Beating the egg-box effect in plane-wave DFT simulations

IOP Publishing Electron. Struct. 7 (2025) 025004 https://doi.org/10.1088/2516-1075/adc056

Electronic Structure

OPEN ACCESS

RECEIVED
5 January 2025

REVISED
20 February 2025

ACCEPTED FOR PUBLICATION
13 March 2025

PUBLISHED
7 May 2025

PAPER

Beating the egg-box effect in plane-wave DFT simulations

Ben Durham, Matt I J Probert and Phil J Hasnip*

School of Physics, Engineering and Technology, University of York, Heslington, York YO10 5DD, United Kingdom
* Author to whom any correspondence should be addressed.

E-mail: phil.hasnip@york.ac.uk

Keywords: egg-box effect, density functional theory, electronic structure, software, numerical methods

Abstract

The 'egg-box effect' is a known challenge in density functional theory (DFT) calculations which arises from the discretisation of continuous quantities, e.g. the electron density. This effect is observed when the system is moved relative to the underlying computational grid and causes an unphysical change in the system's total energy, violating translational invariance. The pattern of energy change with translation is reminiscent of an egg-box. This effect can cause unphysical results such as geometry relaxations finding incorrect crystal symmetries or imaginary phonon modes in vibrational calculations. The egg-box effect can be mitigated by using finer grids for all continuous quantities, but this greatly increases computational cost. For plane-wave DFT the effect's origin is the evaluation of the exchange and correlation (XC) energy (E_{XC}). We present a novel technique for estimating the violation of translational invariance in E_{XC}, using a Fourier interpolation scheme, providing an estimate of the uncertainty. Our results also show that the numerical behaviour of the XC approximations is strongly linked to the magnitude of the violation of the translation invariance. The more numerically ill behaved XC functionals, such as the more advanced meta-generalised-gradient approximations functionals, exhibit changes in E_{XC} (due to violation of translation invariance) that are orders of magnitude worse than less advanced functionals. Performing this analysis at an early stage of a workflow can inform the user about the expected accuracy of subsequent calculations. Further, our results demonstrate that by selectively computing E_{XC} and its corresponding potential (V_{XC}) on a finer grid, the egg-box effect can be significantly reduced. Coupled with our uncertainty quantification method, egg-box related inaccuracies can be avoided more conveniently and efficiently than just increasing grid resolution until inaccuracies appear suppressed. This work offers a promising pathway towards mitigating the egg-box effect in a diverse range of materials modelling applications.

1. Introduction

Significant progress has been made over recent years in improving the accuracy of density functional approximations (DFAs) for electronic structure simulations. The strongly constrained and appropriately normed (SCAN) functional was the first meta-generalised-gradient approximation (meta-GGA) to adhere to all the known constraints for a meta-GGA functional [1]. This was a major achievement for chemical accuracy; however, it was quickly observed that SCAN, as with many other meta-GGA functionals, suffered from severe numerical instability [2–8]. This led to the creation of several related functionals [2, 9–11], which attempted to regularise the numerical behaviour of SCAN, while maintaining SCAN's physical accuracy. Considerable improvements have been achieved with respect to SCAN, but, as recent work has shown, many modern DFAs including the SCAN family are numerically ill behaved [8].

The root of this poor numerical behaviour is in the integration of the contributions to the exchange and correlation (XC) energy, E_{XC} ; see more details in section 2. The choice of basis set, the number of sampling points, and the position of those sampling points can greatly affect the calculated value of E_{XC}. This can lead to simulations which violate translational symmetry, with the energy of the system oscillating if the system is translated relative to the discrete sampling points. This is sometimes referred to as the egg-box effect, due to

the characteristic shape of the energy landscape (see figure 1). The numerical integration of the contributions to E_{XC} in real space will almost always cause some violation of the translational symmetry. However, the choice of basis set functions or representation of potentials can lead to additional causes of the violation of translational symmetry [12–14].

The numerical integration of the contributions to E_{XC} is highly dependent on the choice of numerical sampling and can have significant effects on the physical properties predicted by simulation packages. Sitkiewicz *et al* [15, 16] have made several studies of how this numerical ill-behaviour can lead to spurious oscillations in the energy as the atoms are displaced with respect to the grid, due to numerical integration errors. These spurious oscillations have been shown to cause particular difficulty in calculating vibrational spectra [17], particularly for molecules that have low energy vibrational modes, where these oscillations cause significant noise in the spectrum.

For users of electronic structure software, it can be difficult to find a balance between the desire for high physical accuracy and the need for good numerical behaviour. There is very little indication to a user that there may be numerical problems arising from the combination of their choice of system, XC functional and basis set parameters. The work of Sitkiewicz *et al* [15, 16] provides useful insight into the relative size of these errors in physical properties for different DFAs on the specific molecular data set tested. The lack of any efficient and transferable method for quantifying errors related to E_{XC} means that, in general, users must perform careful and expensive testing for their system and functional choice.

In this work, we present an approach to quantifying the error in E_{XC} due to the effect of the violation of translational symmetry in XC integration. We show that this method can be used to accurately estimate the change in E_{XC} as the discrete grid is translated relative to the system. Importantly, this method of error estimation is efficient to compute, relies on no other data than the ground state density, and is general to all XC functionals.

Further, we show that the integration of the XC energy contributions can be performed more accurately by computing only the XC contribution on a finer grid to correctly sample the energy and potential. By then truncating the XC potential (V_{XC}) in Fourier space, the method reduces the aliasing of high-frequency components into the potential. This can greatly reduce the computational expense compared to computing every contribution to the total energy on this finer grid, and hence allows properties to be calculated more accurately and more efficiently than would otherwise be possible.

This article is set out as follows; in section 2 we provide a brief description of the different DFAs and how the numerical integration error arises in the plane-wave basis set and the consequences this can have for computed properties. We then describe the uncertainty quantification (UQ) method in section 3 and finally in section 4 we demonstrate its use as an indication that the calculations are under-converged and likely to produce inaccurate results.

2. Theory

2.1. Plane-wave basis set

In plane-wave density functional theory (DFT) codes, such as CASTEP [18], periodic boundary conditions are used. Bloch's theorem states that the wavefunction can be written as the product of a cell-periodic part $u_b(\mathbf{r})$ and a wavelike part,

$$\psi_{b,\mathbf{k}}(\mathbf{r}) = u_b(\mathbf{r})\, e^{i\mathbf{k}\cdot\mathbf{r}} \tag{1}$$

where $\mathbf{k}$ becomes a free parameter that must be sampled over the 1st Brillouin zone. The cell-periodic part can be expressed as a discrete expansion of plane-waves. The plane-waves have wave-vectors with spatial frequencies corresponding to a reciprocal lattice vector $\mathbf{G}$. The expansion is restricted such that only plane-waves that have a kinetic energy less than some cut-off energy E_{cut} are considered,

$$\psi_{b,\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{G}}^{|\mathbf{k}+\mathbf{G}|\leqslant G_{cut}} c_{b,\mathbf{k},\mathbf{G}} e^{i(\mathbf{k}+\mathbf{G})\cdot\mathbf{r}} \tag{2}$$

where $\psi_{b,\mathbf{k}}$ is the wavefunction for a particular k-point $\mathbf{k}$ and state b, $G_{cut} = \sqrt{\frac{2m_e E_{cut}}{\hbar^2}}$, and $c_{b,\mathbf{k},\mathbf{G}}$ is the Fourier expansion coefficient for that particular plane-wave. The use of the plane-wave basis set means that fast Fourier transforms (FFTs) can be used to quickly move between reciprocal space and real space representations of the Kohn–Sham wavefunctions.

