



Cite this: *Mater. Horiz.*, 2026, 13, 1993Received 22nd July 2025,
Accepted 1st December 2025

DOI: 10.1039/d5mh01404b

rsc.li/materials-horizons

The carbon cost of materials discovery: Can machine learning really accelerate the discovery of new photovoltaics?

Matthew Walker  and Keith T. Butler *

Computational screening has become a powerful complement to experimental efforts in the discovery of high-performance photovoltaic (PV) materials. Most workflows rely on density functional theory (DFT) to estimate electronic and optical properties relevant to solar energy conversion. Although more efficient than laboratory-based methods, DFT calculations still entail substantial computational and environmental costs. Machine learning (ML) models have recently gained attention as surrogates for DFT, offering drastic reductions in resource use with competitive predictive performance. In this study, we reproduce a canonical DFT-based workflow to estimate the maximum efficiency limit and progressively replace its components with ML surrogates. By quantifying the CO₂ emissions associated with each computational strategy, we evaluate the trade-offs between predictive efficacy and environmental cost. Our results reveal multiple hybrid ML/DFT strategies that optimize different points along the accuracy–emissions front. We find that direct prediction of scalar quantities, such as maximum efficiency, is significantly more tractable than using predicted absorption spectra as an intermediate step. Interestingly, ML models trained on DFT data can outperform DFT workflows using alternative exchange–correlation functionals in screening applications, highlighting the consistency and utility of data-driven approaches. We also assess strategies to improve ML-driven screening through expanded datasets and improved model architectures tailored to PV-relevant features. This work provides a quantitative framework for building low-emission, high-throughput discovery pipelines.

1. Introduction

The development of new material functionalities has historically been a time-consuming process, with new materials or materials' applications often discovered serendipitously or taking many decades of careful synthesis and characterization

New concepts

This work introduces a novel framework for evaluating materials discovery strategies that explicitly balances predictive performance with environmental impact—specifically carbon emissions. While machine learning (ML) has been widely heralded as a route to accelerate computational screening, our approach is the first to rigorously benchmark the carbon cost of ML-augmented workflows against traditional density functional theory (DFT) pipelines. By treating emissions as a quantifiable design parameter, we reveal trade-offs and “sweet spots” along an accuracy–emissions Pareto front, challenging the prevailing assumption that model accuracy alone should guide methodological choice. Our framework enables materials scientists to make evidence-based decisions about when and how to incorporate ML into discovery campaigns. This concept represents a shift in how computational efficiency is defined—from a purely time- or resource-based metric to one that incorporates sustainability. The additional insight is twofold: (1) we show that certain ML surrogates not only reduce emissions but also outperform higher-fidelity DFT calculations in screening contexts, and (2) we identify clear priorities for future model and dataset development to maximize impact per carbon emitted. As AI becomes increasingly integrated into materials research, our contribution lays the groundwork for responsible, low-emission innovation in computational materials science.

before realizing a real-world application.¹ In recent years, with the increase in computing power and the sophistication of materials modelling software, computational chemistry has promised to accelerate this process by providing qualitative insight and design rules as well as quantitative predictions allowing virtual screening of new materials for given applications. However, atomistic modelling has known limitations for the discovery of new materials. In recent years, the emergence of data-driven approaches, notably machine learning (ML), and the availability of large, high-quality annotated datasets of material properties have been predicted to be a route to accelerate computational materials design. But there are many open questions for applying ML to PV discovery: Are these methods really reliable? Which methods are best suited to the task? How does the quality of the underlying data affect model performance? Finally, what should we actually model? In this

Department of Chemistry, University College London, 20 Gordon Street, London WC1H 0AJ, UK. E-mail: matthew.walker.21@ucl.ac.uk, k.t.butler@ucl.ac.uk



paper, we set out to address some of these questions in the context of photovoltaic (PV) materials discovery and provide some concrete guidelines on how effective ML is for PV discovery currently by estimating the carbon cost of both ML-based and density functional theory (DFT)-based materials discovery. We also provide suggestions on how to advance in the future.

Global PV capacity reached approximately 1.6 TW in 2023,² and a future push toward 30–70 TW by 2050 could see PVs meeting most of the world's energy requirements.³ Achieving this target requires the development of new materials as well as the optimization of existing ones.⁴ While crystalline and multi-crystalline Si modules remain the industrial standard,⁵ alternative materials such as amorphous Si,⁶ CIGS,⁷ CdTe,⁸ organic photovoltaics,⁹ and dye-sensitized solar cells¹⁰ have been commercialized to varying degrees of success. A number of perovskites have also emerged as promising candidates in the last decade.¹¹ However, established technologies often rely on critical raw materials, toxic elements, or suffer from long-term stability issues, conversion efficiency limitations, or low technological flexibility; overcoming these challenges is essential for reaching TW-level production of PV energy.¹²

New inorganic materials offer significant promise as future PV absorbers due to their potential for low-cost fabrication, defect tolerance, earth abundance, and facile synthesis *via* various techniques such as sol-gel processing or sputtering.^{13–16} These materials exhibit stability across a wide range of thermal, chemical, and mechanical conditions and are compatible with device architectures that may offer lower capital costs, enabling rapid scale-up.¹³

Computational modelling has played an important role in the development of new inorganic photovoltaic materials such as CZTS,^{17,18} SnS,¹⁹ BiSI,²⁰ Sb₂Se₃,²¹ CdTe²² and many others. Typically, these studies have been DFT calculations allowing accurate estimation of optical absorption, carrier transport and defect properties.²³ Although these DFT calculations are more efficient than experimental synthesis and characterisation, they nonetheless have a non-negligible energy cost. In recent years there has been a trend to replace some of the costly DFT calculations with ML surrogate models. However, the questions previously raised about the veracity of these models remain largely unanswered.

To address these questions, we have developed a framework that enables the joint assessment of both predictive accuracy and carbon emissions associated with different computational approaches for estimating PV performance in novel inorganic crystalline materials. These approaches span from hybrid-functional DFT (the most computationally expensive) to direct ML estimation of maximum PV efficiency (the least expensive), and include intermediate strategies such as predicting optical absorption profiles or applying corrections to low-fidelity DFT calculations based on the generalized gradient approximation (GGA).²⁴ The paper begins with a detailed outline of our evaluation methodology, covering both PV property estimation and carbon emission quantification. We then compare these approaches in terms of predictive efficacy and environmental

cost. Our analysis allows us to propose optimal trade-offs, highlight important limitations, and suggest promising directions for future research aimed at improving the effectiveness and sustainability of computational PV screening. More broadly, our framework offers a template for evaluating computational discovery pipelines in which resource intensity is considered alongside predictive performance—a consideration we believe will be increasingly important across many areas of energy research.

II. Evaluation methodology

We provide details of the different design choices in our evaluation protocol. Covering approaches to obtain the carbon emissions of calculations, the optical absorption spectra and the maximum PV efficiency.

A. The carbon cost of discovery

Ultimately, we are interested in developing new photovoltaic materials as a renewable energy technology. Therefore, it is important to consider the energy cost of a discovery campaign. While computational discovery is less resource intensive than experimental programmes, it is not carbon neutral. One promise of ML is that it can reduce the computational cost and ultimately the resource required for discovery. To assess the trade-off between prediction accuracy and carbon cost of the calculations involved, we have used the CodeCarbon package,²⁵ which integrates into computational workflows to estimate the CO₂ emissions associated with running a given job by monitoring total energy usage across all processing units. This enables facile comparison of computational chemistry calculations, typically CPU-based, and ML inferences, which mostly use GPUs. CO₂ emissions are then estimated based on the sources of energy for the grid in the location of the computer, in our case the UK. The specific numbers for these emissions are thus very sensitive to change, so we mostly give relative emissions. However, we provide some raw numbers to give context for the scale of the CO₂ emissions associated with the calculations in this work.

