 (2009).
[38] J. P. Perdew, K. Burke, and M. Ernzerhof, Generalized gradient approximation made simple, Phys. Rev. Lett. **77**, 3865 (1996).
[39] A. Tkatchenko and M. Scheffler, Accurate molecular van der Waals interactions from ground-state electron density and free-atom reference data, Phys. Rev. Lett. **102**, 073005 (2009).
[40] H. J. Monkhorst and J. D. Pack, Special points for Brillouin-zone integrations, Phys. Rev. B **13**, 5188 (1976).
[41] J. Moreno and J. M. Soler, Optimal meshes for integrals in real- and reciprocal-space unit cells, Phys. Rev. B **45**, 13891 (1992).
[42] P. Wisesa, K. A. McGill, and T. Mueller, Efficient generation of generalized Monkhorst-Pack grids through the use of informatics, Phys. Rev. B **93**, 155109 (2016).
[43] J. C. Meyer, A. K. Geim, M. I. Katsnelson, K. S. Novoselov, T. J. Booth, and S. Roth, The structure of suspended graphene sheets, Nature (London) **446**, 60 (2007).
[44] K. Hermann, *Crystallography and Surface Structure: An Introduction for Surface Scientists and Nanoscientists* (Wiley, Hoboken, NJ, 2017).
[45] B. Jeong, H. G. Abbas, B. P. Klein, G. Bae, A. R. Velmurugan, C. H. Choi, G. Kim, D. Kim, K.-J. Kim, B. J. Cha, Y. D. Kim, F. Jaouen, R. J. Maurer, and S. Ringe, CO cryo-sorption as a surface-sensitive spectroscopic probe of the active site density of single-atom catalysts, Angew. Chem. Int. Ed. **64**, e202420673 (2025).
[46] S. Ullah, A. Hussain, and F. Sato, Rectangular and hexagonal doping of graphene with B, N, and O: A DFT study, RSC Adv. **7**, 16064 (2017).
[47] E. B. Yutomo, F. A. Noor, and T. Winata, Effect of the number of nitrogen dopants on the electronic and magnetic properties of graphitic and pyridinic N-doped graphene–a density-functional study, RSC Adv. **11**, 18371 (2021).
[48] N. Dimov, A. Staykov, M. I. M. Kusdhany, and S. M. Lyth, Tailoring the work function of graphene via defects, nitrogen-doping and hydrogenation: A first principles study, Nanotechnology **34**, 415001 (2023).
[49] Z. Khan, S. O. Park, J. Yang, S. Park, R. Shanker, H.-K. Song, Y. Kim, S. K. Kwak, and H. Ko, Binary *N*, *S*-doped carbon nanospheres from bio-inspired artificial melanosomes: A route to efficient air electrodes for seawater batteries, J. Mater. Chem. A **6**, 24459 (2018).
[50] M. Yang, L. Wang, M. Li, T. Hou, and Y. Li, Structural stability and O2 dissociation on nitrogen-doped graphene with transition metal atoms embedded: A first-principles study, AIP Adv. **5**, 067136 (2015).
[51] T. T. Pham, T. N. Pham, V. Chihaia, Q. A. Vu, T. T. Trinh, T. T. Pham, L. Van Thang, and D. N. Son, How do the doping concentrations of N and B in graphene modify the water adsorption? RSC Adv. **11**, 19560 (2021).
[52] H. Wang, T. Maiyalagan, and X. Wang, Review on recent progress in nitrogen-doped graphene: Synthesis, characterization, and its potential applications, ACS Catal. **2**, 781 (2012).
[53] L. Pauling, The nature of the chemical bond. IV. The energy of single bonds and the relative electronegativity of atoms, J. Am. Chem. Soc. **54**, 3570 (1932).
[54] F. Pukelsheim, *Optimal Design of Experiments* (SIAM, Philadelphia, 2006).
[55] B. Saunders, Predicting the structure and stability of compositionally diverse materials, Ph.D. thesis, University of Warwick, 2025.
[56] I. Choudhuri, N. Patra, A. Mahata, R. Ahuja, and B. Pathak, B–N@graphene: Highly sensitive and selective gas sensor, J. Phys. Chem. C **119**, 24827 (2015).
