

Numerical Quality Control for DFT–based Materials Databases

Christian Carbogno,¹ Kristian Sommer Thygesen,² Björn Bieniek,¹ Claudia Draxl,^{3,1} Luca M. Ghiringhelli,¹ Andris Gulans,³ Oliver T. Hofmann,⁴ Karsten W. Jacobsen,² Sven Lubeck,³ Jens Jørgen Mortensen,² Mikkel Strange,² Elisabeth Wruss,⁴ and Matthias Scheffler¹

¹*Fritz-Haber-Institut der Max-Planck-Gesellschaft, Faradayweg 4–6, D-14195 Berlin, Germany*

²*Center for Atomic-scale Materials Design (CAMD), Department of Physics, Technical University of Denmark, Fysikvej 1 2800 Kgs. Lyngby, Denmark*

³*Physics Department and IRIS Adlershof, Humboldt-Universität zu Berlin, Zum Großen Windkanal 6, D-12489 Berlin, Germany*

⁴*Institute of Solid State Physics, Graz University of Technology, NAWI Graz, Petergasse 16, 8010 Graz, Austria*

(Dated: 25 August 2020)

Electronic-structure theory is a strong pillar of materials science. Many different computer codes that employ different approaches are used by the community to solve various scientific problems. Still, the precision of different packages has only recently been scrutinized thoroughly, focusing on a specific task, namely selecting a popular density functional, and using unusually high, extremely precise numerical settings for investigating 71 monoatomic crystals¹. Little is known, however, about method- and code-specific uncertainties that arise under numerical settings that are commonly used in practice. We shed light on this issue by investigating the deviations in total and relative energies as a function of computational parameters. Using typical settings for basis sets and \mathbf{k} -grids, we compare results for 71 elemental¹ and 63 binary solids obtained by three different electronic-structure codes that employ fundamentally different strategies. On the basis of the observed trends, we propose a simple, analytical model for the estimation of the errors associated with the basis-set incompleteness. We cross-validate this model using ternary systems obtained from the NOMAD Repository and discuss how our approach enables the comparison of the heterogeneous data present in computational materials databases.

Over the last decades, computational materials science has evolved as a paradigm of materials science, complementing theory and experiment with *computer experiments*.² In particular, density-functional theory (DFT) has become the workhorse for a plenitude of computational investigations, representing a good compromise between precision and computational expense, thus allowing for the investigation of realistic systems with affordable numerical effort.³ The widespread application of electronic-structure theory was especially fueled by the development and distribution of many user-friendly and computationally efficient simulation packages (termed *codes* in the following) based on DFT. Essentially all these codes rely on the same fundamental physical concept and solve the Kohn-Sham (KS) equations⁴ of DFT self-consistently by expanding the Kohn-Sham states in a finite basis set. Moreover, apart from the choice of the basis set, different approximations and various numerical techniques and algorithms are employed. Inherently, this raises the question how consistent, and hence, how comparable, results from different codes are.

Only recently, a synergistic community effort led by K. Lejaeghere and S. Cottenier¹ has shed light on these issues, essentially concluding that “most recent codes and methods converge toward a single value”. This concerns, however, only the investigated relatively robust case of computing the equation of states for elemental solids^{1,5} using the PBE exchange-correlation (xc) functional. In this context, it has to be noted that such a close agreement across codes and methods was only achieved by us-

ing *safe* numerical settings that guaranteed highest precision and that are rarely used in routine DFT calculations. In practice, such settings are often not even necessary as long as only data obtained by the same methodology, code, and settings are used, because then one benefits from error cancellation, and trends are described reliably.

Over the last decade, the increased amount of available computational power as well as the maturity of existing first-principles materials-science codes made it possible to perform computational studies in a “high-throughput” fashion by scanning the compositional and structural space in an almost automated manner.^{6–9} In such a case, the numerical settings have to be chosen *a priori* in such a way that the trends of the properties of interest are captured. Often, this is achieved via educated guesses, sometimes via (semi-)automatic algorithms.^{10,11} Since the properties of interest differ in different investigations, also the numerical settings can vary quite significantly.^{12–14} This has some impact on the possibility of reusing data beyond its original scope and purpose. Also, comparing data from different sources – created using different methodologies and settings or focusing on different properties – is not risk-free, in spite of the fact that the data may be publicly available in databases and repositories, as for instance, in the NOMAD Repository,¹⁵ AFLOW,^{16,17} Materials Project,¹⁸ OQMD,¹⁹ Materials Cloud,²⁰ and alike. In a nutshell, using data from different sources that are based on different numerical settings implies potentially uncontrollable uncertainties. This is a pressing and severe issue,

given that the sheer amount of calculations existing to date prevents a human, case-by-case check of the data.