B. Calculating maximum efficiencies

The computational design of materials typically relies on the availability of a readily computable figure of merit (FoM), which provides a measure of how good a given material is for an application. In photovoltaics, the detailed balance limit²⁶ gives the maximum achievable power conversion efficiency of a single-junction PV cell as a function of the band gap. This simple FoM assumes a step-function absorptance ($A(E)$), which is particularly inaccurate for indirect band-gap absorbers, which typically show a more gradual onset of absorption.

Instead, the efficiency of potential PV absorber materials can be estimated using the spectroscopic limited maximum efficiency (SLME).²⁷ The theory and practical details of calculating SLMEs are discussed in the following section. To distinguish between these methods (since both use detailed balance), we



shall henceforth refer to methods using a step-function approximation of the absorptance spectrum as ‘step-function methods’ and those that use the calculated/predicted spectrum as ‘SLME methods’.

The SLME of a material requires an absorption profile $\alpha(E)$, usually in units of cm^{-1} , and an ‘offset’ (our taxonomy):

$$\Delta = E_g^{\text{da}} - E_g,$$

where E_g^{da} is the minimum direct, dipole-allowed band gap and E_g is the fundamental band gap which, in contrast, may be indirect and dipole-forbidden. The absorptance, $A(E)$, for a material of thickness d is calculated from the absorption, $\alpha(E)$, using a Lambert–Beer approximation:

$$A(E) = 1 - \exp(-2d\alpha(E)).$$

This is used to calculate the short-circuit current (density):

$$J_{\text{SC}} = e \int_0^\infty A(E) \phi_{\text{AM1.5G}}(E) dE,$$

where $\phi_{\text{AM1.5G}}(E)$ is the spectrum of solar radiation received at ground level on Earth when the Sun is perpendicular and e is the charge on an electron. The internal quantum efficiency is assumed to be one: that is, all photons absorbed contribute to the current.

Detailed balance says that the rate of radiative emission must equal that of photon absorption from the surroundings, which can be quantified using the black-body spectrum at the temperature, T , of the solar device. This gives the reverse saturation current density (or recombination current density) as

$$J_0^{\text{rad}} = \frac{e\pi}{f_r} \int_0^\infty A(E) \phi_{\text{BB}}(E) dE,$$

with the black-body spectrum given by

$$\phi_{\text{BB}}(E) = \frac{2e\pi}{h^3 c^2} \int_0^\infty \frac{A(E) E^2}{\exp\left(\frac{E}{k_B T}\right) - 1} dE,$$

and

$$f_r = \exp\left(\frac{-\Delta}{k_B T}\right),$$

which uses the offset defined above and represents the fraction of recombination due to radiative processes.

The voltage-dependent total current density is then multiplied by the voltage to give the power:

$$P = VJ = V \left\{ J_{\text{SC}} - J_0^{\text{rad}} \left[\exp\left(\frac{eV}{k_B T}\right) - 1 \right] \right\}.$$

The maximum value of this power, P_{max} , will be found at some balance of V and J , giving the optimal efficiency as

$$\eta = \frac{P_{\text{max}}}{P_{\text{in}}}$$

where the denominator is given by integrating over the AM1.5G spectrum, which has by convention been normalised to integrate to approximately 1000 W m^{-2} .

C. Optical absorption calculations

The SLME is calculated using the absorption profile of the material, which is directly accessed from electronic structure calculations, such as DFT. The reliability of the calculated SLME thus depends on the reliability of the underlying absorption profile. In general, more accurate DFT (or other electronic structure methods) calculations of the absorption profile require more computational resource to calculate. A simple hierarchy of electronic structure methods could include, in increasing accuracy and cost, the generalised gradient approximation²⁴ (GGA) to DFT, hybrid methods such as the Heyd–Scuseria–Ernzerhof²⁸ (HSE) functional, and GW routines. However, intermediate methods exist: particularly relevant to this study is the process of applying a scissor correction to a low-fidelity (e.g. GGA) spectrum using the difference in band gaps calculated at the low-fidelity and a higher-fidelity (e.g. HSE) level. We can consider applying a GGA \rightarrow HSE scissor correction,

$$\Delta E = E_g^{\text{HSE}} - E_g^{\text{GGA}},$$

to a GGA absorption spectrum as an approximation to an HSE-level absorption spectrum:

$$\alpha_{\text{HSE}} \approx \alpha_{\text{GGA}}(E - \Delta E),$$

which has the effect of shifting the spectrum to the right in most cases, since hybrid band gaps are usually larger than their GGA equivalents. This approach to approximating HSE spectra has been shown to be reasonable by Yang *et al.*,²⁹ who also showed that the independent particle approximation (IPA) to optics calculations produces spectra that generally agree well with those calculated using the more rigorous (and expensive) random-phase approximation (RPA). The fundamental band gaps needed for the scissor correction and offset are typically calculated using band structure calculations, while optics calculations provide both the dielectric tensor and the transition dipole matrix required for $\alpha(E)$ and E_g^{da} , respectively.

III. Results

In order to accelerate the identification of new materials with promising SLME values, one can propose replacing computationally demanding electronic structure calculations with surrogate models trained on existing data and capable of making predictions at a fraction of the cost; indeed, this is done quite routinely.^{30–32} We have trained surrogate models for each of the electronic structure steps in the SLME calculation workflow and now assess (i) how accurate these models are and (ii) how the errors in the model predictions propagate and affect the final ranking of new materials in the calculated SLME and (iii) the relative carbon cost of the different approaches.



Table 1 Table outlining the methods considered for estimating SLMEs, either directly or via the properties required to calculate SLMEs, with properties originating from DFT calculations at GGA and HSE level or ML predictions

Method	E_g	E_g^{da}	$\alpha(E)$	$A(E)$	ΔE	Δ	SLME
I	—	—	—	—	—	—	ML
II	—	—	ML	Py	—	ML	Py
III	GGA	GGA	GGA	Py	ML	Py	Py
IV	ML	—	—	Py	—	—	Py
V	HSE	—	—	Py	—	—	Py
VI	GGA	—	—	Py	ML	—	Py
VII	GGA	GGA	GGA	Py	—	Py	Py
VIII	HSE	GGA	GGA	Py	Py	Py	Py

In Table 1 we provide a list of potential workflows where electronic structure calculations are replaced with ML surrogates, where DFT calculations are indicated by their functional class, GGA or HSE. We have also provided a number of workflows that include step-function approaches and purely GGA-level properties to provide context for the accuracies and costs of the ML-based approaches. Note that in method II the model has been trained on scissor-corrected spectra, so a subsequent calculated or predicted scissor correction is not necessary. Those properties that are calculated using straightforward and negligibly expensive operations, in our case executed in Python, are represented as 'Py' in the table. For instance, step-function approaches use Python to estimate an absorbance profile from the material's band gap.

We use GGA-level band gaps for the offset calculation because this requires an optics calculation that inherently produces the data required for a GGA absorption spectrum, making it inefficient to predict one without the other. GGA offsets will introduce some error, though both gaps in the equation will be wrong by similar amounts, cancelling out some of this error. However, the test dataset used GGA-level offsets, so this source of error was not examined in this work.

Fig. 1 shows the relative cost of the calculations and predictions used in this work. Note that the area of a circle is proportional to the natural logarithm of its relative carbon cost, so the difference is even more stark than it appears. The negligible Python calculations are given as crosses to emphasise their low cost. ML inferences are also extremely inexpensive, though can be more meaningfully quantified as incurring a carbon cost around $1/2000\times$ that of a static GGA calculation, which is itself around an order of magnitude cheaper than a similar HSE calculation. In terms of energy, this single ML inference used around 1.9×10^{-3} Wh (around 7 J), which CodeCarbon²⁵ estimates as producing 4.5×10^{-4} g of CO₂: equivalent to driving a typical diesel transit bus 0.3 mm.³³

Finally, optics and band structure calculations are more expensive than static calculations, making an accurate absorption spectrum predicting model all the more promising. The figure does in some ways under-represent the cost of machine learning approaches, since training (and hyperparameter tuning, though this was not performed in this work) is not included. Training model 1 on the 4.8k dataset for 300 epochs




Fig. 1 Plot of the methods for estimating SLMEs outlined in Table 1, with crosses representing Python calculations and circles representing more costly calculations, with area $A \propto \ln\left(\frac{C}{C_{\min}}\right)$ for cost C . The absorbance column from Table 1 has been excluded for brevity.

was equivalent to 1.7×10^5 times that of a single inference. This is more indicative of how small the inference costs are than how large the training costs are. Moreover, these are one-off costs that would become negligible when considering the application of these models on vast datasets, and would not be incurred by future users of these models.

