

Predicting Lattice Parameters from Atomic-Scale Images of Two Dimensional Materials Using Deep Learning

Sayak Chakrabarty¹, Kamal Choudhary^{2,3,4}, Daniel Wines², Youjia Li⁵, Vishu Gupta^{5,6,7},
Muhammed Nur Talha Kilic⁵, Alok Choudhary⁵, and Ankit Agrawal*⁵

¹*Department of Computer Science, Northwestern University, Evanston, IL 60208, USA*

²*Material Measurement Laboratory, National Institute of Standards and Technology,
Gaithersburg, MD 20899, USA*

³*Department of Materials Science and Engineering, Whiting School of Engineering, The
Johns Hopkins University, Baltimore, Maryland 21218, USA*

⁴*Department of Electrical and Computer Engineering, Whiting School of Engineering, The
Johns Hopkins University, Baltimore, Maryland 21218, USA*

⁵*Department of Electrical and Computer Engineering, Northwestern University, Evanston,
IL 60208, USA*

⁶*Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, NJ 08544,
USA*

⁷*Cancer Research Institute and Institute of Integrative Genomics, Princeton University, NJ
08544, USA*

April 17, 2026

Abstract

Determining lattice parameters in two-dimensional (2D) materials is essential for materials characterization and discovery. In this work, we propose a deep-learning-driven pipeline that addresses the regression task of estimating the lattice constants a and b and the angle γ directly from 2D images using computer vision. We evaluate our approach on three different 2D material datasets: JARVIS-2D (JV2D) and Computational 2D Materials Database (C2DB), and a newly created dataset derived from the Alexandria database. Multiple architectures are compared, including DenseNet121, Vision Transformers (ViT-L/14) paired with Multi Layer Perceptron (MLP) heads, and GoogleNet. DenseNet121 achieves accurate performance, with mean absolute errors as low as 0.18 Å for the Alexandria-based dataset and 0.17 Å for C2DB and 0.59 Å for JV2D, as well as up to 96 % accuracy in classifying Bravais lattice types for the Alexandria-based dataset.

*Corresponding author

1 Introduction

Two-dimensional (2D) crystalline layers with atomic thickness have emerged as a transformative class of materials over the past two decades. Notably, the discovery of graphene demonstrated the remarkable properties of one-atom-thick structures, motivating extensive exploration of similarly layered crystals such as hexagonal boron nitride (h-BN) and transition metal dichalcogenides (e.g., MoS₂) [1, 2, 3]. Many 2D materials exhibit exceptional characteristics, including high electrical conductivity, unique optical responses, and mechanical robustness, underscoring their significance for next-generation technologies in optoelectronics, energy storage, catalysis, and beyond [4, 5].

Accurately determining lattice parameters (i.e., lattice constants and angles) in 2D materials is crucial for both fundamental studies and applied design. Conventional methods often involve time-intensive, manual analyses of atomic-scale images obtained through scanning transmission electron microscopy (STEM) or other atomistic imaging techniques, such as scanning tunneling microscopy (STM) and atomic force microscopy (AFM) [6]. A few examples of experimentally recorded STEM images are provided in the supplementary information as can be seen in Figure S7.

Due to rapid advancements in computer vision [7], automated approaches to analyzing these images have become increasingly appealing. Microscopy data, whether from experimental or computational sources, is now more abundant than ever [8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23]. However, a unified library that centralizes dataset collection, generation, and data-analytic methods remains in high demand, especially for materials characterization.

Such libraries would prove invaluable for various microscopy image tasks: classification [11, 24, 25, 26], pixelwise learning (e.g., semantic segmentation) [27, 28, 29, 30, 31, 32, 33, 34], object recognition, localization, and super-resolution [35, 36, 37, 38, 39]. Their applicability also extends to multiple scientific domains such as materials science, condensed matter physics, and biology [9, 12, 40].