The ground state wavefunctions are found by minimising the total energy functional

$$E = \sum_{b,\mathbf{k}}^{occ} \langle \psi_{b,\mathbf{k}} | \hat{T} + \hat{V}_{\text{ext}} | \psi_{b,\mathbf{k}} \rangle + E_{\text{H}}[n] + E_{\text{XC}}[n] + E_{\text{Ion-Ion}}, \tag{3}$$

where the sum is over the occupied states, $\hat{T}$ is the kinetic energy operator, $\hat{V}_{\text{ext}}$ is the external potential, $E_{\text{Ion-Ion}}$ is the energy of the interaction of the ions in the system, E_{H} is the energy of the electrostatic interaction of the electrons and E_{XC} is the exchange and correlation energy functional of the electron density, n. The wavefunctions are constrained to be orthonormal, $\langle \psi_b \rangle \psi_{b'} = \delta_{bb'}$. The minimised total energy is found by iterative solving the Kohn–Sham equations [19]

$$\hat{H}|\psi_{b,\mathbf{k}}\rangle = \left(\hat{T} + \hat{V}_{\text{ext}} + \frac{\delta E_H}{\delta n} + \frac{\delta E_{\text{XC}}}{\delta n} \right) |\psi_{b,\mathbf{k}}\rangle = \epsilon_{b,\mathbf{k}} |\psi_{b,\mathbf{k}}\rangle, \tag{4}$$

where $\hat{H}$ is the Kohn–Sham Hamiltonian of the system and $\epsilon_{b\mathbf{k}}$ are the Kohn–Sham eigenvalues.

The wavefunctions are used to calculate the electron density:

$$n(\mathbf{r}) = \sum_{b,\mathbf{k}} f_{b,\mathbf{k}} |\psi_{b,\mathbf{k}}(\mathbf{r})|^2 \tag{5}$$

where $f_{b,\mathbf{k}}$ is the occupancy of that state b at k-point $\mathbf{k}$. If the representations of the wavefunctions are restricted to reciprocal-lattice vectors with magnitude $|\mathbf{G}| \leqslant G_{\text{cut}}$ then the Fourier components of the charge density are restricted to $|\mathbf{G}| \leqslant 2G_{\text{cut}}$ such that:

$$n(\mathbf{r}) = \sum_{\mathbf{G}}^{|\mathbf{G}| \leqslant 2G_{\text{cut}}} a_{\mathbf{G}} e^{i\mathbf{G}\cdot\mathbf{r}} \tag{6}$$

where $a_{\mathbf{G}}$ is the coefficient in the Fourier expansion of the charge density. This means that for a particular cut-off for the wavefunction states, G_{cut}, there is a minimum FFT grid size required to exactly represent the resultant density. The minimum FFT grid required to represent all Fourier components of the density is a grid that is twice as dense as that of the plane-wave basis used for the wavefunctions, which is the default choice of CASTEP. CASTEP uses the `FINE_GRID_SCALE` parameter to control the scale of the grid used to represent the density compared to the wavefunction basis.

The FFT grid forms a regular grid for the real space sampling points. For most terms in the Hamiltonian, the choice of origin is arbitrary and the use of the default FFT grid, i.e. `FINE_GRID_SCALE = 2.0`, is sufficient to calculate the energy without integration error or aliasing. The major exception is E_{XC}.

2.2. XC functionals

The exact XC functional is unknown, and so we must use an approximation. There are various approaches to approximating E_{XC}, often classified as different rungs on a 'Jacob's ladder' [20] of accuracy, based on the required information for each functional. In general, it is observed that each step up the ladder results in a step towards increased physical accuracy at the cost of increased calculation time. The simplest approximations (lowest rung of the ladder) are the local density approximations (LDAs) [21, 22] which describe the contribution to E_{XC} at a specific point $\mathbf{r}$ using only the density, $n(\mathbf{r})$, at that specific point. LDAs are therefore purely local functionals.

The next class of approximations (rung up the ladder) is the generalised-gradient approximations (GGAs) [23], which include information about the surrounding electron density via the gradient of the charge density $|\nabla n|$ at each point. GGAs are often referred to as semi-local functionals.

Next are the so-called 'meta-GGAs' (mGGAs), which also depend on the Laplacian of the charge density, $\nabla^2 n$, and/or the orbital kinetic energy density (KED), τ. They are still semi-local functionals, but the inclusion of the KED allows further non-local information to be included, particularly giving an indication of the electron localisation.

The expression for the XC energy of the system calculated from a local/semi-local functional is:

$$E_{\text{XC}} = F\left[n, |\nabla n|, \nabla^2 n, \tau, \ldots\right] = \int_{\Omega} n(\mathbf{r})\, \epsilon_{\text{XC}}\left(n(\mathbf{r}), |\nabla n(\mathbf{r})|, \nabla^2 n(\mathbf{r}), \tau(\mathbf{r}), \ldots\right) \mathrm{d}\mathbf{r} \tag{7}$$

where ϵ_{XC} is the XC energy density and Ω is the volume of the system. Equation (7) shows E_{XC} evaluated as a continuous integral of the density and other ingredients. However, in practice, the evaluation of E_{XC} is carried out as a sum over a discrete set of N_{points}. The value of the density (and other ingredients to the XC

energy density) varies between the sampled points according to the basis functions, and so the contribution to the XC energy will also vary between the sampled points. The errors are generally assumed to be small enough to be negligible; however, as shown by Lehtola and Marques [8], this depends on the numerical behaviour of the XC functional. For many published XC functionals, the errors can be far from small.

2.3. 'Direct' vs. 'indirect' integration errors

A change in sampling affects the calculated energy in two ways. The first is as described above, a change in the sampling of the density leads to a change in the integration and hence has a direct effect on E_{XC}. The second effect comes via the XC potential, V_{XC}, through the minimisation of the total energy. The change in sampling leads to different values for V_{XC}, so that as the density and wavefunction states are improved to minimise the energy, a different ground-state density will be found, leading to different total energies. Here, the change in sampling indirectly affects the total energy through the other energy contributions in the Hamiltonian.

The calculated $V_{\mathrm{XC}}(\mathbf{r})$ is the derivative of E_{XC} w.r.t. a change in the density at that specific point $\mathbf{r}$. In order to calculate the sampling-error-free ground state, what is really required is the derivative of the contributions to E_{XC} across the grid point volume w.r.t. the change of the density across the grid point volume. With a sufficiently fine grid point sampling, the difference between the continuous integral and the discrete would be negligible, but with the default grids in most codes, this is not usually the case. An alternative interpretation is that discretised V_{XC} includes high-frequency elements that cannot be represented on the grid, resulting in aliasing. The change in V_{XC} due to the sampling will lead to a change in the final converged density, and so will result in the other energy components of the Hamiltonian changing as well.

Sitkiewicz *et al* show that the choice of sampling in atomic basis set calculations can cause spurious oscillations in the total energy and its derivative as atoms are moved relative to each other [15]. The size of the oscillations is entirely dependent on the choice of XC functional, and for some functionals, the oscillations can only be mitigated by using grid sizes that are very expensive and far exceed those used in most routine calculations [16]. They show how this can lead to difficulties in calculating accurate molecular properties such as vibrational spectra. Similar issues occur in a plane-wave basis set.

2.4. XC integration errors in the plane-wave basis

The plane-wave basis set has a more straightforward expression for the evaluation of the XC energy than some of the schemes in atom-centred basis sets. Each of the N_{points} grid points are given equal weight and the sum is calculated as:

$$E_{\mathrm{XC}} = \frac{1}{N_{\mathrm{points}}} \sum_{\mathbf{r}} n(\mathbf{r})\, \epsilon_{\mathrm{XC}}\left(n(\mathbf{r}), \ldots\right). \tag{8}$$

The value of $\epsilon_{\mathrm{XC}}(\mathbf{r})$ is evaluated at point $\mathbf{r}$ and is assumed to be constant across the volume of the grid point. This is inaccurate since the calculation involves non-linear functions of the density that must introduce Fourier components with magnitude $|G| > 2G_{\mathrm{cut}}$ which are therefore aliased to lower frequencies. A different choice of sampling points, that is, a constant offset applied to the position of all points in the grid, would lead to a different E_{XC} and V_{XC} due to a change in the sampling of aliased frequencies [24]. The calculations therefore lose translational invariance.