In this work, we describe a first step for overcoming this unsatisfactory situation and show how errors for data stemming from DFT computations can be estimated. We emphasize that we do not investigate errors that originate from the use of approximate *physical* equations, e.g., the use of a particular xc-functional. We rather focus on *numerical* aspects, i.e., on errors arising from the fact that different codes solve the same equation by employing different numerical approximations and techniques. To this end, we systematically investigate the numerical errors that arise in total energies and energy differences in three different DFT codes, the two all-electron, full-potential packages `exciting`²¹ and `FHI-aims`^{22,23} as well as the projector-augmented wave formalism implemented in the package `GPAW`.^{24–27} In the two all-electron, full-potential codes, all electrons are accounted for in the self-consistency cycle on the same footing. The `exciting`²¹ code uses linearized augmented plane-waves plus local orbitals, whereas real-valued, numeric atom-centered orbitals (NAOs) are used as a basis set both for core and valence electrons in `FHI-aims`^{22,23}. Conversely, in the projector augmented wave (PAW) method²⁸ employed in `GPAW`^{24,25}, core states are frozen whereas valence states are mapped onto smooth pseudo-valence states using a linear transformation involving atom-centered partial wave expansions.²⁸ The pseudo-states are smooth and represented in a plane-wave expansion.²⁹ In the following, we evaluate and analyze the numerical errors arising in these different formalisms at various levels of precision and then suggest how to estimate the errors associated with the basis-set incompleteness and, consequently, get access to the complete-basis-set limit for total energies and energy differences.

I. METHODS

To perform the DFT calculations with these three codes in a systematic manner, the atomic simulation environment `ASE`^{26,27} was used to generate the code-specific input files and to store the results using `ASE`'s lightweight database module. In this paper, we focus on the two main numerical approximations that are used to discretize and represent the electron density $n(\mathbf{r}) = \sum_{l\mathbf{k}} |\psi_{l\mathbf{k}}(\mathbf{r})|^2$ via the Kohn-Sham wavefunctions. These are the density of the reciprocal-space grid (\mathbf{k} grid) for Brillouin-zone (BZ) integrations and the finite basis set $\phi_{j\mathbf{k}}(\mathbf{r})$. The Kohn-Sham wavefunctions are written as

$$\psi_{l\mathbf{k}}(\mathbf{r}) = u_{l\mathbf{k}}(\mathbf{r}) \exp(i\mathbf{k}\mathbf{r}) \quad \text{with} \quad u_{l\mathbf{k}}(\mathbf{r}) = \sum_j c_{lj\mathbf{k}} \phi_{j\mathbf{k}}(\mathbf{r}). \quad (1)$$

For the BZ sampling, we use a Γ -centered \mathbf{k} -grid characterized by a uniform \mathbf{k} -point density

$$\rho_{\mathbf{k}} = (N_{\mathbf{k}}/V_{\text{BZ}})^{\frac{1}{3}}, \quad (2)$$

where $N_{\mathbf{k}}$ is the total number of \mathbf{k} -points and V_{BZ} the BZ volume.

To discuss and analyze numerical errors, we performed total-energy calculations for fixed geometries, i.e., without any relaxation, using a representative set of numerical settings. These are \mathbf{k} -point densities of 2, 4, and 8 \AA^{-3} , respectively, and choices of basis sets that are described in detail in the Supplemental Material. They reflect settings typically used in production calculations and also include extremely precise numerical settings that ensure convergence in total energy of < 0.001 eV/atom. The latter are termed “fully converged” reference when we evaluate the error occurring with less precise (typical) settings. To make sure that no other numerical errors cloud the ones stemming from the \mathbf{k} -grid and the basis set, all other computational parameters—for example, the convergence thresholds for self-consistency—were chosen in an extremely conservative way, as detailed in the Supplemental Material.

To cover the chemical space in these benchmark calculations, a set of representative materials has been chosen. This includes the 71 elemental solids that have been studied in the aforementioned work by Lejaeghere and coworkers¹ and also includes prototypical binary materials (one for each element with atomic number ≤ 71 ; noble gases excluded). The atomic structures and detailed geometries were taken from the experimental Springer materials database³⁰ by selecting the energetically most stable binary structure for each particular element.³¹ On top of that, 10 ternary materials were chosen from the NOMAD Repository.³² A detailed list including space groups, stoichiometric formulae, structures, and references to the original scientific publications is given in the Supplemental Material.