On the computational front, various frameworks exist for simulating STM and STEM images. STM images can be generated using the Bardeen [41], Tersoff-Hamman [42], or Chen [43] approaches, produced in constant-height or constant-current modes. High-Angle Annular Dark-Field (HAADF)-STEM images can be obtained through convolution, Bloch wave, or multislice approximations [44, 45, 46, 47, 48], with the convolution approximation being particularly efficient for thin films. Though more computationally demanding, Bloch wave and multislice methods offer broader generalizability.

Several notable libraries focus on experimental and computational microscopy images, including abTEM [49], EXtraction, Separation, and Caption-based natural Language Annotation of IMages [26], AtomAI [50], Prismatic [51], QSTEM [52], and AtomVision: A machine vision library for atomistic images [53]. Deep learning (DL), especially convolutional neural networks (CNNs), has become a standard approach for analyzing atomistic image data [9]. Numerous investigations have applied DL methods to tasks including pixelwise atom detection in simulated atomic-resolution TEM images of graphene [28] or WS₂ [29], dopant and vacancy identification in STEM images of WSe₂ [30], and automatic defect detection (e.g., dislocations, precipitates, voids) in steel [31]. Ziatdinov *et al.* [33] also employed DL to infer surface molecular structures from microscopy data.

Recent advancements further underscore deep learning’s benefits in various microscopy contexts. Segment Anything for Microscopy [54] uses vision-transformer models for image segmentation, while AI-integrated scanning probe microscopy [55] accelerates atomic-level characterization. Jacobs *et al.* [56] discussed semantic segmentation methods for detecting defects in TEM images, and Lobato *et al.* [57] demonstrated the efficacy of CNNs for enhancing single-shot electron microscopy images. Deep-learning-based noise filtering has enabled millisecond-order STEM imaging [58], and real-time defect tracking was showcased in DefectTrack [59]. Other efforts include automated dislocation loop analysis [60], STM-based atomic identification [61], and the integration of large vision-language models [62] to extend microscopy data interpretation. Single-molecule identification in zeolite structures with deep-learning-enabled STEM [63] further highlights these methods’ versatility. Moreover, large language models (LLMs) have begun to play a substantial role in automating and unifying multimodal microscopy datasets [64, 65, 66, 67], with Mandal *et al.* [68] demonstrating an LLM-based agent for automating AFM experiments. Lei *et al.* [67] provided a broader overview of LLMs’ transformative potential across materials science.

As large-scale, high-resolution data becomes more readily available, there is an increasing need for automated, robust, and scalable methods to extract structural parameters from the images, classify lattice types, and predict lattice constants. Recent research [69] has been done where neural networks have been trained on basic atomic descriptors to classify ABO₃ crystal structures and regress lattice constants, reporting 88 % classification accuracy and strong lattice-parameter regression. One other related study used modified Hirshfeld surface fingerprints of inorganic crystals as structure-aware inputs to a deep network to predict lattice parameters of cubic perovskites [70], while another study introduced a composition-only random forest model with new descriptors to predict a, b, c (and angles) across many space groups (an R^2 close to 0.97 was achieved for cubic crystals) [71].

We propose a deep-learning-based pipeline capable of learning from diverse 2D material images and accurately predicting key lattice parameters, namely, a , b , and γ . To the best of our knowledge, similar studies have not been done before and there are no state-of-the-art models in the literature that perform this task. In our classification task, the labels are the following: We have label 0 for hexagonal (hp), label 1 for square/tetragonal (tp), label 2 for rectangle/orthorhombic (op), label 3 for rhombus/centered orthorhombic (oc), and label 4 for parallelogram/monoclinic (mp). Our work aims to streamline the pipeline (as shown in Figure 1) from raw atomic-scale imaging data to actionable structural insights.

Recent developments in crystal-graph attention neural networks have yielded notable advances in predicting thermodynamic stability by leveraging high-quality datasets with balanced chemical and crystal-symmetry coverage [72]. Through a systematic symmetry-based approach, researchers have also uncovered a vast range of novel binary and ternary two-dimensional (2D) materials, employing combinatorial engines that populate Wyckoff positions with chemical elements to generate previously uncharted structures [73]. In another effort, crystal graph attention models circumvent the need for precise atomic positions by embedding graph distances, enabling both composition- and prototype-based high-throughput searches of stable crystal structures [74].