This is illustrated in figure 1 which shows the change in the energy of an isolated helium atom in a periodic box as it is moved around the unit cell [25]. The results were obtained with the CASTEP code using the 'on the fly' norm-conserving pseudo-potentials generated using the descriptors in the CASTEP database. The density and wavefunctions were converged until the change in the forces between successive iterations was less 1×10^{-8} eV Å^{-1}. This is excessive for most routine calculations, but was done to ensure the accuracy of the change in forces under displacement.

A change in sampling points by adding a constant offset to the grid positions is equivalent to each atom in the cell having the opposite shift applied while the grid remains stationary. It can be seen from figure 1 that the energy is not constant as might be expected. Figures 2 and 3 show how the energy and force on the helium atom change as the atom is moved relative to the grid for different XC functionals. When no XC functional is used (or the XC contribution is zeroed), then there is no change in energy, i.e. the simulations are translationally invariant. However, when an XC functional is used, the energy changes and translational invariance is lost. This spurious change in energy gives rise to the rather odd predicted result of an isolated helium atom in its ground state experiencing a small, spontaneous, position-dependent force.

The results shown in figure 2 are in keeping with the investigation of Lehtola and Marques [8] into the numerical behaviour of different functionals. The change in energy of the Perdew–Wang parametrisation of LDA [22] is practically negligible, but for PBE [23] and r2SCAN [9], GGA and meta-GGA functionals, respectively, the energy change is much greater. Lehtola and Marques also found the Perdew–Zunger (PZ) parametrisation of the LDA functional [21] to have large numerical instability [8], failing to converge the

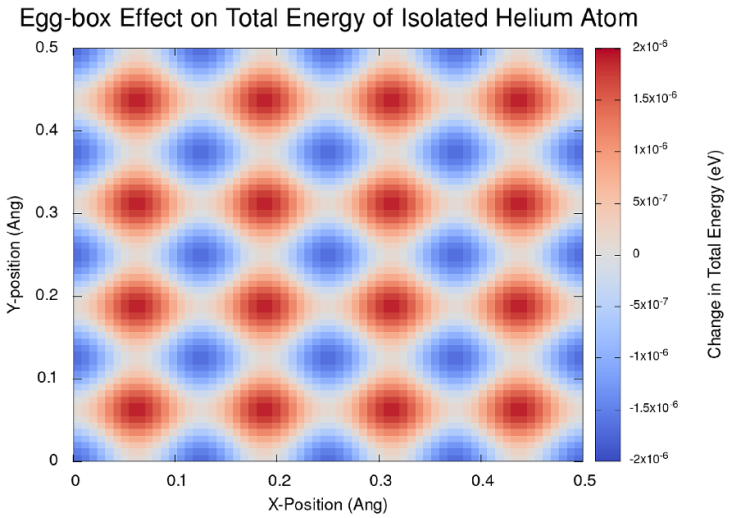

Figure 1. Change in the total energy of an isolated helium atom as the atom is moved around a $(10\ \text{Å})^3$ cubic unit cell. An $(80)^3$ FFT grid was used to represent the density and a single k-point at the Gamma point was used for Brillouin zone integration. The calculation was performed with the PBE functional [23].

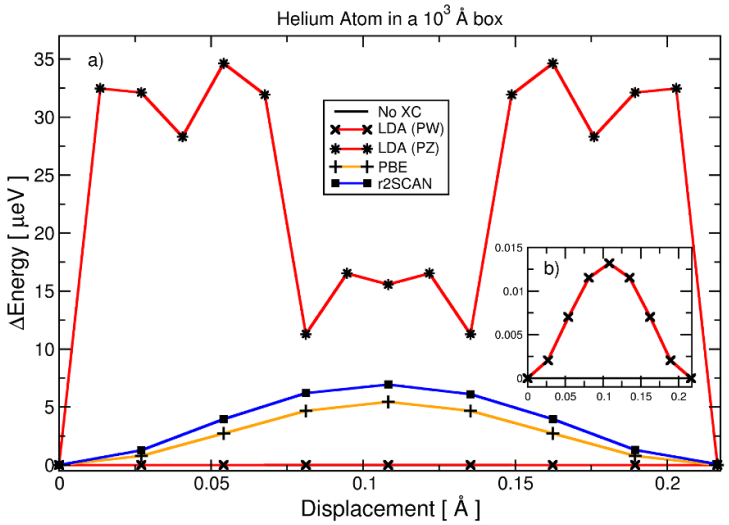

Figure 2. Change in the total energy of an isolated helium atom as the atom is displaced along the (111) direction of an $(10\ \text{Å})^3$ cubic unit cell. A 80^3 FFT grid to represent the density and a single k-point at the Gamma point was used for Brillouin zone integration. Sub-figure (b) is the same plot with the force axis scaled to allow the LDA (PW) [22] result to be seen.

energy w.r.t. grid resolution below $\approx 10\ \mu$eV. The results shown in figure 2 supports their findings as the LDA-PZ parametrisation has the largest variation in energy of all the functionals tested and is seen to be the most non-linear in behaviour. This ill-behaviour is attributed to the functional form of LDA-PZ, which is a piecewise function of the density with a cusp.

There have been various methods of dealing with the consequences of the integration errors that have developed over the years, perhaps without it always being clear why these corrective schemes are necessary. The most obvious consequence is the violation of translational symmetry leading, as seen above, to net forces on the system. The spurious net force is clearly of concern in both isolated systems and crystal structures. This is often assumed to be the result of numerical noise and under-converged calculations. However, as is shown in figure 3, even very well converged systems can still have forces on the centre of mass (CoM). This force on the CoM is consistent with the energy change resulting from a change in sampling as the system moves relative to the discrete grid.

Permitting forces on the CoM mean that in a geometry optimisation or a molecular dynamics (MD) calculation, the entire system can spontaneously begin to drift which is non-physical. In most cases this is corrected by determining the average force acting on all atoms and removing it from each individual atomic

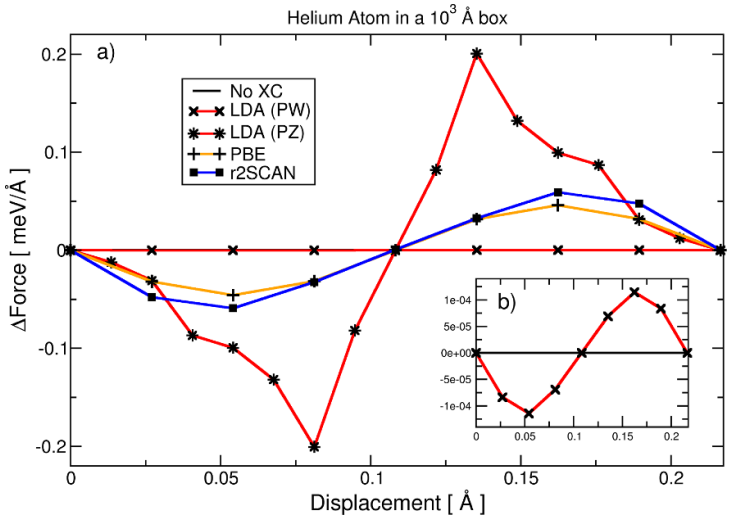

Figure 3. Change in the spurious force on an isolated helium atom as the atom is displaced along the (111) direction of an $(10\ \text{Å})^3$ cubic unit cell. A 80^3 FFT grid to represent the density and a single k-point at the Gamma point was used for Brillouin zone integration. Sub-figure (b) is the same plot with the force axis scaled to allow the LDA (PW) result to be seen.

force, which therefore zeros the net force on the CoM. Although this fixes the CoM position, it fails to address the *cause* of the force on the CoM.