In the following, we focus on the convergence and related errors of two fundamental properties, i.e., the absolute total energies E_{tot} and relative energies E_{rel} . The latter were computed as the total-energy difference between the original unit cell and an expanded cell, with 5% larger volume and scaled internal atomic positions. While E_{tot} includes both the energetic contribution from core and valence electrons, E_{rel} is less sensitive to contributions from the core and semi-core electrons due to benign error cancellation. Accordingly, E_{rel} is a good metric to quantify the typically needed numerical precision for energy differences as well as potential-energy surfaces. It also sheds light on the errors that would occur in properties derived from the total energy, like elastic constants, vibrational properties, and alike. In our evaluations, the error for one material i in a data set x_i is always defined with respect to the “fully converged” reference value c_i , as indicated by the notation $\Delta x_i = x_i - c_i$, e.g., $\Delta E_{\text{tot},i}$ for the total energy error of material i . To statistically analyze the errors across the full set of materials with N entries, we report the mean

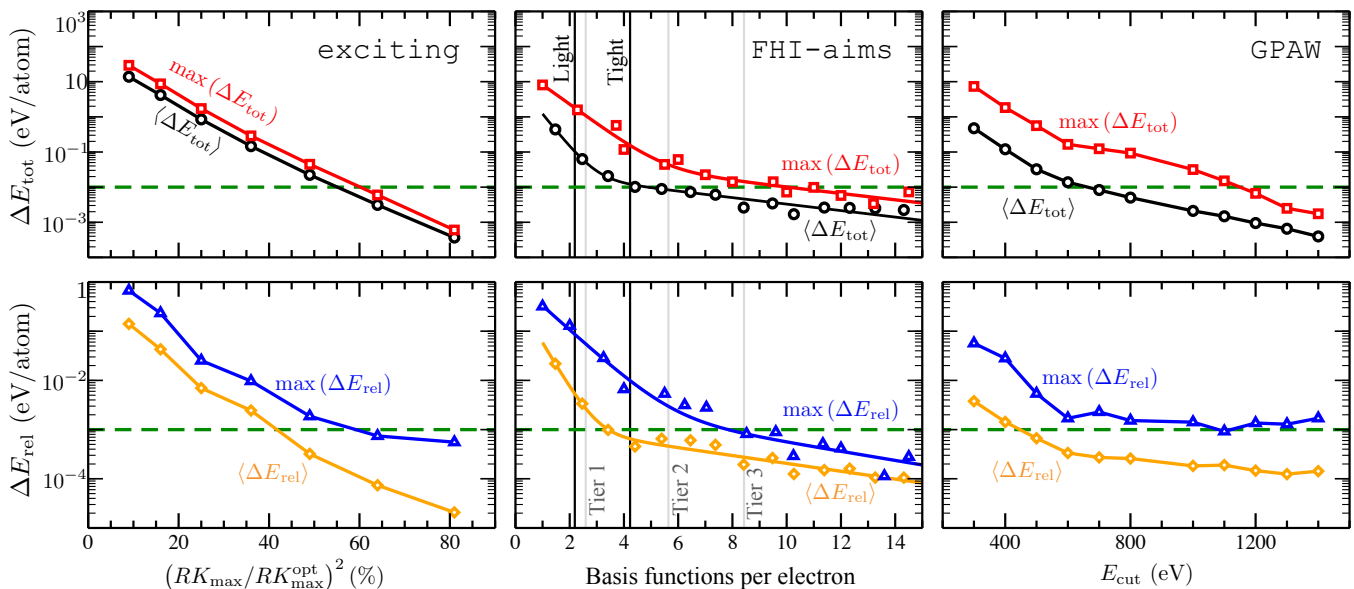


FIG. 1. Average (black/orange) and maximum (red/blue) error in E_{tot} (upper row) and E_{rel} (lower row) as a function of basis-set size as observed in `exciting`, `FHI-aims`, and `GPAW` for the 71 elemental solids (PBE, 8\AA \mathbf{k} -point density). Note that we use a logarithmic scale for the energy axis and different energy windows in the upper and lower row. See text for a discussion of the different metrics used in the different codes to quantify basis set quality. The green dashed horizontal lines indicate values of typically wanted precision.

absolute error

$$\langle \Delta x \rangle = \frac{1}{N} \sum_i^N |\Delta x_i| \quad (3)$$

and the maximum error

$$\max(\Delta x) = \max_i |\Delta x_i|. \quad (4)$$

Here, we limit the discussion to data computed with the PBE xc-functional. The numerical errors occurring with a different type of generalized gradient approximation (GGA) or the local-density approximation (LDA) show the same qualitative behavior and only minor quantitative differences (see Supplemental Material).

II. RESULTS

In the following, we first summarize the trends observed for the elemental solids (Sec. II A). When discussing errors related to the basis set, we always compare to calculations that are “fully converged” with respect to \mathbf{k} -points. Likewise, errors arising from an insufficient \mathbf{k} -point density are discussed for “fully converged” basis sets, since the errors arising from either source can be considered independent of each other. In all cases, a simple summation approach with a Fermi-function smearing of 100 meV is used for the BZ integration. The observed trends allow us to propose a simple, but versatile mathematical model to estimate the error associated to the

basis set for *any* compound and *any* of the investigated codes, as exemplified in Sec. II B for binary and ternary materials.