A. We haven't reached data saturation

Before analysing the effects of different ML interventions in the estimation of SLMEs, we first look at how accurate the various ML models are and how their performance scales with training data. It is a well-known phenomenon that the performance of deep learning models generally scales very well with data^{34–38} and therefore we investigate how the models we use scale with the available data.

Fig. 2(a) shows how the performance of the various ML models on a held-out test set evolves as the size of the training data increases. The dataset size is truncated at just under 5000: the number of materials in both the band gap and absorption spectra datasets (after a test/training split), since both are required to calculate a scissor-corrected SLME. From this plot, it is quite clear that all of the property models are still improving with more training data, and we have not reached data saturation. In the SI we show how the predicted absorption spectrum of GaAs (not in the training set) improves with more training data: in particular, point-to-point correlation is achieved at around 1k training data points, with the curve becoming smooth.



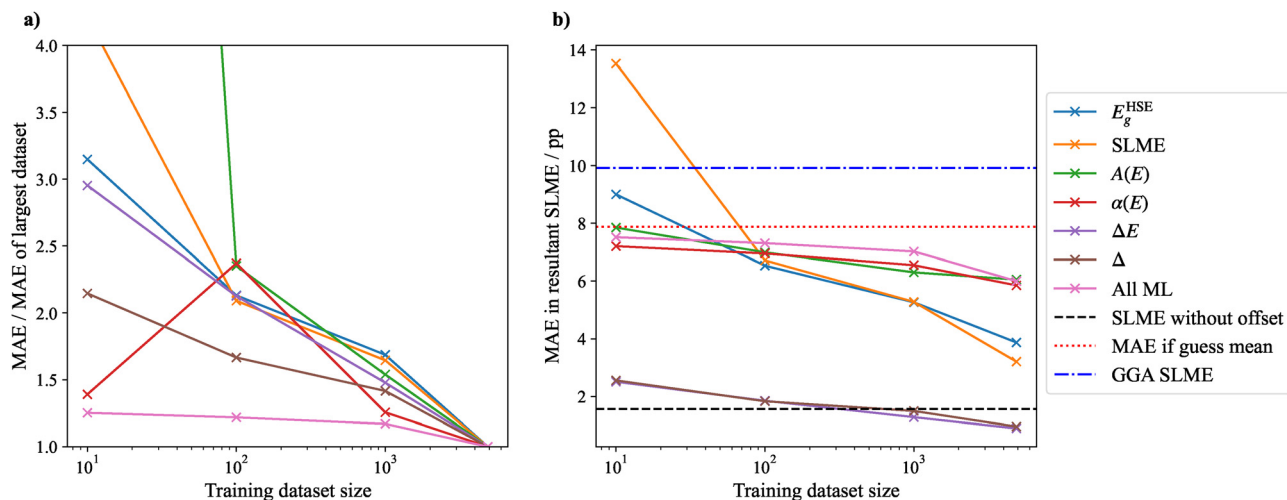


Fig. 2 Learning curves for each property, looking at (a) relative errors for the given property and (b) the resultant error when this learned property is used to calculate the SLME. For all but the full dataset, the values are averaged over three random sub-samples of the dataset. The errors are calculated using the test dataset used throughout the paper.

In the context of the final target (accurate SLMEs for PV screening), we show the effect of dataset size in Fig. 2(b). Here, the abscissa is the error in the final estimated SLME when a particular ML model is used in the workflow. The dotted red line shows a null hypothesis, where our “model” simply predicts the mean SLME of the training data. Clearly, with a few training data (≤ 100) all models exceed this baseline, even the model that predicts the high-dimensional absorption spectrum. The plot also demonstrates how with ~ 100 data points all workflows incorporating ML perform favourably when compared to calculating an SLME from a low-cost, low-fidelity DFT optical absorption profile obtained from a GGA calculation (without a scissor correction).

Perhaps more important than the absolute values in Fig. 2(b) are the gradients. The gradients give us an indication of how the predictions may be improved with additional data collection. The gradient of the direct prediction of SLME (with no DFT intermediates) shows the steepest gradient and extrapolation at the current rate of model improvement suggests that, with several tens of thousands of high-quality estimates of SLME, a model with negligible errors is possible.

If the absorption spectrum is known but the offset is not, Fig. 2 suggests that the inclusion of an ML-predicted offset is worthwhile (rather than a semi-SLME approach with $f_r = 0$), provided that the training dataset size exceeds $\sim 10^3$.

Predicting the absorption spectra and calculating the absorbance from them gives errors very similar to predicting the absorbance directly. This is perhaps surprising as absorbance spectra are naturally scaled to be between 0 and 1, and are relatively featureless (all more or less sigmoid-shaped), whereas absorption values may be anywhere between 0 and 10^7 , and the overall profile is generally more irregular. One possible explanation is that, when the absorbance is calculated from the predicted absorption spectrum, small discrepancies are smoothed out by the exponential function, thereby reducing




Fig. 3 Violin plot comparing the success of the methods outlined in Table 1 in recreating the test set's SLMEs, in terms of raw accuracy (LH axis) and ranking order when the materials are ranked by their SLME (RH axis). Note that the numerical difference is $\eta_{\text{pred}} - \eta_{\text{true}}$ so a positive difference is an overestimate.

the propagated error in the final SLME, whereas direct absorbance prediction has no such advantage. Given the better performance with the full training set, the methods that included spectral prediction predicted absorption rather than absorbance.

B. Some methods are more worth pursuing than others

Fig. 3 shows the performance of each method when predicting the SLMEs of the materials in the test set. This performance is quantified by the numerical difference between the target and predicted SLME on the left-hand axis, and the resultant difference in the ranking of these materials when sorted by their



SLMEs on the right. This latter distinction is arguably more important when filtering large databases for candidate materials. The difference between these performance quantifications is discussed in more depth in the following section, but we first consider the numerical accuracy.

Comparing the seven methods considered (method VIII is how the test set is calculated), we see some common trends. Scalar properties (SLME and scissor correction, methods I and III) are easier to predict than high-dimensional properties (the absorption spectrum as part of method II). Method II also suffers from the combination of errors, using predictions for the offset (by itself rather well predicted, see Fig. 2) and the absorption spectrum. This inaccuracy leads the step function-based approaches (methods IV–VI) to outperform method II. Otherwise, these approaches struggle compared to direct SLME prediction. Method V, wherein the band gap is calculated at the HSE level, does the best of these approaches, but the cost of this calculation is significantly higher than that of the ML inference in method I, as discussed in Section III D.

Finally, method VII, based on all GGA-level calculations (without any kind of scissor correction), is the poorest-performing approach. Interestingly, this approach gives the most clearly systematic error, with the vast majority being overestimates. GGA is known to underestimate band gaps due to the self-interaction error, so the absorption profiles will have an earlier offset, and thus we would expect larger short-circuit currents, but not necessarily larger efficiencies due to the voltage–current trade-off: smaller band gaps mean each excited carrier has less energy. We also see some systematic behaviour in method II, where SLME overestimates are limited to around 5 percentage points, while underestimates can be much more significant. The step function approaches also tend to overestimate SLMEs: this is likely because real absorptance spectra have more gradual onsets than step functions, especially for materials with indirect band gaps.