This spurious position-dependent effect is caused by the existence of preferred low-energy locations relative to the grid for each atom. This introduces systematic errors in the energy landscape, in turn creating systematic errors in forces and stresses, which in turn leads to complications in MD and/or geometry optimisation calculations. For example, the preferred low-energy locations might not respect crystal symmetry, leading to difficulties in finding the correct geometries. Enforcing the crystal symmetry for known structures can aid in geometry optimisation, but if there is little symmetry or the symmetry is unknown, then the systematic errors must be overcome in some other way.

When performing lattice dynamics (phonon) calculations, further issues are caused by numerical errors in E_{XC}. Phonon calculations are calculating the derivative of the atomic forces due to an atomic displacement in the system, from which the dynamical matrix can be constructed and diagonalised, yielding the phonon modes. The necessity for an acoustic sum rule (ASR) correction in almost all DFT phonon calculations to force the acoustic modes to have zero energy at Γ, is due to the discretisation of the contributions to E_{XC} on a real space grid. This was noted by Gonze and Lee [26] when first suggesting their original ASR correction.

ASR corrections are similar to zeroing the CoM force, and some ASR corrections can be interpreted as accounting for the CoM force contribution to the force derivatives. An ASR correction is useful when there is slight numerical noise in the forces but does not adequately address the cause of the discrepancies [27]. When the error due to the XC integration is large, the energy landscape becomes very complicated, leading to severe problems when trying to take numerical derivatives of the forces. An ASR correction can still be applied to force the acoustic modes to zero at Γ, but the correction cannot account for the modes away from the Γ point as seen in figure 4.

It is often found that phonon calculations require higher cut-off energies or finer grid sampling for the density, but the reason for this is not always properly appreciated. Finite difference schemes that make use of supercells suffer particularly from the density sampling problem as the different discretisation in different supercells can lead to large changes in the XC integration error in each supercell, such that the overall error no longer cancels. For example, consider a system where the ground-state of the primitive cell is found using a particular cut-off wavevector G_{cut} of the wavefunction, for which the corresponding FFT grid spacing would require a $9.2 \times 9.2 \times 9.2$ grid. Since every grid must have an integer number of points, the dimensions of the FFT grid for the wavefunction are rounded up to give a $10 \times 10 \times 10$ grid. The density grid is twice the size, so will require a $20 \times 20 \times 20$ FFT grid. For a $3 \times 3 \times 3$ supercell constructed from the primitive cell using the same wavefunction cut-off G_{cut}, the grid spacing now corresponds to a $27.6 \times 27.6 \times 27.6$ FFT grid. This is again rounded up, this time to a $28 \times 28 \times 28$ grid for the wavefunction and a $56 \times 56 \times 56$ FFT grid for the density. The rounding of the grids means the supercell's grid is *not* simply three times the original cell's grid dimensions (which would have been $30 \times 30 \times 30$), and the spacing between points will be different between the two cases. This change in grid spacing causes all real-space energy contributions to be altered slightly. However, because the XC energy is particularly sensitive to the real space grid used, there can

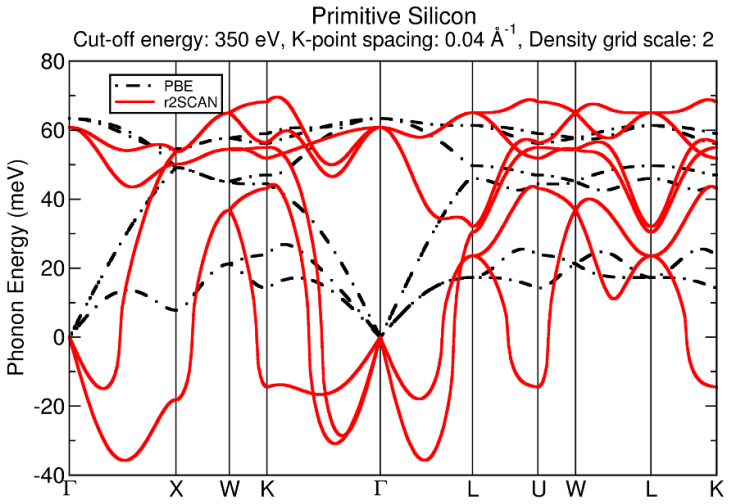

Figure 4. Phonon spectra calculated for a silicon primitive cell using the PBE and r2SCAN functionals. Calculated via finite displacement using the non-diagonal supercell method [28] equivalent to a $3 \times 3 \times 3$ q-point sampling grid.

be a very large change in E_{XC}. These changes can mean that the ground-state structure found for the primitive cell may not be equivalent to the ground-state structure in the supercell. This is particularly a problem for the non-diagonal supercell method as this method can generate supercells that involve rotations of the lattice vectors, which can therefore produce supercells where constructing an identical grid sampling to the primitive cell is impossible. In the case of diagonal supercells, the change in sampling can be avoided by ensuring that the same scaling is applied to the density grid as is applied to the lattice vectors. However, this does not avoid the systematic errors in the energy landscape that are induced by computing E_{XC} on a discrete grid. Displacing the atoms to perform a finite difference in the forces will mean that the finite derivative includes the error due to the sampling of E_{XC}. Simply increasing the cut-off energy and density grid may mask this issue, but at the cost of making phonon calculations even more computationally expensive.

The nature of phonon calculations involves perturbing the ground state structure, which will break some of the crystal symmetries. If symmetry has been enforced during the geometry optimisation, the structure will have been optimised until the symmetry-constrained forces are zero, regardless of whether these symmetries are obeyed by the underlying density grid. Turning off the symmetry-constraints may therefore lead to non-zero forces even in the unperturbed system, meaning that the 'relaxed' geometry is not in fact the ground state structure. This issue is exacerbated if the CoM force is also zeroed, which is equivalent to enforcing continuous translational symmetry for the overall system.

Ordinarily, if the energy of the system decreases under an atomic perturbation, this indicates an imaginary phonon frequency and that the system is mechanically unstable to this perturbation. However, if the perturbation breaks symmetries, thus removing constraints on the energy landscape, this can allow the perturbed energy to be lowered due to numerical artefacts from the XC energy integration, regardless of the physical stability. This can lead to spurious imaginary phonon modes appearing in the calculated spectrum. These consequences can be seen in figure 4 where the phonon dispersion plot has been calculated for silicon using the PBE functional and the r2SCAN functional using the non-diagonal supercell method [28] equivalent to a $3 \times 3 \times 3$ q-point sampling grid.

These are relatively naïve calculations in terms of the basis set used, yet the PBE calculation still managed to produce physically reasonable results. However, the r2SCAN calculation returns a phonon spectrum with several imaginary modes (plotted as negative modes for convenience), suggesting that the FCC primitive cell for silicon is unstable, which we know to be physically incorrect. In this case the apparently unstable modes are not caused by structural instability, they are caused by a change in discretisation between primitive cell and super-cell and breaking of the crystal symmetry from the perturbation. Note that the only reason the modes are zero at the gamma point is due to the CoM force being zeroed. In this case, applying an ASR correction has no effect on the modes.

In most cases, the errors due to the numerical integration error in E_{XC} are small enough to be negligible, particularly with a more simple XC functional such as the Perdew-Wang (PW) parametrisation of the LDA. When the XC integration error is large enough to cause issues, it can be mitigated by using higher cut-off energies or finer grids to represent the electron density. As noted in the literature [5, 6, 8], several modern

density functionals have particular numerical pathologies. This means that it can be very difficult to eliminate numerical errors by increasing the basis set until results converge.