The results of this work are available for further analysis, both as raw data in the NOMAD Repository³² and as a Jupyter notebook in the NOMAD Data-Analytics Toolkit (<https://analytics-toolkit.nomad-coe.eu/tutorial-error-estimates>). Therein, errors for arbitrary systems can be calculated via an easy-to-use interface for various numerical settings for `exciting`, `FHI-aims`, and `GPAW`. The corresponding Python code can be modified and extended for custom purposes.

A. Elemental Solids

First, we address the convergence with respect to the size viz. degree of completeness of the basis set. The results are shown in Fig. 1. In the case of `exciting`, the atom-specific settings, which are kept fixed in all calculations, correspond to a sizable number of local orbitals that ensure well-converged ground-state calculations and transferability between different compounds. The remaining (and most widely used) parameter to judge the quality of the plane-wave basis is RK_{max} , which is the product of the radius of the smallest atomic sphere and the plane-wave cutoff (for details, see Ref. 21). Choosing the *optimal* value $RK_{\text{max}}^{\text{opt}}$ such that it corresponds to a convergence of the total energy of about 0.1 meV/atom (for details, see the Supplemental Mate-

rial), we use the squared fraction $(RK_{\max}/RK_{\max}^{\text{opt}})^2$ to label the basis set quality, see Supplemental Material for details. For FHI-aims, which uses tabulated, chemical-species-specific sets of NAOs, the number of NAOs per electron is used as metric. Note that these NAOs come in *tiers* that group different angular momenta.²² The average number of basis functions per electron present in these tiers and in the species-specific suggested settings (“light”, “tight”) provided by the FHI-aims developers are shown as black and gray vertical lines in the figures. Since the “translation” from the number of NAOs into this metric requires binning (not all elemental solids appear for all values of the x-axis), the reported errors do not decrease monotonically. It is important to note that tier 4 sets are not provided for all elements, but only for those species for which such an additional set of basis functions improved the description of the electronic structure during the basis-set construction procedure.²² Accordingly, only these problematic elements determine the errors shown for 9 and more NAOs per electron. The more benign elements, that are already fully converged in this limit, no longer enter the shown average error, since the developer-suggested settings do not allow for more than 8 NAOs per electrons for these species. In the plane-wave code GPAW the basis set is characterized by the cut-off energy E_{cut} , i.e., all plane waves with a kinetic energy smaller than E_{cut} are included in the basis set. Note that this affects the convergence of relative energies, since, for the same value of E_{cut} , cells with different volume contain different number of plane waves.

As evident from Fig. 1, the errors in the total energy exhibit a systematic convergence with increasing basis-set size for all three codes. Generally, the maximum error in the total energy can be even roughly one order of magnitude larger than the average error. This is due to the fact that numerical errors are species specific, i.e., some chemical species require a large basis set to be described precisely. This is reflected by the fact that the difference between average and maximum error is more pronounced in the results for FHI-aims and GPAW (Fig. 1) due to the metric chosen to quantify the basis-set completeness, i.e., the x -axis in this figure. While FHI-aims and GPAW use an absolute metric, exciting uses a relative one, i.e., fractions of species-specific values RK_{\max}^{opt} . In this case, the fact that the developers provide well-balanced, species-specific values for RK_{\max}^{opt} ensures that a similar precision is achieved for all species at a specific fraction of $(RK_{\max}/RK_{\max}^{\text{opt}})^2$. In turn, this leads to a more consistent precision across material space and thus to smaller maximum errors at a given value of $(RK_{\max}/RK_{\max}^{\text{opt}})^2$. For all three codes, the average and maximum errors in total energies are roughly one to two orders of magnitudes larger than the ones for relative energies. Again, this finding reflects that the main source for imprecisions in the total energy is species specific and leads to a beneficial error cancellation in energy differences. We note in passing that this aspect is directly exploited in the VASP code^{33,34} for an auto-

matic convergence correction^{35,36}. In Sec. II B, we will exploit this fact for the three codes exciting, FHI-aims, and GPAW to predict errors *a priori* for multi-component systems using information from the elemental solids, as shown in Fig. 1.

Let us now inspect the errors in total energies that arise due to the finite reciprocal-space grid. Figure 2 shows results for \mathbf{k} -point densities of 2 \AA and 4 \AA . Data obtained with a \mathbf{k} -point density of 8 \AA serves as “fully converged” reference. The rather large observed errors result from the fact that many elemental solids are metallic with a more involved shape of the Fermi-surface, so that a substantial number of \mathbf{k} -points is required to reach convergence. Quite consistently, all codes yield average errors of the same order of magnitude if the same \mathbf{k} -point densities are used, despite the fact that the three codes handle the numerical details of the reciprocal-space integration differently. This is reflected in the maximum errors, which vary slightly more between codes than the average ones. Again, we observe that the maximum error is approximately one order of magnitude larger than the average error.