C. Better accuracy doesn't always give better screening

The ultimate goal in terms of PV materials discovery might be an accurate direct estimate of SLME from material structure and our previous analysis makes it clear that there is still significant room for improvement. We now consider how the errors in prediction accuracy relate to ranking errors and how this changes for different ML interventions.

We can see from Fig. 3 that different ML interventions introduce errors with different degrees of systematicity. This is a reminder that training objectives and benchmarks commonly used to compare ML models are not always appropriate for a given task.³⁹ More specifically related to ML for PV screening, this shows that trying to learn SLME directly is probably preferable to prediction of an absorption profile and using that to calculate the SLME. The direct SLME prediction is both more likely to improve with more data and gives more systematic errors. Any effort to generate more high-quality absorption profiles could be trivially translated to SLMEs; therefore, this is the most promising path for the screening of PV materials.

D. There will always be a cost-accuracy trade-off, but there are some sweet spots

Fig. 4(a) shows the 8 methods compared by their cost and accuracy, where the latter is measured as MAE from the held-out test set, so an ideal technique would be close to the origin. Three approaches (I, III, and VIII) stand out as significantly better than the rest, forming a roughly straight line when the cost is logarithmically scaled. Unsurprisingly, the machine learning approaches that predict scalar quantities are similar in cost but much more accurate than the approaches predicting high-dimensional spectra. However, the plot demonstrates the unsuitability of step-function approaches when there exist ML models that can predict SLMEs cheaply and more accurately. Method III, using a learned scissor correction alongside GGA-level absorption spectra and offsets, emerges as a viable intermediate method, nearly two orders of magnitude cheaper than fully hybrid calculations with MAEs under 1.0 percentage points. It should be emphasised that these errors are relative to the fully hybrid approach, which is itself limited. In the SI, we provide a plot of the method I model applied to a set of 29 high-efficiency materials, whose SLMEs were calculated using GW routines by Yu and Zunger²⁷ and were not in the training set; emphasising the need for high fidelity and large volumes of training data.

Another consideration when comparing the accuracy of different machine learning approaches is interpretability: the direct SLME prediction is rather a black box, where predicting the absorption spectrum and offset gives us better insight into why a given material is a good absorber. It also allows us to calculate properties like the short-circuit current and photo-voltage of a material, extending the applicability of this approach beyond traditional solar cells. Moreover, calculated SLMEs have the temperature, material thickness, and incoming radiation profile (typically the AM1.5G spectrum) implicit in their value, whereas predicting the spectra allows the user to alter these parameters for their application. This could be particularly useful when looking for materials for solar cells used on satellites or in indoor lighting. However, the distance of the Pareto front from the other 5 methods makes it hard to justify this approach. Method III is perhaps the best compromise between interpretability and accuracy.

There is also a large gulf in interpretability between all ML-based approaches and computational chemistry calculations. Even a static energy calculation provides a wealth of information compared to a single scalar from an ML model. This is an advantage of the computational methods that is difficult to quantify, but should be considered when deciding between methods. With this in mind, the scissor-correction approach, method III, is even more powerful, providing additional information (albeit at a GGA-level) compared to more ML-based approaches, while leveraging the low-cost, high-accuracy ML prediction of the scissor correction.

A final additional factor that could be considered is domain expertise. For instance, comparing V and VIII, we see that if a hybrid band gap is being calculated, it is only slightly more expensive to calculate the GGA absorption spectrum and offset



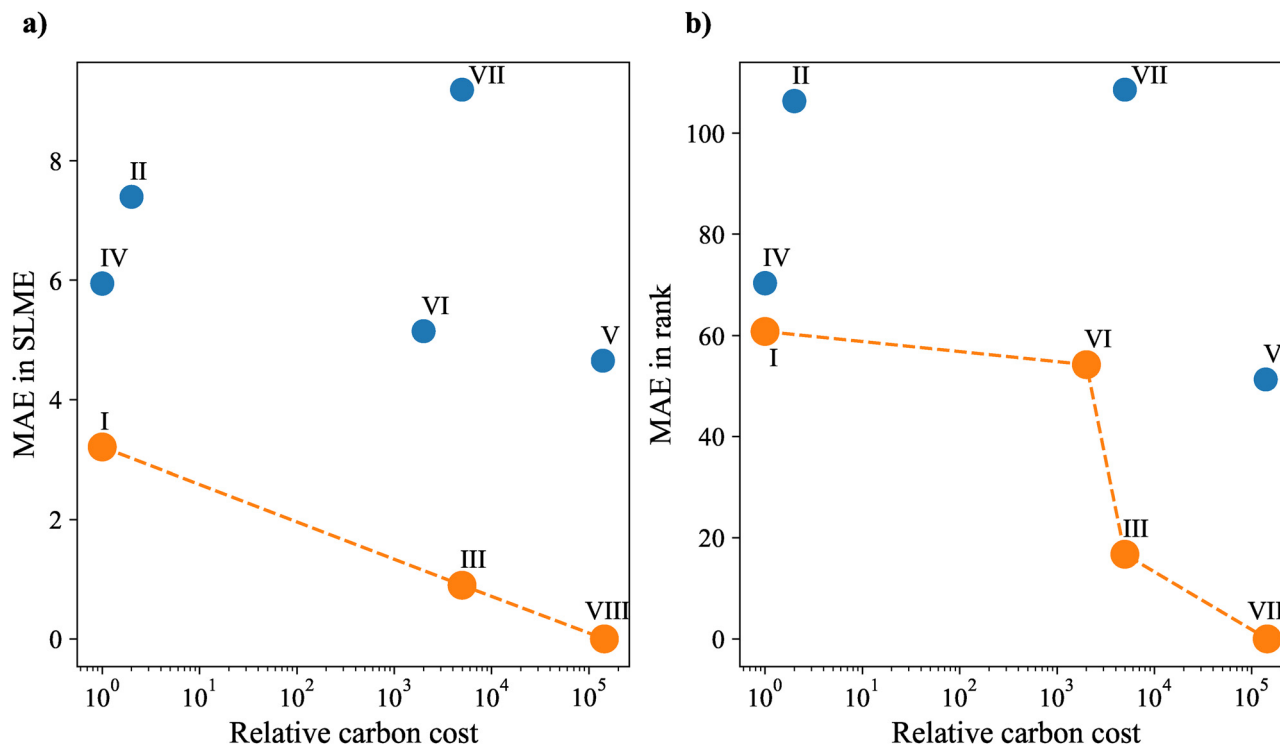


Fig. 4 Pareto front for performance vs. cost for the methods outlined in Table 1, where performance is measured as (a) MAE in SLME and (b) MAE in rank when the test set and the predictions are sorted by their SLMEs.

to facilitate an SLME rather than just detailed balance calculation. However, it requires the user to have experience with optics calculations in, for instance, DFT. Packages like Atomate2,⁴⁰ used for some example calculations in this report, make this very straightforward, while ML models like the atomistic line graph neural network (ALIGNN)⁴¹ used in this work are increasingly easy to use out-of-the-box.

Fig. 4(b) tells a similar story, although the difference in numerical accuracy and ranking accuracy is highlighted by method VI becoming part of the Pareto front. This seems to be a combination of method I being relatively poor at accurate ranking and method VI relatively good. However, VI is only narrowly better than I and is over 3 orders of magnitude more expensive, while III is much more accurate at less than $10\times$ the expense, making VI difficult to justify in most instances.

E. Machine learning isn't perfect – but neither is DFT

Next, we evaluate the performance of the direct SLME prediction model (method I) on an external test set drawn from the work of Fabini *et al.*,⁴³ which applies the Δ -sol correction scheme of Chan and Ceder⁴⁴ to GGA-level calculations. This test set is entirely independent from the training data used in this work, as it originates from a different set of DFT calculations and computational parameters. As such, it provides a robust test of how well our ML model generalises beyond the specific data distribution upon which it was trained (Fig. 5).