It is rarely obvious that a particular combination of system, functional, and basis set will cause numerical issues due to XC energy integration until long dynamical or after structural calculations have been run. This is especially true when the CoM force has been zeroed and crystal symmetry enforced, which can easily hide such problems until dynamical calculations of the system are performed or the crystal symmetry is broken by some perturbation, as in a phonon calculation. Clearly, it would be advantageous to know the uncertainty that might be introduced from the numerical integration error before running computationally expensive simulations.

3. Uncertainty quantification (UQ)

In this section we present our approach to quantifying the error related to the integration of the contributions to the XC energy. As already demonstrated, these errors have important implications for the calculation of material properties. It would therefore be useful for the user to have an indication of the level of error in the calculation of the total energy due to XC energy integration before completing long structural, dynamical, or phonon calculations. Our aim is to provide an 'on the fly' UQ method that is; (i) highly accurate, (ii) general to all XC functional approximations and all systems, i.e. requires little to no fine tuning, (iii) computationally cheap enough to allow the method to be used in live MD calculations and phonon calculations, warning a user when results may be affected.

The approach we take for our UQ method is to focus on the breaking of translation invariance, i.e. change in E_{XC} with the position of the FFT grid. The ground state energy of the system, E_0, can be determined as follows:

$$E_0 = \langle\psi|\hat{H}|\psi\rangle, \tag{9}$$

where $\hat{H}$ is the Kohn–Sham Hamiltonian as defined in equation (4). Under some perturbation of strength λ, the energy will change. The new energy of the system, E_λ can be expressed as,

$$E_\lambda = \langle\psi_\lambda|\hat{H}_\lambda|\psi_\lambda\rangle \approx \left\langle \psi + \lambda\frac{\partial\psi}{\partial\lambda}\middle|\hat{H} + \lambda\frac{\partial\hat{H}}{\partial\lambda}\middle|\psi + \lambda\frac{\partial\psi}{\partial\lambda}\right\rangle. \tag{10}$$

At the ground state of the unperturbed system, the energy is minimised w.r.t. all possible variations of the wavefunctions. Therefore, provided there is no explicit dependence of ψ on λ, using the same arguments as in Hellmann–Feynman [29], we can eliminate the $\frac{\partial\psi}{\partial\lambda}$ terms. This leaves us with,

$$E_\lambda = \langle\psi|\hat{H} + \lambda\frac{\partial\hat{H}}{\partial\lambda}|\psi\rangle = \langle\psi|\hat{H}_\lambda|\psi\rangle. \tag{11}$$

In order to determine the energy of the perturbed system, we only have to apply the perturbed Hamiltonian to the unperturbed ground state.

The perturbation we are interested in this case is a change in the position of real-space FFT sampling points. In plane-wave DFT, almost all the energy contributions are handled in reciprocal space. The XC energy contribution is the only contribution that must be evaluated explicitly at the real space grid points, and therefore is sensitive to the relative position of the grid. The energy change due to the real space position of the FFT grid being shifted by an offset $\Delta\mathbf{r}$ can therefore be calculated by Fourier interpolating the ground-state charge density onto an FFT grid offset by $\Delta\mathbf{r}$ and recalculating the XC energy. The change in XC energy between the shifted grid positions will be the same as the change in total energy were all the atoms in the system moved relative to the real space grid.

The charge density is transformed from the original FFT grid to the offset FFT grid by applying a phase factor to the terms of the Fourier expansion of the charge density and then performing an inverse transform. The value of the charge density at a particular point, $\mathbf{r} + \Delta\mathbf{r}$ can be written in terms of the original FFT grid expansion as follows:

$$n(\mathbf{r} + \Delta\mathbf{r}) = \sum_{\mathbf{G}} \left(a_{\mathbf{G}} e^{i\mathbf{G}\cdot\mathbf{r}}\right) e^{i\mathbf{G}\cdot\Delta\mathbf{r}}. \tag{12}$$

In figure 5 there is a comparison of the change in energy found from rigidly shifting the atoms of various systems by $\Delta\mathbf{r}$ and reminimising the total energy, and the change in E_{XC} found from interpolating the density to an FFT grid that has had a $-\Delta\mathbf{r}$ rigid shift applied (which is an equivalent transformation). It can be clearly seen that the calculated changes in energy from both approaches are in excellent agreement across a

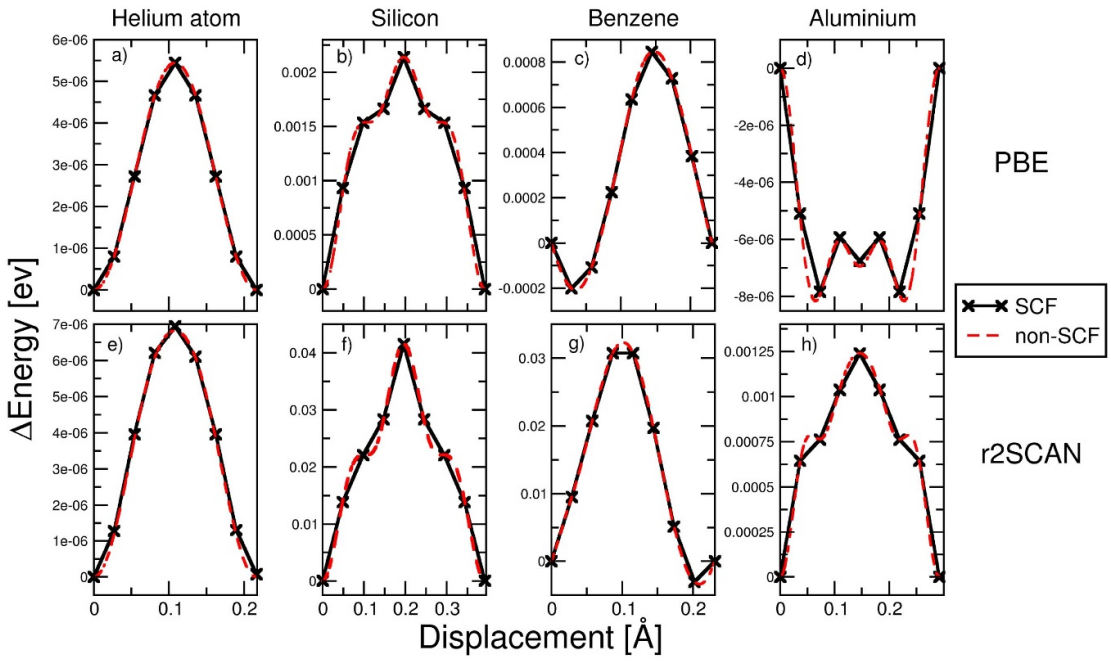

Figure 5. The change in energy, ΔE, calculated for various systems as the atoms are uniformly displaced by $\Delta\mathbf{r}$ along the (111) direction the width of a grid point. In each column we look at a different system, an isolated helium atom (graphs (a) and (e)), a primitive silicon unit cell (graphs (b) and (f)), a Benzene molecule (graphs (c) and (g)) and aluminium (graphs (d) and (h)). For each system two XC functionals were tested, the top row shows the PBE results and the bottom row shows the r2SCAN results. Calculations for the same systems used consistent basis-sets and the norm-conserving pseudopotentials obtained using the same XC functional. The solid black line denotes values the change in the total energy post minimisation. The red dashed lines denotes change in XC energy from applying the offset to the density via interpolation as shown in equation (12) and non-self consistently recalculating the XC energy.

range of different systems. This result means that we are able to probe the precision of the XC integration grid and provide an estimate of the uncertainty.