B. Predicting Errors for Binary and Ternary Systems

Following our discussion of the errors in total and relative energies of elemental solids stemming from the basis-set incompleteness, we propose to estimate the corresponding errors for multicomponent systems by linearly combining the respective errors observed for the constituents in the elemental-solids calculations at the same settings. This follows the above discussed observation that there are chemical species that require larger basis sets to reach convergence. This is in fact independent of




FIG. 2. Average (left) and maximum error (right) observed for the elemental solids for two different \mathbf{k} -point densities $\rho_{\mathbf{k}}$ with exciting, FHI-aims, and GPAW. In all three codes, the calculations were carried out using the PBE xc-functional. A simple summation with a Fermi-function smearing of 100 meV is used for the BZ integration to facilitate comparison between codes.

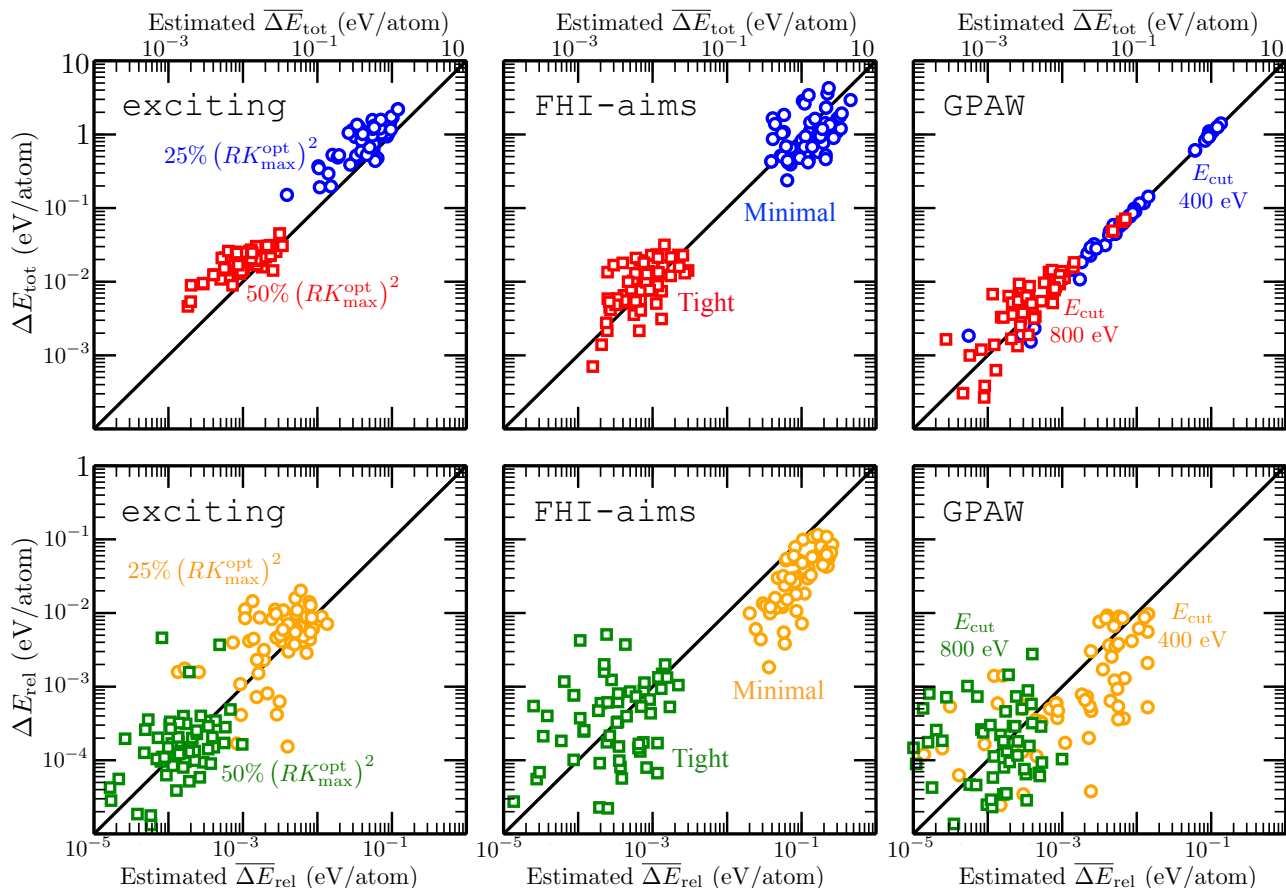


FIG. 3. Estimated vs. actual numerical errors in E_{tot} (upper row) and E_{rel} (lower row) for 63 binary systems, considering two basis-set sizes for each of the three employed codes. The structures were chosen from the experimental Springer materials database³⁰ by selecting the energetically most stable binary structure for each particular element. Note the logarithmic scales and the different energy windows in the upper and lower row.

the employed code. For the error in the total energy,³⁷ we simply assume:

$$\overline{\Delta E}_{\text{tot}} = \frac{1}{N} \sum_I N_I \Delta E_{\text{tot},I} \quad (5)$$

N_I being the number of atoms of species I . For $\overline{\Delta E}_{\text{rel}}$ we proceed analogously.