Unsurprisingly, the model's performance on this external dataset is somewhat worse than on the internal test set sampled from the same DFT workflow as the training data.




Fig. 5 Violin plot comparing the successes of the SLME-predicting ML model (method I) and the Choudhary *et al.*⁴² TB-mBJ dataset in reproducing the SLMEs of an external test set: the Fabini Δ -sol set.

This degradation is expected, as discrepancies between the DFT methodologies used to generate training and test labels introduce additional sources of error, which compound with those from the ML model itself. Nevertheless, the model maintains a reasonable ability to rank materials by predicted SLME, as shown in the rank correlation plots (Fig. 4).

To contextualise these errors, we also compared SLME values for the same materials computed using two different



DFT approaches: Δ -sol-corrected GGA (from Fabini *et al.*⁴³) and the Tran–Blaha modified Becke–Johnson (TB-mBJ) potential⁴⁵ (dataset from Choudhary *et al.*⁴²). Interestingly, the absolute and ranking errors between these two DFT methods are comparable in magnitude to the errors observed between the ML predictions and the Δ -sol data. For example, the mean absolute error (MAE) in SLME values between TB-mBJ and Δ -sol is 7.2 percentage points, *versus* 6.8 percentage points for the ML predictions; similarly, ranking errors are also of similar scale.

These results highlight two important conclusions. First, they demonstrate that the predictive performance of ML models trained on high-fidelity data can approach the level of variability introduced by changes in the DFT methodology itself. Second, they emphasise that the generation of consistent, high-fidelity SLME datasets remains a major bottleneck in data-driven PV discovery. For SLME prediction tasks, our findings suggest that investing in better-quality training data may yield greater improvements than simply expanding the size of existing datasets. In contrast, for absorption spectrum prediction—where model errors remain large even on consistent data—improvements in model architecture and training volume may be the more effective path forward.

F. Ways forward: more data, better data, or better models?

We now consider how we could push the boundary of efficient PV materials discovery. One promising avenue for improving ML predictions is through the development of models that more effectively capture the underlying structure–property relationships in the data. In the field of machine-learned interatomic potentials, for instance, the inclusion of physically motivated inductive biases—such as equivariance—has enabled models to achieve high accuracy with relatively modest training datasets.^{46–53} Meanwhile, recent attempts⁵⁴ to optimise the connectivity in the chemical graph used in a GNN have been shown to improve performance relative to the ALIGNN model used herein. Similar ideas are beginning to be explored in the prediction of optical properties.

While these recent efforts show encouraging progress, there are still important limitations. For example, two recent studies^{55,56} have proposed neural network (NN) models for predicting absorption spectra, both demonstrating reasonable accuracy. However, these models were trained and tested on more constrained datasets than those used in this work, and their performance may degrade when applied to more chemically diverse materials such as those in the W-R dataset. Grunert *et al.*⁵⁵ limited their materials to main-group elements from the first five rows of the periodic table, while Hung *et al.*⁵⁶ allowed a broader range of elements but restricted their dataset to structures with nine atoms or fewer per unit cell. Such constraints significantly reduce the overlap with the datasets used here, particularly where both band gap and spectrum data are needed. When trained on the dataset used in this work, the GNNOpt model from Hung *et al.*, based on the equivariant e3nn,^{46–50} predicts spectra that give better SLMEs than ALIGNN (see SI), but not enough to become a viable strategy, especially when the errors are confounded with those of the offsets.

This suggests that developments in model architectures such as these will continue to drive improved predictions of PV-relevant properties.

Another key challenge lies in the availability of consistent, high-quality training data. Both of the recently proposed neural network models for spectral prediction were trained on data generated using generalized gradient approximation (GGA) functionals, which—as we have shown—can lead to suboptimal screening performance. The reliance on GGA is largely driven by its relative abundance compared to more accurate methods, such as hybrid-DFT. However, progress in data infrastructure and learning techniques offers promising ways forward. Initiatives such as the novel materials discovery (NOMAD) program⁵⁷ and MPContribs (the platform for contributing to the Materials Project⁵⁸) are enabling the sharing of curated, high-quality computational datasets in line with FAIR data principles.⁵⁹

At the same time, recent advances in multi-fidelity machine learning^{60–62} allow models to be trained on datasets that combine varying levels of theoretical accuracy. By leveraging correlations between low- and high-fidelity data, these methods enable the use of larger training sets without sacrificing predictive reliability, thereby offering a practical route to more robust and generalizable ML models for materials discovery.

For traditional computational chemistry calculations, we note that plane-wave codes such as VASP^{63–65} are not the most efficient approach to hybrid DFT calculations. Atom-centred basis sets, such as those used in CRYSTAL⁶⁶ and CP2K,⁶⁷ are more efficient because many of the 4-electron Hartree–Fock integrals over real space decay rapidly, whereas the reciprocal space equivalents (used in plane-wave DFT) do not.⁶⁸ VASP was used in this study due to its widespread use in materials science (including for the generation of large datasets) and its ease of use *via* workflow managers like Atomate2⁴⁰ – we aim to replicate the most common workflows rather than necessarily the most efficient. However, we advise future hybrid DFT-based studies on photovoltaic materials to consider the more efficient atom-centred methods.

Another efficiency improvement could come from intermediate methods between GGA and HSE, such as those considered in a recent review by Janesko:⁶⁹ DFT+U, self-interaction corrections, localized orbital scaling corrections, local hybrid functionals, real-space nondynamical correlation, and their Rung-3.5 approach. Several of these approaches can approach hybrid accuracy at a fraction of the cost and are routinely used for systems where a full hybrid treatment would be prohibitively expensive.^{70,71} These methods have limitations of their own: DFT+U, for instance, requires optimisation on a case-by-case basis. The comparison between Δ -sol-corrected GGA and TB-mBJ in Section III E highlights the inconsistencies in these approaches.

Closely tied to the development of better data and models is the need for high-quality community benchmarks. As our results demonstrate, benchmarking efforts should not only assess predictive performance, but also account for the environmental cost of computation—such as carbon emissions—which can meaningfully influence the practicality of



different approaches. Evaluation choices fundamentally shape not only our measurements but also research priorities and scientific progress. Ensuring transparency and reproducibility in benchmarking is therefore critical. Recent proposals, such as evaluation cards, offer a structured means of documenting the assumptions, metrics, and limitations that underpin model assessments.^{39,72} By adopting such practices in the context of materials discovery, the community can move toward more robust, equitable, and environmentally conscious progress in the development of machine learning for photovoltaics and beyond.

A final consideration for improvement is the SLME metric itself. The Blank selection metric⁷³ has emerged as a more accurate computational characterisation of photovoltaic efficiency. However, it requires additional data such as the refractive index $n(E)$, of which there are currently no large datasets. A more rigorous computational study of a candidate photovoltaic would go even further, considering factors such as defects, dopants, and stability under real operating conditions. However, as a heuristic for filtering large areas of chemical space for intrinsically good PV absorbers, the SLME should be sufficient, hence its use in this work. As we have emphasised with the ranking plots, exact numbers for efficiency are less important than identifying the best materials.

IV. Conclusion

ML has the potential to dramatically accelerate the discovery of new PV materials. However, as we demonstrate here, it is not (currently at least) a panacea. Current limitations mean that for successful materials discovery campaigns for thin-film PV a combination of ML surrogates and electronic structure calculations is required. Our findings suggest that direct prediction of the SLME offers the most cost-effective approach to obtain reliable estimates of photovoltaic performance. Similarly, learned scissor corrections can substantially improve the accuracy of GGA absorption spectra at a fraction of the computational cost required for HSE band gap or optics calculations. However, direct spectral prediction currently introduces too much error to be practically useful for discovering novel photovoltaic materials, despite the appeal of the flexibility of this approach.