The interpolation method is general to all XC functionals, as all that is required is the translation of the density (and KED if necessary) to an offset FFT grid, from which any gradient terms can be calculated and E_{XC} re-evaluated. The interpolation method comes at the cost of a few additional Fourier transforms and the evaluation of the XC energy over the grid, but this is significantly cheaper than reminimising calculation with the atoms at different positions. As shown in figure 5, the method has been tested on a variety of systems, and in all cases the agreement between the methods is excellent.

So far, the methods described are not a complete UQ method. The issues from numerical errors in E_{XC} arise when an offset is applied to either the atoms or the real-space grid, violating translational symmetry. A useful heuristic test would therefore be the size of this violation, i.e. an estimation of the maximum change in energy under translation. With periodic boundary conditions, applying a grid offset that translates the grid points exactly so that they sit on another grid point will give the same discretisation of the density and therefore the same energy. It can be seen in figure 1 that the oscillations in energy as the grids/atoms are moved have periodicity of the distance between grid points. Therefore, to sample the possible changes in energy under any arbitrary translation, only translations that span the volume of a single grid point must be considered. By sampling across the volume of the grid point using a regular grid mesh, we can sample the oscillations in energy, i.e. the violation of translational symmetry.

In figure 6, this approach is used to estimate the maximum violation of translational symmetry found when using UQ sampling grids of different sizes. The same test system is used as in figure 5(f), where the difference between the peak and the trough in that graph is approximately 41 meV, which is also the value at which the UQ test converges. It can also be seen from figure 6 that the UQ sampling grid does not have to be very fine to give a reasonable estimate for the violation of translational symmetry and therefore the error in E_{XC}. Knowing the order of magnitude of the error is useful in deciding whether the error is at an acceptable level. Although the cost of determining the energy change from a particular translation is relatively inexpensive, using a very fine sampling for the UQ sampling grid can quickly become prohibitively expensive, involving hundreds of FFTs and evaluations of the XC energy. In practice, a $4 \times 4 \times 4$ UQ sampling grid was found to be sufficient to obtain a useful idea for the magnitude of the violation of translational symmetry. This allows the UQ test to remain inexpensive whilst still giving useful heuristic information.

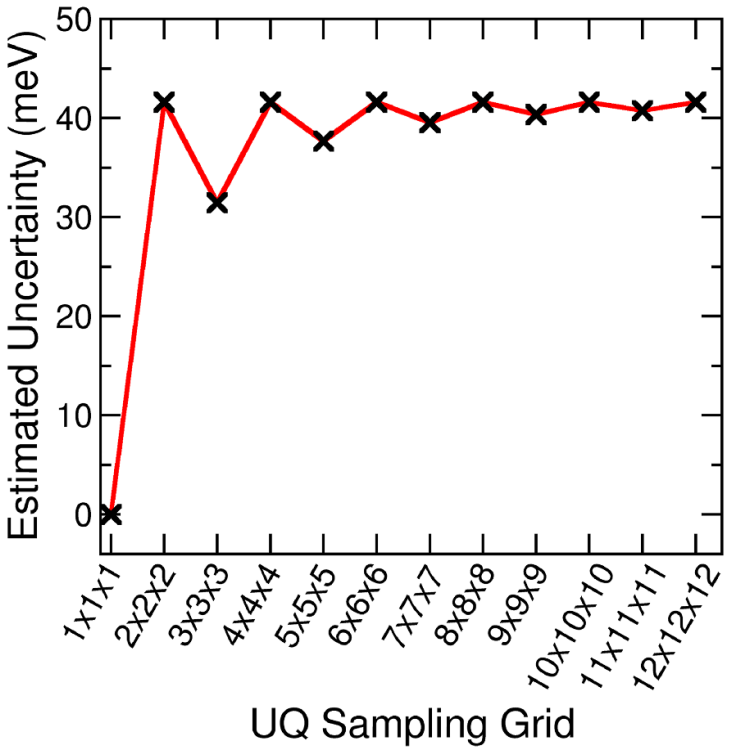

Figure 6. The convergence of the estimated uncertainty due to the XC energy integration as the sampling of each grid point volume is increase. The test system was a primitive unit cell of silicon using the r2SCAN functional with a cut-off energy of 350 eV, a *k*-point spacing of 0.04 and density grid scale of 2.

4. Correcting the XC energy error

The UQ test described in the previous section can be used to provide the user of a plane-wave DFT software with information on how large the violation of translation symmetry is for the given calculation. The advantage of the UQ test is that the user now has a single metric from a relatively cheap test, which can then be used to tune their choice of basis set and XC functional without needing to run expensive calculations, which may eventually fail and need to be repeated. Different material properties of interest may require different levels of accuracy, and it is not always obvious whether that level of accuracy can be achieved. For example, the magnetic anisotropy energy is often $\mathcal{O}(10^{-6})$ eV for weakly anisotropic magnetic materials, and we have already seen (e.g. figure 6) that the errors due to the egg-box effect can be four orders of magnitude greater than this. This UQ can inform the user if such a level of accuracy has been achieved, and if not, then the user can then take action to address the cause of the error in the XC energy integral and thus reduce the violation of translational invariance.

Since in this work we focus on plane-wave approaches, and therefore periodic boundary conditions, the XC energy integral is performed by Fourier-expanding the XC energy density, and the integral over the simulation cell is simply the zeroth order term of this Fourier expansion. Increasing the order of the expansion simply requires increasing the number of points in the Fourier-space sampling grid, which is equivalent to using a finer real-space sampling. As finer and finer sampling is used, the numerical evaluation of the integral, and hence the value of E_{XC}, approaches the exact result. Moreover, $V_{\mathrm{XC}}(\mathbf{r})$ at each point will be more accurate and better represent the *change* in energy with respect to the density, as more high-frequency components of $V_{\mathrm{XC}}(\mathbf{r})$ will be correctly sampled, with correspondingly reduced aliasing in lower frequencies. The combination of these effects should lead to improved representation of the ground state and the physics of the system. It will reduce the magnitude of symmetry violations, including both crystal symmetry and translational symmetry, generating more accurate results.

In plane-wave codes, there are a couple of common approaches that one can take to increase the sampling resolution of E_{XC}, neither of which is ideal, as we will explain. The first common approach is to increase the plane-wave cut-off energy, i.e. increase the number of basis functions in the Fourier expansion of the wavefunction. This has the side effect of significantly increasing the computational cost of every operation in the Hamiltonian, increasing the memory requirements for the calculation, and adding additional Fourier components to the density, and therefore to the XC integral. These extra components will

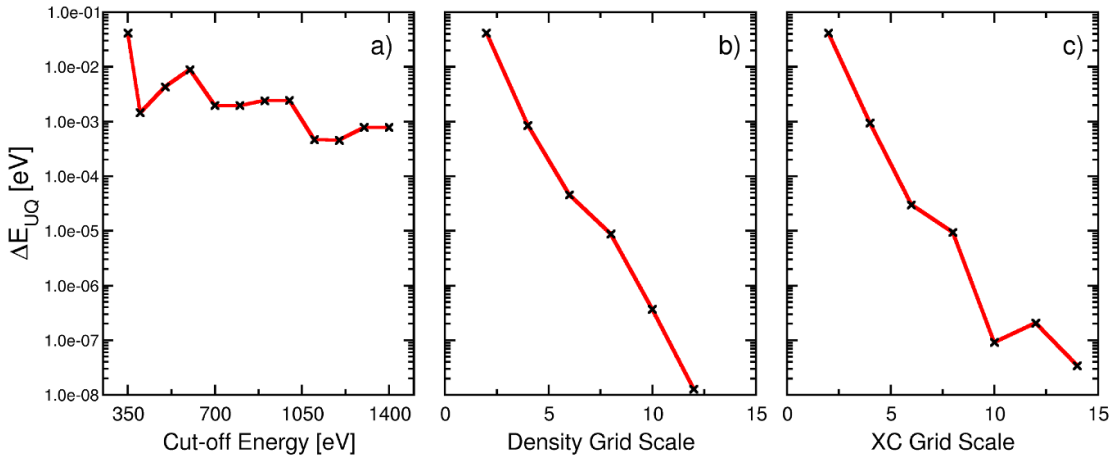

Figure 7. The change in the estimated uncertainty of the XC energy integration as the grid is altered in different ways. The base system is the primitive unit cell of silicon using; a cut-off energy of 350 eV, a density grid scale of 2 and a XC grid scale of 2. In each of the subplots we show the change in estimated uncertainty as; (a) the cut-off energy is increased; (b) the density grid scale is increased; (c) the XC grid scale is increased. In each case the other parameters have been fixed at those used for the first data point.

add higher frequency terms, which should reduce in magnitude, but may also make the reduction of numerical artefacts more difficult.