To validate the ansatz of Eq. (5), we have computed the total and relative energy errors for 63 binary solids using the exact same strategies used for the elemental solids in Sec. II A. In Fig. 3, we then compare these real errors observed in the calculations for binary systems for two basis-set sizes for each of the three codes to the estimated errors obtained via Eq. (5). As shown in these plots, we generally get quite reliable total energy predictions for all three codes by this means. For the total energies (top panels), we observe better predictions when an “unbiased” and smooth metric is used to characterize the basis-set completeness. For instance, GPAW, which uses the atom-independent plane-wave cutoff E_{cut} , yields an almost perfect correlation between predicted and actual total energy errors. Conversely, more scat-

tering is observed for FHI-aims, which uses an atom-specific, granular metric with different NAOs for each atom. Nonetheless, we find a clear correlation between the predicted, $\overline{\Delta E}_{\text{tot}}$, and the actual errors, ΔE_{tot} , for all codes. In particular, this holds for absolute energy errors larger than > 10 meV/atom. This demonstrates that the relatively intuitive relation formulated in Eq. (5) can serve as a reliable estimate for the error associated to a particular total-energy calculation.

For the relative-energy errors shown in the lower half of Fig. 3, we observe more scattering and a less neat correlation between predicted and actual errors. The reason for that is twofold: First, benign error cancellation reduces numerical errors in relative energies, since total energy differences are inspected. In other words, a large portion of the species-specific errors described by Eq. (5) cancel each other out for when computing relative energies as a difference. For this exact reason, relative energies are generally less affected by numerical errors (see Fig. 1 and its discussion). Second, relative errors are –in contrast to total energies– non-variational, i.e., they do not necessarily decrease monotonically with basis-set size. The reason




FIG. 4. Estimated vs. actual numerical error in the total energy of ternary systems, using two basis-set sizes in each of the employed codes. Note that Ca_2CdP , Na_2IrZn , and Li_2SnCl exhibit very similar errors and are thus not fully visible for all codes and levels of precision.

is that the errors associated to the two total energies entering the relative energies typically do not decrease at the exact same rate. Still, the relative energy error estimates for all codes are reliable enough in the respective energy window of interest, hence allowing us to compare relative energies obtained from different codes with different settings.

The data shown and discussed for the binary materials suggest that Eq. (5) can be used to estimate the total energy errors for *any* multi-component system. As an example, we demonstrate this in Fig. 4, in which the same comparison between predicted and actual total energy errors is made for ten ternary systems, which were selected from the huge pool of compounds available in the NOMAD Repository³² so to cover material and structural space. Also in this case, the same quantitative and qualitative behavior as discussed for Fig. 3 is observed. The relatively simple approach of Eq. (5) is able to correctly predict the numerical errors also in these ternary systems. This further substantiates that the described approach is not only applicable to the relatively simple binary systems discussed in Fig. 3, but also to more complex systems, as the ones found in electronic-structure materials databases.

III. CONCLUSIONS

In this study, we have presented an extensive, curated data set obtained by three conceptually very different electronic-structure codes. This set contains elemental solids, binary, and ternary materials for various combinations of computational parameters. The data have been used to understand and predict the errors of calculations with respect to the basis-set quality. More specific, we have shown that the errors for arbitrary systems can be estimated from the errors obtained from systematic calculations for related elemental solids, as exemplified for

63 binaries and 10 ternary systems covering 13 different space groups. The developed formalism, which gives surprisingly good results for total energies despite its conceptual simplicity, can be incorporated into computational materials databases to estimate errors of stored data. This is a prerequisite for operating on data collections that originate from different computations, performed with different computer codes and/or different precision. Our work may serve as a starting point for more sophisticated concepts to quantify numerical errors and uncertainties, especially for more complex materials properties that do not necessarily depend monotonically on the basis-set size, e.g., band gaps, forces, vibrational frequencies, and the relative energies discussed in this work.

IV. DATA AVAILABILITY

All presented data, i.e., in- and output files for all electronic-structure theory codes, is available at the NOMAD Repository³² for `exciting`, `FHI-aims`, `GPAW`, as well as `VASP`. Additionally, the data can be explored interactively using the Jupyter notebook in the NOMAD Data-Analytics Toolkit (<https://analytics-toolkit.nomad-coe.eu/tutorial-error-estimates>).