We have also identified clear pathways to improve ML surrogate models. Enhanced performance will likely require either substantially larger datasets of high-fidelity calculations than are presently available, or the implementation of transfer learning approaches that leverage extensive low-fidelity datasets alongside smaller, high-accuracy training sets.

More broadly, our study highlights the fundamental trade-off between computational cost and the efficacy of data-driven screening in materials design. We have outlined a blueprint for jointly evaluating the carbon cost and discovery performance of such campaigns. Embedding carbon cost reporting into computational discovery workflows is, we argue, a vital step toward ensuring that AI-, ML-, and simulation-driven approaches deliver truly beneficial and socially responsible innovation.

V. Methods

The atomistic line graph neural network (ALIGNN) model⁴¹ was used for ML predictions, with output dimensions of 1 or 100 for the scalar and spectral properties, respectively. Spectral data were represented by binning into 100 dimensions using the `numpy.interpolate()` function rather than compression into latent dimensions using, for example, a variational auto-encoder.⁷⁴ This decision was largely based on a paper from Kaundinya *et al.*⁷⁵ that used ALIGNN to predict the electron density of states of inorganic materials and found the binning approach (into 300 bins in their case) to be the more successful of the two.

Z-Score normalisation was used to scale labels for a more stable gradient descent; spectral properties were normalised per bin. Each model was trained for 300 epochs with a batch size of 64 and the rest of the hyperparameters in line with the model's original paper for consistency across the various properties predicted. A batch size of 2 was used for the learning curves as this enabled each dataset size to be trained with the same batch size.

Datasets from Woods-Robinson *et al.*,³² Kim *et al.*,⁷⁶ Fabiani *et al.*,⁴³ Yu and Zunger,²⁷ and Choudhary *et al.*⁴² were used, all accessed from freely available sources. The main dataset (the ~ 5.3 k overlapping materials from W-R and Kim) was split into an 80:10:10 ratio of training:validation:test data; the test materials were kept the same for all models for fairer comparison.

Some examples of DFT calculations at GGA and HSE levels were performed using the projected augmented wave (PAW) method^{77,78} within the Vienna *ab initio* Simulation Package (VASP),^{63–65} with CodeCarbon²⁵ monitoring the energy (and thus carbon) cost of each calculation. Atomate2⁴⁰ was used to generate the input files for these calculations, with structure files from the Materials Project,⁵⁸ to simulate a high-throughput workflow rather than bespoke calculations for each material. The raw numbers for these costs are available in the SI. CodeCarbon was also used for some ML training and inferences. Matplotlib was used for plotting.

Conflicts of interest

There are no conflicts to declare.

Data availability

The code and data used in this work are provided at https://github.com/mattheww98/PV_paper.git and <https://doi.org/10.5281/zenodo.17609133>.

The data supporting this article have been included as part of the supplementary information (SI). The supplementary information contains additional data to support the arguments made in the paper. These are, parity plots for Model I predictions *versus* various out of distribution test sets; the effect of training set size on predicted spectrum smoothness; full information on calculations time and carbon cost of DFT



calculations; a comparison of how even quite good predictions of spectra lead to large errors in derived SLMEs. See DOI: <https://doi.org/10.1039/d5mh01404b>.

Acknowledgements

We acknowledge support from EPSRC project EP/Y000552/1 and EP/Y014405/1. Via our membership of the UK's HEC Materials Chemistry Consortium, which is funded by EPSRC (EP/X035859/1), this work used the ARCHER2 UK National Supercomputing Service <https://www.archer2.ac.uk>. The authors acknowledge the use of the UCL Myriad High Performance Computing Facility (Myriad@UCL), and associated support services, in the completion of this work. We acknowledge Young HPC access which is partially funded by EPSRC (EP/T022213/1, EP/W032260/1 and EP/P020194/1). We acknowledge the assistance of the MPcontri team, in particular Patrick Huck for assisting us in assembling our initial datasets.

References

- 1 A. K. Cheetham, R. Seshadri and F. Wudl, *Nat. Synth.*, 2022, **1**, 514.
- 2 G. Masson, E. Bosch, A. Van Rechem and M. de l'Epine, *IEA Photovoltaic Power Systems Programme*, 2024, DOI: [10.69766/VHRF4040](https://doi.org/10.69766/VHRF4040).
- 3 N. M. Haegel, H. Atwater, T. Barnes, C. Breyer, A. Burrell, Y.-M. Chiang, S. De Wolf, B. Dimmler, D. Feldman, S. Glunz, J. C. Goldschmidt, D. Hochschild, R. Inzunza, I. Kaizuka, B. Kroposki, S. Kurtz, S. Leu, R. Margolis, K. Matsubara, A. Metz, W. K. Metzger, M. Morjaria, S. Niki, S. Nowak, I. M. Peters, S. Philipps, T. Reindl, A. Richter, D. Rose, K. Sakurai, R. Schlatmann, M. Shikano, W. Sinke, R. Sinton, B. Stanbery, M. Topic, W. Tumas, Y. Ueda, J. van de Lagemaat, P. Verlinden, M. Vetter, E. Warren, M. Werner, M. Yamaguchi and A. W. Bett, *Science*, 2019, **364**, 836.
- 4 J. C. Blakesley, R. S. Bonilla, M. Freitag, A. M. Ganose, N. Gasparini, P. Kaienburg, G. Koutsourakis, J. D. Major, J. Nelson, N. K. Noel, B. Roose, J. S. Yun, S. Aliwell, P. P. Altermatt, T. Ameri, V. Andrei, A. Armin, D. Bagnis, J. Baker, H. Beath, M. Bellanger, P. Berrouard, J. Blumberger, S. A. Boden, H. Bronstein, M. J. Carnie, C. Case, F. A. Castro, Y.-M. Chang, E. Chao, T. M. Clarke, G. Cooke, P. Docampo, K. Durose, J. R. Durrant, M. R. Filip, R. H. Friend, J. M. Frost, E. A. Gibson, A. J. Gillett, P. Goddard, S. N. Habisreutinger, M. Heeney, A. D. Hendsbee, L. C. Hirst, M. S. Islam, K. D. G. I. Jayawardena, M. B. Johnston, M. Kauer, J. Kettle, J.-S. Kim, D. Lamb, D. Lidzey, J. Lim, R. MacKenzie, N. Mason, I. McCulloch, K. P. McKenna, S. B. Meier, P. Meredith, G. Morse, J. D. Murphy, C. Nicklin, P. Ortega-Arriaga, T. Osterberg, J. B. Patel, A. Peaker, M. Riede, M. Rush, J. W. Ryan, D. O. Scanlon, P. J. Skabara, F. So, H. J. Snaith, L. Steier, J. Thiesbrummel, A. Troisi, C. Underwood, K. Walzer, T. Watson, J. M. Walls, A. Walsh, L. D. Whalley, B. Winchester, S. D. Stranks and R. L. Z. Hoyer, *J. Phys.: Energy*, 2024, **6**, 041501.
- 5 C. Battaglia, A. Cuevas and S. D. Wolf, *Energy Environ. Sci.*, 2016, **9**, 1552.
- 6 H. Sayed, A. M. Ahmed, A. Hajjiah, M. A. Abdelkawy and A. H. Aly, *Sci. Rep.*, 2025, **15**, 16529.
- 7 J. Ramanujam and U. P. Singh, *Energy Environ. Sci.*, 2017, **10**, 1306.
- 8 M. A. Scarpulla, B. McCandless, A. B. Phillips, Y. Yan, M. J. Heben, C. Wolden, G. Xiong, W. K. Metzger, D. Mao, D. Krasikov, I. Sankin, S. Grover, A. Munshi, W. Sampath, J. R. Sites, A. Bothwell, D. Albin, M. O. Reese, A. Romeo, M. Nardone, R. Klie, J. M. Walls, T. Fiducia, A. Abbas and S. M. Hayes, *Sol. Energy Mater. Sol. Cells*, 2023, **255**, 112289.
- 9 E. K. Solak and E. Irmak, *RSC Adv.*, 2023, **13**, 12244.
- 10 B. O'Regan and M. Grätzel, *Nature*, 1991, **353**, 737.
- 11 M. A. Green, A. Ho-Baillie and H. J. Snaith, *Nat. Photonics*, 2014, **8**, 506.
- 12 A. Zakutayev, J. D. Major, X. Hao, A. Walsh, J. Tang, T. K. Todorov, L. H. Wong and E. Saucedo, *J. Phys.: Energy*, 2021, **3**, 032003.
- 13 D. B. Needleman, J. R. Poindexter, R. C. Kurchin, I. M. Peters, G. Wilson and T. Buonassisi, *Energy Environ. Sci.*, 2016, **9**, 2122.
- 14 S. Y. Yang, J. Seidel, S. J. Byrnes, P. Shafer, C.-H. Yang, M. D. Rossell, P. Yu, Y.-H. Chu, J. F. Scott, J. W. Ager, L. W. Martin and R. Ramesh, *Nat. Nanotechnol.*, 2010, **5**, 143.
- 15 L. Wu and Y. Yang, *Adv. Mater. Interfaces*, 2022, **9**, 2201415.
- 16 A. W. Welch, L. L. Baranowski, H. Peng, H. Hempel, R. Eichberger, T. Unold, S. Lany, C. Wolden and A. Zakutayev, *Adv. Energy Mater.*, 2017, **7**, 1601935.
- 17 S. Yadav, R. K. Chauhan and R. Mishra, *Renewable Energy*, 2025, **255**, 123810.
- 18 A. A. Ahmad, A. B. Migdadi, A. M. Alsaad, I. A. Qattan, Q. M. Al-Bataineh and A. Telfah, *Heliyon*, 2022, **8**, e08683.
- 19 J. Vidal, S. Lany, M. d'Avezac, A. Zunger, A. Zakutayev, J. Francis and J. Tate, *Appl. Phys. Lett.*, 2012, **100**, 032104.
- 20 A. M. Ganose, K. T. Butler, A. Walsh and D. O. Scanlon, *J. Mater. Chem. A*, 2016, **4**, 2060.
- 21 X. Wang, S. R. Kavanagh, D. O. Scanlon and A. Walsh, *Joule*, 2024, **8**, 2105.
- 22 J.-H. Yang, W.-J. Yin, J.-S. Park, J. Ma and S.-H. Wei, *Semi-cond. Sci. Technol.*, 2016, **31**, 083002.
- 23 Z. Yuan, D. Dahliah, M. R. Hasan, G. Kassa, A. Pike, S. Quadir, R. Claes, C. Chandler, Y. Xiong, V. Kyverga, P. Yox, G.-M. Rignanese, I. Dabo, A. Zakutayev, D. P. Fenning, O. G. Reid, S. Bauers, J. Liu, K. Kovnir and G. Hautier, *Joule*, 2024, **8**, 1412.
- 24 J. P. Perdew, K. Burke and M. Ernzerhof, *Phys. Rev. Lett.*, 1996, **77**, 3865.
- 25 B. Courty, V. Schmidt, S. Luccioni, Goyal-Kamal, Marion-Coutarel, B. Feld, J. Lecourt, LiamConnell, A. Saboni, Inimaz, supatomic, M. Léval, L. Blanche, A. Cruveiller, ouminasara, F. Zhao, A. Joshi, A. Bogroff, H. D. Lavoireille, N. Laskaris, E. Abati, D. Blank, Z. Wang, A. Catovic, M. Alencon, Michał Stęchły, C. Bauer, L. O. N. D. Araújo, JPW, and MinervaBooks, mlco2/codecarbon: v2.4.1 (2024), version Number:v2.4.1.