The other common approach is to increase the fineness of the grid used for the density, leaving the basis functions for the wavefunctions untouched. This only increases the cost of operations involving the density and the construction of local potentials, which is computationally cheaper in both number of operations and the memory required. The caveat is that, although this will improve the XC terms, it costs a considerable amount of wasted effort/memory for many of the other terms in the Hamiltonian. There are also issues with this approach arising from using pseudo-potentials. Both the projector augmented wave method (PAW) [30] and ultrasoft pseudo-potential method (USP) [31] add additional 'augmentation charge' terms to the valence density which can require the density grid to be more than twice as fine as the wavefunction basis in order to be accurately represented (traditional norm-conserving pseudo-potentials [32] do not have these augmentation terms). Many pseudo-potentials also make use of nonlinear core corrections (NLCCs) [33] when computing the XC energy integral for a better representation of the XC interaction between the core charges and the electrons treated as valence. The core charges are interpolated from a spherically symmetric radial function onto the density grid. Both of these pseudo-potential features can mean that increasing the scale of the density grid will have a similar effect to increasing the cut-off energy, i.e. increasing the number of terms in the density expansion and therefore the XC energy, making the reduction of numerical artefacts harder.

The alternative approach we propose here, is to increase only the fineness of the grid for the XC part of the calculation. Every other term in the Hamiltonian can be handled on a grid which will accurately represent the density (including any augmentation from the PAW/USP methods). When computing the XC terms, the NLCC terms should be interpolated onto the density grid, and then we can use Fourier extrapolation to move the real-space density to a much finer grid, i.e. Fourier transform the density to reciprocal space, pad the Fourier expansion with extra zeros up to the highest frequency we want to use in the E_{XC} computation, and then back transform the density to real space on a much finer grid. We then compute the contributions to E_{XC} and V_{XC} on this finer grid in real-space, transform to Fourier space to truncate the potential beyond the highest frequency components representable on the original density grid, and then back transform it to real space on the original density grid. This truncation has no effect on the XC energy integral computed by multiplying the density and the potential compared to when they are computed on the finer XC grid, but it removes any aliasing of high-frequency components represented on the XC grid into the potential on the original density grid.

A comparison of these approaches is shown in figure 7 where we show how the error estimated by the UQ test, ΔE_{UQ}, changes with the fineness of the grid by the methods described above. For the UQ tests we use a $4 \times 4 \times 4$ UQ sampling grid to estimate the error. All approaches were found to be effective in reducing the estimated uncertainty. The approaches of increasing the density grid scale and the XC grid scale are equally and highly effective in reducing the estimated uncertainty. Increasing the cut-off energy does cause a reduction in the estimated error; however, figure 7(a) shows that the cut-off energy must be quadrupled to reduce the error by one order of magnitude. The cut-off wave vector for the plane-waves, G_{cut}, scales as the

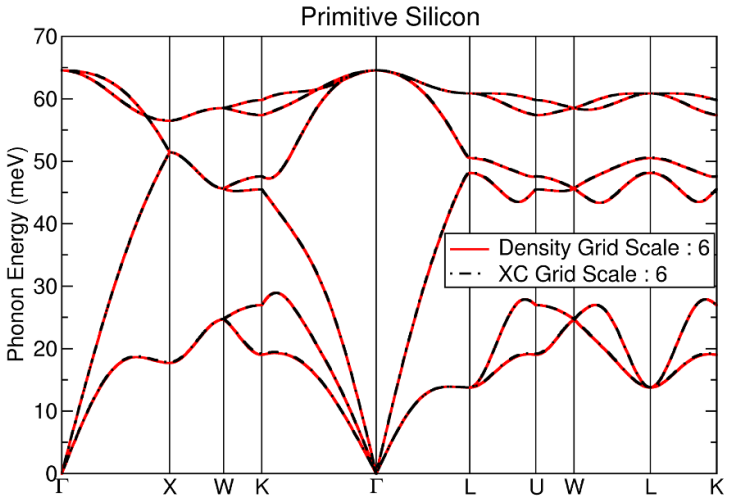

Figure 8. Phonon spectrum calculated for a silicon primitive cell with the r2SCAN functional. Calculated via finite displacement using the non-diagonal supercell method [28] equivalent to a $3 \times 3 \times 3$ q-point sampling grid. All other parameters are the same as those used in figure 4.

root of the cut-off energy, E_{cut}, therefore quadrupling the cut-off energy doubles the size of the density grid. The calculation with a cut-off energy of 1400 eV with a density grid scale of 2 will use the same actual density grid for the XC calculation as that used with a cut-off energy of 350 eV but with a density grid scale or XC grid scale of 4. The estimated uncertainty in these calculations is very similar, but increasing the cut-off energy by a factor of 4 causes the calculation to take ∼8 times as long to complete. Increasing the density grid or the XC grid scale by an extra factor of 2 only increased the cost by a factor of ∼40% and ∼20% respectively, and so clearly these are much more efficient approaches at reducing the uncertainty compared to the cut-off energy. Figure 7(a) only shows the estimated uncertainty for cut-off energies in the range of 350–1400 eV. Further increases in the cut-off energy would further reduce the uncertainty; however, the additional computational cost is dramatic and unjustified.

Increasing the density grid scale compared to increasing the XC grid scale has a comparable effect in terms of accuracy but also comes with additional computational cost. We have already shown in section 3 figure 6 that the effect of the violation of translation symmetry was ≈41 meV in the original calculation without a change in the grid scaling. In figure 7, it can be seen that a grid scale of 8 reduces the error to $\mathcal{O}(10^{-5})$ eV for both the density grid scale calculation and the XC grid scale. Increasing the scale of the density grid to 8 increased the time required for the energy minimisation calculation by a factor of ∼5.5×, whilst increasing the scale of the XC grid to 8 increased the time required by only a factor of ∼2.4×.

Increasing either the cut-off energy, the density grid scale or the XC grid scale leads to the total energy, forces and stresses converging to *different* values. Initially, this might seem worrying, but it is to be expected. As explained above, increasing the XC grid scale *only* improves the representation of the XC energy contributions and the XC potential. Increasing the density grid scale improves the representation of all the density and potential terms in the calculation, including the NLCC and augmentation charge terms, which will include additional contributions to the energy. Increasing the cut-off energy introduces additional plane-waves into the wavefunction and the density, waves which were previously assumed to have zero weighting. This will add extra contributions to every term in the Hamiltonian and will therefore change the total energy calculated.