[57] O. Olaniyan, R. Maphasha, M. Madito, A. Khaleed, E. Igumbor, and N. Manyala, A systematic study of the stability, electronic and optical properties of beryllium and nitrogen co-doped graphene, Carbon **129**, 207 (2018).
[58] S. Kattel, P. Atanassov, and B. Kiefer, Stability, electronic and magnetic properties of in-plane defects in graphene: A first-principles study, J. Phys. Chem. C **116**, 8161 (2012).
[59] N. E. Benti, G. A. Tiruye, and Y. S. Mekonnen, Boron and pyridinic nitrogen-doped graphene as potential catalysts for rechargeable non-aqueous sodium–air batteries, RSC Adv. **10**, 21387 (2020).
[60] J. Lee, W.-C. Tian, W.-L. Wang, and D.-X. Yao, Two-dimensional pnictogen honeycomb lattice: Structure, on-site spin-orbit coupling and spin polarization, Sci. Rep. **5**, 11512 (2015).
[61] F. Nogueira, Bayesian optimization: Open source constrained global optimization tool for Python (2014), https://github.com/bayesian-optimization/BayesianOptimization.
[62] A. P. Bartók, R. Kondor, and G. Csányi, On representing chemical environments, Phys. Rev. B **87**, 184115 (2013).
[63] S. De, A. P. Bartók, G. Csányi, and M. Ceriotti, Comparing molecules and solids across structural and alchemical space, Phys. Chem. Chem. Phys. **18**, 13754 (2016).
[64] K. K. Huguenin-Dumittan, P. Loche, N. Haoran, and M. Ceriotti, Physics-inspired equivariant descriptors of nonbonded interactions, J. Phys. Chem. Lett. **14**, 9612 (2023).
[65] Z. Hou, X. Wang, T. Ikeda, K. Terakura, M. Oshima, M.-A. Kakimoto, and S. Miyata, Interplay between nitrogen dopants and native point defects in graphene, Phys. Rev. B **85**, 165439 (2012).
[66] V. V. Chaban and O. V. Prezhdo, Nitrogen–nitrogen bonds undermine stability of N-doped graphene, J. Am. Chem. Soc. **137**, 11688 (2015).
[67] Y. Miyazaki, M. Nakano, A. I. Krivchikov, O. A. Koroyuk, J. F. Gebbia, C. Cazorla, and J. L. Tamarit, Low-temperature heat capacity anomalies in ordered and disordered phases of normal and deuterated thiophene, J. Phys. Chem. Lett. **12**, 2112 (2021).
[68] D. Szewczyk, J. F. Gebbia, A. Jeżowski, A. I. Krivchikov, T. Guidi, C. Cazorla, and J.-L. Tamarit, Heat capacity anomalies of the molecular crystal 1-fluoro-adamantane at low temperatures, Sci. Rep. **11**, 18640 (2021).
[69] B. C. Melot, R. Tackett, J. O'Brien, A. L. Hector, G. Lawes, R. Seshadri, and A. P. Ramirez, Large low-temperature specific heat in pyrochlore $Bi_2Ti_2O_7$, Phys. Rev. B **79**, 224111 (2009).
[70] L. Regeciová and K. J. Kapcia, Localized fermions on the triangular lattice with Ising-like interactions, Phys. Rev. E **111**, 064132 (2025).
[71] L. Di Cairano and A. Tkatchenko, The specific heat anomaly stems from a third-order phase transition in the 2D lattice sine-Gordon model, arXiv:2504.20766.

[72] A. Jeindl, J. Domke, L. Hörmann, F. Sojka, R. Forker, T. Fritz, and O. T. Hofmann, Nonintuitive surface self-assembly of functionalized molecules on Ag(111), ACS Nano **15**, 6723 (2021).

[73] M. Scherbela, L. Hörmann, A. Jeindl, V. Obersteiner, and O. T. Hofmann, Charting the energy landscape of metal/organic interfaces via machine learning, Phys. Rev. Mater. **2**, 043803 (2018).

[74] L. Hörmann, A. Jeindl, and O. T. Hofmann, From a bistable adsorbate to a switchable interface: Tetrachloropyrazine on Pt(111), Nanoscale **14**, 5154 (2022).

[75] B. Saunders, L. Hörmann, and R. J. Maurer, Data for: Comprehensive structure exploration and thermodynamics of graphitic heteroatom doped graphene superstructures, doi: 10.17172/NOMAD/2025.09.13-1.

[76] https://gitlab.tugraz.at/othgroup/sample.