V. AUTHOR CONTRIBUTIONS

SL and AG performed the `exciting`, BB and CC the `FHI-aims`, MSt and JJM the `GPAW`, and EW and OTH the `VASP` calculations. BB and MSt developed the ASE-based scripts for setting up and performing the calculations; BB, MSt, SL, LMG and CC wrote the Jupyter notebook to evaluate and analyze the data. CD, KST, and MSc ideated the project that was led and coordi-

nated by CC. All authors contributed to the discussion of the results and to the writing of the manuscript.

ACKNOWLEDGMENTS

This project has received funding from the European Unions Horizon 2020 research and innovation program under grant agreement No. 676580 and No. 740233 (TEC1p). OTH and EW gratefully acknowledge funding by the Austrian Science Fund, FWF, under the project P27868-N36.

- ¹K. Lejaeghere, G. Bihlmayer, T. Björkman, P. Blaha, S. Blügel, V. Blum, D. Caliste, I. E. Castelli, S. J. Clark, A. Dal Corso, S. de Gironcoli, T. Deutsch, J. K. Dewhurst, I. Di Marco, C. Draxl, M. Dulak, O. Eriksson, J. A. Flores-Livas, K. F. Garrity, L. Genovese, P. Giannozzi, M. Giantomassi, S. Goedecker, X. Gonze, O. Grånäs, E. K. U. Gross, A. Gulans, F. Gygi, D. R. Hamann, P. J. Hasnip, N. A. W. Holzwarth, D. Iuşan, D. B. Jochym, F. Jollet, D. Jones, G. Kresse, K. Koepnik, E. Küçükbenli, Y. O. Kvashnin, I. L. M. Locht, S. Lubeck, M. Marsman, N. Marzari, U. Nitzsche, L. Nordström, T. Ozaki, L. Paulatto, C. J. Pickard, W. Poelmans, M. I. J. Probert, K. Refson, M. Richter, G.-M. Rignanese, S. Saha, M. Scheffler, M. Schlipf, K. Schwarz, S. Sharma, F. Tavazza, P. Thunström, A. Tkatchenko, M. Torrent, D. Vanderbilt, M. J. van Setten, V. Van Speybroeck, J. M. Wills, J. R. Yates, G.-X. Zhang, and S. Cottenier, *Science* **351** (2016), 10.1126/science.aad3000, <http://science.sciencemag.org/content/351/6280/aad3000.full.pdf>.
- ²C. Draxl and M. Scheffler, “Big-Data-Driven Materials Science and its FAIR Data Infrastructure,” in *Handbook of Materials Modeling*, edited by S. Yip and W. Andreoni (Springer, 2019).
- ³R. O. Jones, *Rev. Mod. Phys.* **87**, 897 (2015).
- ⁴W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
- ⁵K. Lejaeghere, V. V. Speybroeck, G. V. Oost, and S. Cottenier, *Critical Reviews in Solid State and Materials Sciences* **39**, 1 (2014), <http://dx.doi.org/10.1080/10408436.2013.772503>.
- ⁶Curtarolo Stefano, Hart Gus L. W., Nardelli Marco Buongiorno Mingo Natalio, Sanvito Stefano, and Levy Ohad, *Nat Mater* **12**, 191 (2013), 10.1038/nmat3568.
- ⁷M. W. Gaultois, T. D. Sparks, C. K. H. Borg, R. Seshadri, W. D. Bonificio, and D. R. Clarke, *Chemistry of Materials* **25**, 2911 (2013), <http://dx.doi.org/10.1021/cm400893e>.
- ⁸J. E. Saal, S. Kirklin, M. Aykol, B. Meredig, and C. Wolverton, *JOM* **65**, 1501 (2013).
- ⁹J. Hachmann, R. Olivares-Amaya, A. Jinich, A. L. Appleton, M. A. Blood-Forsythe, L. R. Seress, C. Roman-Salgado, K. Treppe, S. Atahan-Evrenk, S. Er, S. Shrestha, R. Mondal, A. Sokolov, Z. Bao, and A. Aspuru-Guzik, *Energy Environ. Sci.* **7**, 698 (2014).
- ¹⁰A. Jain, S. P. Ong, W. Chen, B. Medasani, X. Qu, M. Kocher, M. Brafman, G. Petretto, G.-M. Rignanese, G. Hautier, D. Gunter, and K. A. Persson, *Concurrency and Computation: Practice and Experience* **27**, 5037 (2015), cPE-14-0307.R2.
- ¹¹G. Pizzi, A. Cepellotti, R. Sabatini, N. Marzari, and B. Kozinsky, *Computational Materials Science* **111**, 218 (2016).
- ¹²Fischer Christopher C., Tibbetts Kevin J., Morgan Dane, and Ceder Gerbrand, *Nat Mater* **5**, 641 (2006), 10.1038/nmat1691.
- ¹³S. Wang, Z. Wang, W. Setyawan, N. Mingo, and S. Curtarolo, *Phys. Rev. X* **1**, 021012 (2011).