To illustrate the benefit of using a separate XC grid, we have repeated the calculation of the phonon dispersion presented in section 2, using a density grid scale of 6 and also with an XC grid scale of 6. Without either correction, r2SCAN with the same parameters gives imaginary phonon modes with maximum energy of ∼−40 meV, which is the same magnitude as the energy error the UQ test estimates. Increasing the grid scales to 6 can be seen in figure 7 to reduce the estimated error of the energy to be below 10^{-4} eV, i.e. an order of magnitude smaller than the tolerance in the energies of the phonon modes. With either of the corrections, these imaginary modes are now real, and the difference between the phonon dispersion calculated with either approach is negligible; see figure 8. The UQ test allows an informed decision to be made on how fine a grid is needed, in order to reduce the numerical XC integration artefacts enough for accurate phonon calculations.

It is worth noting again that the UQ test relies only on the ground-state density for a particular atomic configuration. Thus, it can be used to tune the fineness of the XC grid for the desired accuracy *before* before performing any dynamical or structural calculations. This step requires multiple energy minimisation

calculations, with the atoms in a specific configuration, to get the ground state at different XC grid scales, but this is no worse than the convergence calculations required for the choice of cut-off energy and k-point sampling.

It is clear from the uncertainty estimates presented in figure 6 that the use of the finer grid, only to calculate E_{XC} and V_{XC} is an effective method to improve the numerical accuracy of calculations while minimising the additional computational cost. We have also shown that when comparing the phonon dispersion curves presented in figure 8, there is no noticeable difference in the physical predictions compared to the existing approach of increasing the grid scale of the density grid for every energy contribution, while the XC grid method is cheaper to compute.

5. Conclusions

We have developed and presented a new UQ approach for plane-wave DFT calculations that allows an estimation of the violation of translational symmetry due to the numerical integration of the XC energy, sometimes referred to as the egg-box effect. Our novel UQ method fulfils all the desired requirements of being an accurate, efficient, and general-purpose metric for all systems and XC functionals. This is done by Fourier interpolating the ground state density to a set of shifted grid positions and re-evaluating the change in XC energy at these grid offsets. This is found to be in very close agreement with the change in total energy from an equivalent displacement of the entire system.

The results of using our UQ approach are in keeping with those seen elsewhere in the literature, i.e. the size of the integration error is highly dependent on the system of interest, the choice and size of the basis set, and the choice of XC functional. Performing this analysis at an early stage of a workflow can inform the user of the expected accuracy of subsequent calculations.

In addition, we suggest a slightly different approach to improve the E_{XC} calculation. By selectively computing the contributions to E_{XC} and its corresponding potential V_{XC} on a finer grid (via Fourier extrapolation), the integration error can be reduced to a level deemed appropriate by the user, without the additional overhead of increasing the cut-off energy and/or the fineness of the density grid. Coupled with the uncertainty quantification method, the inaccuracies related to the XC integration error (a.k.a. the egg-box effect) can confidently be avoided with a minor additional computational cost. This work offers a promising pathway towards improving accuracy in a diverse range of materials modelling applications.

Data availability statement

The data created and analysed during the current study are available from https://doi.org/10.15124/164d96af-872c-4a91-bb5d-ea99e05bd146.

Acknowledgments

B D acknowledges financial support from EPSRC via a DTP studentship (Grant EP/T518025/1) and P J H was supported by an EPSRC Research Software Engineer Fellowship (Grant EP/R025770/1), with additional funding from the UK Car-Parrinello High End Compute Consortium (UKCP; EPSRC Grant EP/X035891/1).

Conflicts of interest

The authors declare no conflicts of interest.

Author contributions

Ben Durham: Investigation (lead); Methodology (lead); Software (lead); Visualisation (lead); Writing—original draft (lead); Writing—review & editing (equal). Matt Probert: Supervision (equal); Writing—review & editing (equal). Phil Hasnip: Conceptualisation (lead); Funding acquisition (lead); Supervision (equal); Writing—review & editing (equal).

ORCID iDs

Ben Durham https://orcid.org/0000-0002-7201-6867
Matt I J Probert https://orcid.org/0000-0002-1130-9316
Phil J Hasnip https://orcid.org/0000-0002-4314-4093

References

[1] Sun J, Ruzsinszky A and Perdew J P 2015 *Phys. Rev. Lett.* **115** 036402
[2] Bartók A P and Yates J R 2019 *J. Chem. Phys.* **150** 161101
[3] Yamamoto Y, Diaz C M, Basurto L, Jackson K A, Baruah T and Zope R R 2019 *J. Chem. Phys.* **151** 154105
[4] Yang Z H, Peng H, Sun J and Perdew J P 2016 *Phys. Rev.* B **93** 205205
[5] Wheeler S E and Houk K N 2010 *J. Chem. Theory Comput.* **6** 395–404
[6] Johnson E R, Becke A D, Sherrill C D and DiLabio G A 2009 *J. Chem. Phys.* **131** 034111
[7] Furness J W and Sun J 2019 *Phys. Rev.* B **99** 041119
[8] Lehtola S and Marques M A L 2022 *J. Chem. Phys.* **157** 174114
[9] Furness J W, Kaplan A D, Ning J, Perdew J P and Sun J 2020 *J. Phys. Chem. Lett.* **11** 8208
[10] Furness J W, Kaplan A D, Ning J, Perdew J P and Sun J 2022 *J. Chem. Phys.* **156** 034109
[11] Mezei P D, Csonka G I and Kállay M 2018 *J. Chem. Theory Comput.* **14** 2469–79
[12] Tafipolsky M and Schmid R 2006 *J. Chem. Phys.* **124** 174102
[13] Artacho E *et al* 2008 *J. Phys.: Condens. Matter* **20** 064208
[14] Ruiz-Serrano A, Hine N D M and Skylaris C K 2012 *J. Chem. Phys.* **136** 234101
[15] Sitkiewicz S P, Ferradás R R, Ramos-Cordoba E, Zaleśny R, Matito E and Luis J M 2024 *J. Chem. Theory Comput.* **20** 3144–53
[16] Sitkiewicz S P, Zaleśny R, Ramos-Cordoba E, Luis J M and Matito E 2022 *J. Phys. Chem. Lett.* **13** 5963–8
[17] Ning J, Furness J W and Sun J 2022 *Chem. Mater.* **34** 2562–8
[18] Clark S J, Segall M D, Pickard C J, Hasnip P J, Probert M J, Refson K and Payne M 2005 *Z. Kristallogr.* **220** 567–70
[19] Kohn W and Sham L J 1965 *Phys. Rev.* **140** A1133–8
[20] Perdew J P and Schmidt K 2001 *AIP Conf. Proc.* **577** 1–20
[21] Perdew J P and Zunger A 1981 *Phys. Rev.* B **23** 5048–79
[22] Perdew J P and Wang Y 1992 *Phys. Rev.* B **45** 13244–9
[23] Perdew J P, Burke K and Ernzerhof M 1996 *Phys. Rev. Lett.* **77** 3865–8
[24] White J A and Bird D M 1994 *Phys. Rev.* B **50** 4954–7
[25] Durham B, Probert M and Hasnip P J 2024 EggBox paper (University of York) (https://doi.org/10.15124/164d96af-872c-4a91-bb5d-ea99e05bd146)
[26] Gonze X and Lee C 1997 *Phys. Rev.* B **55** 10355–68
[27] Pallikara I, Kayastha P, Skelton J M and Whalley L D 2022 *Electron. Struct.* **4** 033002
[28] Lloyd-Williams J H and Monserrat B 2015 *Phys. Rev.* B **92** 184301
[29] Feynman R P 1939 *Phys. Rev.* **56** 340–3
[30] Blöchl P E 1994 *Phys. Rev.* B **50** 17953–79
[31] Vanderbilt D 1990 *Phys. Rev.* B **41** 7892–5
[32] Hamann D R, Schlüter M and Chiang C 1979 *Phys. Rev. Lett.* **43** 1494–7
[33] Louie S G, Froyen S and Cohen M L 1982 *Phys. Rev.* B **26** 1738–42