- 26 W. Shockley and H. J. Queisser, *J. Appl. Phys.*, 1961, **32**, 510.
- 27 L. Yu and A. Zunger, *Phys. Rev. Lett.*, 2012, **108**, 068701.
- 28 J. Heyd, G. E. Scuseria and M. Ernzerhof, *J. Chem. Phys.*, 2003, **118**, 8207.
- 29 R. X. Yang, M. K. Horton, J. Munro and K. A. Persson, High-throughput optical absorption spectra for inorganic semiconductors, *arXiv*, 2022, preprint, arXiv:2209.02918, DOI: [10.48550/arXiv.2209.02918](https://doi.org/10.48550/arXiv.2209.02918).
- 30 F. De Angelis, *ACS Energy Lett.*, 2023, **8**, 1270.
- 31 M. D. Witman, A. Goyal, T. Ogitsu, A. H. McDaniel and S. Lany, *Nat. Comput. Sci.*, 2023, **3**, 675.
- 32 R. Woods-Robinson, Y. Xiong, J.-X. Shen, N. Winner, M. K. Horton, M. Asta, A. M. Ganose, G. Hautier and K. A. Persson, *Matter*, 2023, **6**, 3021.
- 33 T. W. Hesterberg, C. A. Lapin and W. B. Bunn, *Environ. Sci. Technol.*, 2008, **42**, 6437.
- 34 M. Banko and E. Brill, in Proceedings of the 39th Annual Meeting on Association for Computational Linguistics, ACL '01, Association for Computational Linguistics, USA, 2001, pp. 26–33.
- 35 D. Amodei, R. Anubhai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, J. Chen, M. Chrzanowski, A. Coates, G. Diamos, E. Elsen, J. Engel, L. Fan, C. Fougner, T. Han, A. Hannun, B. Jun, P. LeGresley, L. Lin, S. Narang, A. Ng, S. Ozair, R. Prenger, J. Raiman, S. Satheesh, D. Seetapun, S. Sengupta, Y. Wang, Z. Wang, C. Wang, B. Xiao, D. Yogatama, J. Zhan and Z. Zhu, *Deep Speech 2: End-to-End Speech Recognition in English and Mandarin*, 2015.
- 36 J. Hestness, S. Narang, N. Ardalani, G. Diamos, H. Jun, H. Kianinejad, M. M. A. Patwary, Y. Yang and Y. Zhou, *Deep Learning Scaling is Predictable, Empirically*, 2017.
- 37 C. Sun, A. Shrivastava, S. Singh and A. Gupta, Revisiting Unreasonable Effectiveness of Data in Deep Learning Era, *arXiv*, 2017, preprint, arXiv:1707.02968, DOI: [10.48550/arXiv.1707.02968](https://doi.org/10.48550/arXiv.1707.02968).
- 38 A. Bailly, C. Blanc, É. Francis, T. Guillotin, F. Jamal, B. Wakim and P. Roy, *Comput. Methods Programs Biomed.*, 2022, **213**, 106504.
- 39 N. Alampara, M. Schilling-Wilhelmi and K. M. Jablonka, Lessons from the trenches on evaluating machine-learning systems in materials science, *arXiv*, 2025, preprint, arXiv:2503.10837, DOI: [10.48550/arXiv.2503.10837](https://doi.org/10.48550/arXiv.2503.10837).
- 40 A. M. Ganose, H. Sahasrabudde, M. Asta, K. Beck, T. Biswas, A. Bonkowski, J. Bustamante, X. Chen, Y. Chiang, D. C. Chrzan, J. Clary, O. A. Cohen, C. Ertural, M. C. Gallant, J. George, S. Gerits, R. E. A. Goodall, R. D. Guha, G. Hautier, M. Horton, T. J. Inizan, A. D. Kaplan, R. S. Kingsbury, M. C. Kuner, B. Li, X. Linn, M. J. McDermott, R. S. Mohanakrishnan, A. N. Naik, J. B. Neaton, S. M. Parmar, K. A. Persson, G. Petretto, T. A. R. Purcell, F. Ricci, B. Rich, J. Riebesell, G.-M. Rignanese, A. S. Rosen, M. Scheffler, J. Schmidt, J.-X. Shen, A. Sobolev, R. Sundararaman, C. Tezak, V. Trinquet, J. B. Varley, D. Vigil-Fowler, D. Wang, D. Waroquiers, M. Wen, H. Yang, H. Zheng, J. Zheng, Z. Zhu and A. Jain, *Digital Discovery*, 2025, **4**, 1944.
- 41 K. Choudhary and B. DeCost, *npj Comput. Mater.*, 2021, **7**, 1.
- 42 K. Choudhary, Q. Zhang, A. C. E. Reid, S. Chowdhury, N. Van Nguyen, Z. Trautt, M. W. Newrock, F. Y. Congo and F. Tavazza, *Sci. Data*, 2018, **5**, 180082.
- 43 D. H. Fabini, M. Koerner and R. Seshadri, *Chem. Mater.*, 2019, **31**, 1561.
- 44 M. K. Y. Chan and G. Ceder, *Phys. Rev. Lett.*, 2010, **105**, 196403.
- 45 F. Tran and P. Blaha, *Phys. Rev. Lett.*, 2009, **102**, 226401.
- 46 M. Geiger and T. Smidt, e3nn: Euclidean Neural Networks, 2022.
- 47 M. Geiger, T. Smidt, A. M. B. K. Miller, W. Boomsma, B. Dice, K. Lapchevskiy, M. Weiler, M. Tyszkiewicz, S. Batzner, D. Madisetti, M. Uhrin, J. Frellsen, N. Jung, S. Sanborn, M. Wen, J. Rackers, M. Rød and M. Bailey, *Euclidean neural networks: e3nn*, 2022.
- 48 N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff and P. Riley, Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds, 2018, _eprint: 1802.08219.
- 49 M. Weiler, M. Geiger, M. Welling, W. Boomsma and T. Cohen, 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data, 2018, _eprint: 1807.02547.
- 50 R. Kondor, Z. Lin and S. Trivedi, Clebsch-Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network, 2018, _eprint: 1806.09231.
- 51 I. Batatia, S. Batzner, D. P. Kovács, A. Musaelian, G. N. C. Simm, R. Drautz, C. Ortner, B. Kozinsky and G. Csányi, *Nat. Mach. Intell.*, 2025, **7**, 56.
- 52 S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt and B. Kozinsky, *Nat. Commun.*, 2022, **13**, 2453.
- 53 A. Musaelian, S. Batzner, A. Johansson, L. Sun, C. J. Owen, M. Kornbluth and B. Kozinsky, *Nat. Commun.*, 2023, **14**, 579.
- 54 R. Ruff, P. Reiser, J. Stühmer and P. Friederich, *Digital Discovery*, 2024, **3**, 594.
- 55 M. Grunert, M. Großmann and E. Runge, *Phys. Rev. Mater.*, 2024, **8**, L122201.
- 56 N. T. Hung, R. Okabe, A. Chottrattanapituk and M. Li, *Adv. Mater.*, 2024, **36**, 2409175.
- 57 L. Sbailò, Á. Fekete, L. M. Ghiringhelli and M. Scheffler, *npj Comput. Mater.*, 2022, **8**, 250.
- 58 A. 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 Mater.*, 2013, **1**, 011002.
- 59 M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, J. Bouwman, A. J. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C. T. Evelo, R. Finkers, A. Gonzalez-Beltran, A. J. Gray, P. Groth, C. Goble, J. S. Grethe, J. Heringa, P. A. 't Hoen, R. Hooft, T. Kuhn, R. Kok, J. Kok, S. J. Lusher, M. E. Martone, A. Mons, A. L. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S.-A. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M. A. Swertz, M. Thompson, J. van der Lei, E. van Mulligen, J. Velterop,