- ¹⁴I. E. Castelli, T. Olsen, S. Datta, D. D. Landis, S. Dahl, K. S. Thygesen, and K. W. Jacobsen, *Energy Environ. Sci.* **5**, 5814 (2012).
- ¹⁵C. Draxl and M. Scheffler, *MRS Bulletin* (2018).
- ¹⁶C. E. Calderon, J. J. Plata, C. Toher, C. Oses, O. Levy, M. Fornari, A. Natan, M. J. Mehl, G. L. W. Hart, M. Buongiorno Nardelli, and S. Curtarolo, **108 Part A**, 233 (2015).
- ¹⁷S. Curtarolo, W. Setyawan, S. Wang, J. Xue, K. Yang, R. H. Taylor, L. J. Nelson, G. L. W. Hart, S. Sanvito, M. Buongiorno Nardelli, N. Mingo, and O. Levy, **58**, 227 (2012).
- ¹⁸A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. a. Persson, *APL Materials* **1**, 011002 (2013).
- ¹⁹J. E. Saal, S. Kirklin, M. Aykol, B. Meredig, and C. Wolverton, *JOM* **65**, 1501 (2013).
- ²⁰L. Talirz, S. Kumbhar, E. Passaro, A. V. Yakutovich, V. Granata, F. Gargiulo, M. Borelli, M. Uhrin, S. P. Huber, S. Zoupanos, C. S. Adorf, C. W. Andersen, O. Schütt, C. A. Pignedoli, D. Passerone, J. VandeVondele, T. C. Schulthess, B. Smit, G. Pizzi, and N. Marzari, (2020), 2003.12510.
- ²¹A. Gulans, S. Kontur, C. Meisenbichler, D. Nabok, P. Pavone, S. Rigamonti, S. Sagmeister, U. Werner, and C. Draxl, *Journal of Physics: Condensed Matter* **26**, 363202 (2014).
- ²²V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, *Comput. Phys. Commun.* **180**, 2175 (2009).
- ²³V. Havu, V. Blum, P. Havu, and M. Scheffler, *J. Comp. Phys.* **228**, 8367 (2009).
- ²⁴J. J. Mortensen, L. B. Hansen, and K. W. Jacobsen, *Phys. Rev. B* **71**, 035109 (2005).
- ²⁵J. Enkovaara, C. Rostgaard, J. J. Mortensen, J. Chen, M. Duak, L. Ferrighi, J. Gavnholt, C. Glinsvad, V. Haikola, H. A. Hansen, H. H. Kristoffersen, M. Kuisma, A. H. Larsen, L. Lehtovaara, M. Ljungberg, O. Lopez-Acevedo, P. G. Moses, J. Ojanen, T. Olsen, V. Petzold, N. A. Romero, J. Stausholm-Mller, M. Strange, G. A. Tritsarlis, M. Vanin, M. Walter, B. Hammer, H. Hkkinen, G. K. H. Madsen, R. M. Nieminen, J. K. Nrskov, M. Puska, T. T. Rantala, J. Schitz, K. S. Thygesen, and K. W. Jacobsen, *Journal of Physics: Condensed Matter* **22**, 253202 (2010).
- ²⁶A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Duak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schitz, O. Schtt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen, *Journal of Physics: Condensed Matter* **29**, 273002 (2017).
- ²⁷S. R. Bahn and K. W. Jacobsen, *Comput. Sci. Eng.* **4**, 56 (2002).
- ²⁸P. E. Blöchl, *Phys. Rev. B* **50**, 17953 (1994).
- ²⁹Throughout this work, we use the PAW potentials recommended by the GPAW developers.
- ³⁰“Springer materials,” <http://materials.springer.com>.
- ³¹We use the $T = 0K$ experimental geometries. Zero-point vibrational effects are included in these experimental values and are not corrected for in the calculations. This is fine as we only need a consistent treatment for all calculations and materials.
- ³²“The nomad repository,” <https://repository.nomad-coe.eu>.
- ³³G. Kresse and J. Furthmüller, *Comp Mater Sci* **6**, 15 (1996).
- ³⁴G. Kresse and J. Furthmüller, *Phys. Rev. B* **54**, 11169 (1996).
- ³⁵A. M. Rappe, K. M. Rabe, E. Kaxiras, and J. D. Joannopoulos, *Phys. Rev. B* **41**, 1227 (1990).
- ³⁶A complete, consistent VASP dataset covering the materials discussed in this work is available via the NOMAD repository at DOI:10.17172/NOMAD/2020.07.29-1. Its analysis is beyond the scope of this paper due to the automatic convergence corrections.
- ³⁷In the case of O, F, and N, the elemental solid is a molecular crystal that is not a good representative for the binding in oxides, fluorides, and nitrides. For this reason, we determined the values of ΔE for these particular elements from the binaries MgO, NaF, and BN by inverting Eq. (5).