- A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao and B. Mons, *Sci. Data*, 2016, **3**, 160018.
- 60 C. Fare, P. Fenner, M. Benatan, A. Varsi and E. O. Pyzer-Knapp, *npj Comput. Mater.*, 2022, **8**, 257.
- 61 N. Hoffmann, J. Schmidt, S. Botti and M. A. L. Marques, *Digital Discovery*, 2023, **2**, 1368.
- 62 H. Kaur, F. D. Pia, I. Batatia, X. R. Advincula, B. X. Shi, J. Lan, G. Csányi, A. Michaelides and V. Kapil, *Faraday Discuss.*, 2025, **256**, 120.
- 63 G. Kresse and J. Hafner, *Phys. Rev. B:Condens. Matter Mater. Phys.*, 1993, **47**, 558.
- 64 G. Kresse and J. Furthmüller, *Comput. Mater. Sci.*, 1996, **6**, 15.
- 65 G. Kresse and J. Furthmüller, *Phys. Rev. B:Condens. Matter Mater. Phys.*, 1996, **54**, 11169.
- 66 A. Erba, J. K. Desmarais, S. Casassa, B. Civalleri, L. Donà, I. J. Bush, B. Searle, L. Maschio, L. Edith-Daga, A. Cossard, C. Ribaldone, E. Ascrizzi, N. L. Marana, J.-P. Flament and B. Kirtman, *J. Chem. Theory Comput.*, 2023, **19**, 6891.
- 67 T. D. Kühne, M. Iannuzzi, M. Del Ben, V. V. Rybkin, P. Seewald, F. Stein, T. Laino, R. Z. Khaliullin, O. Schütt, F. Schiffmann, D. Golze, J. Wilhelm, S. Chulkov, M. H. Bani-Hashemian, V. Weber, U. Borštnik, M. Taillefumier, A. S. Jakobovits, A. Lazzaro, H. Pabst, T. Müller, R. Schade, M. Guidon, S. Andermatt, N. Holmberg, G. K. Schenter, A. Hehn, A. Bussy, F. Belleflamme, G. Tabacchi, A. Glöß, M. Lass, I. Bethune, C. J. Mundy, C. Plessl, M. Watkins, J. VandeVondele, M. Krack and J. Hutter, *J. Chem. Phys.*, 2020, **152**, 194103, _eprint: https://pubs.aip.org/aip/jcp/articlepdf/doi/10.1063/5.0007045/16718133/194103_1_online.pdf.
- 68 K. Adhikari, A. Chakrabarty, O. Bouhali, N. Mousseau, C. S. Becquart and F. El-Mellouhi, *J. Comput. Sci.*, 2018, **29**, 163.
- 69 B. G. Janesko, *Chem. Soc. Rev.*, 2021, **50**, 8470.
- 70 H. Zeng, D. Wang, C.-J. Yang, C.-L. Dong, W. Lin, X. Sang, B. Yang, Z. Li, S. Yao, Q. Zhang, J. Lu, L. Lei, Y. Li, R. D. Rodriguez and Y. Hou, *ACS Catal.*, 2025, **15**, 9610.
- 71 R. Shinde, S. S. R. K. C. Yamijala and B. M. Wong, *J. Phys.: Condens. Matter*, 2020, **33**, 115501.
- 72 M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji and T. Gebru, 2019, pp. 220–229.
- 73 B. Blank, T. Kirchartz, S. Lany and U. Rau, *Phys. Rev. Appl.*, 2017, **8**, 024032.
- 74 D. P. Kingma and M. Welling, Auto-Encoding Variational Bayes, *arXiv*, 2022, preprint, arXiv:1312.6114, DOI: [10.48550/arXiv.1312.6114](https://doi.org/10.48550/arXiv.1312.6114).
- 75 P. R. Kaundinya, K. Choudhary and S. R. Kalidindi, *JOM*, 2022, **74**, 1395.
- 76 S. Kim, M. Lee, C. Hong, Y. Yoon, H. An, D. Lee, W. Jeong, D. Yoo, Y. Kang, Y. Youn and S. Han, *Sci. Data*, 2020, **7**, 387.
- 77 G. Kresse and J. Hafner, *J. Phys.: Condens. Matter*, 1994, **6**, 8245.
- 78 G. Kresse and D. Joubert, *Phys. Rev. B:Condens. Matter Mater. Phys.*, 1999, **59**, 1758.