These examples illustrate the critical importance of large-scale, high-quality DFT datasets for materials-science machine learning. The Computational 2D Materials Database (C2DB) systematically catalogs electronic, magnetic, optical and structural properties of monolayer

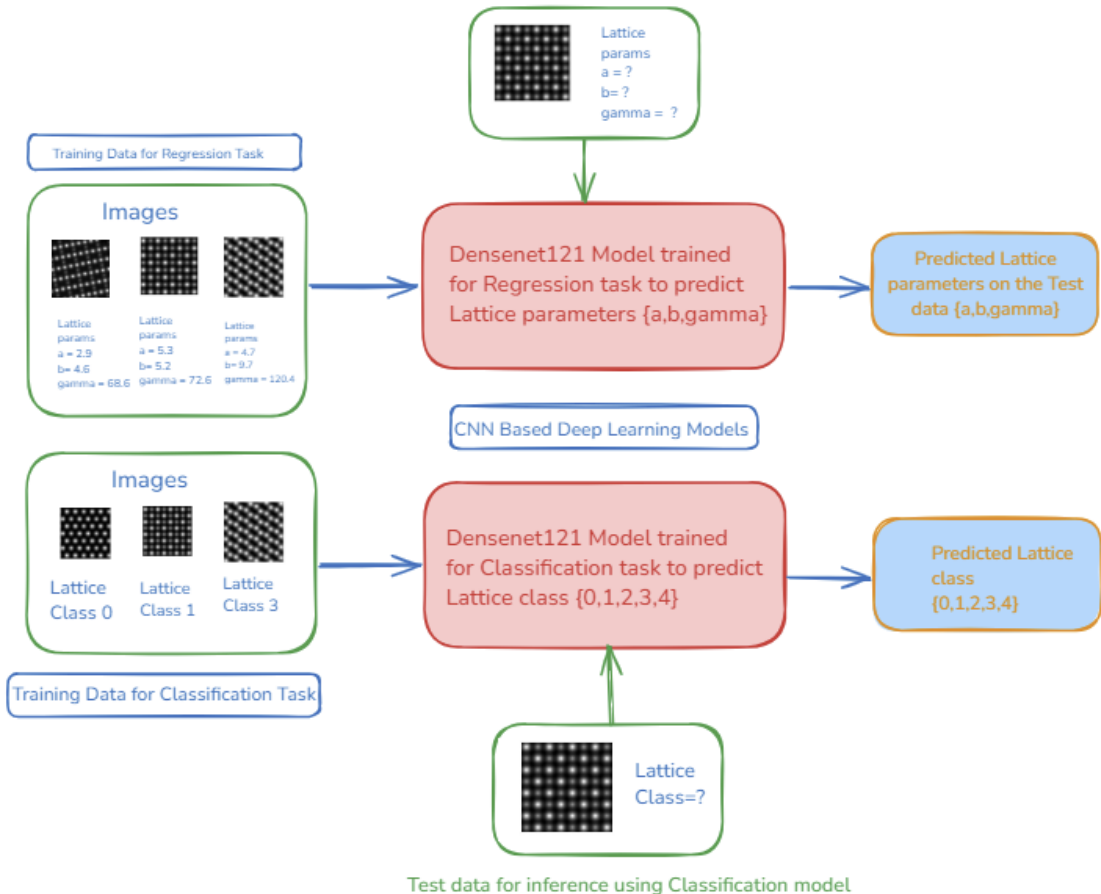


Figure 1: An overview of our workflow where we trained two Densenet models, each with 121 layers. One model is for the regression task to predict the lattice parameters a , b , γ based on the image, and the other Densenet model is trained to predict the lattice class based on the image. The lattice classes can be hexagonal (0), square (1), rectangular (2), rhombus (3), parallelogram (4).

materials to identify spintronic and optoelectronic candidates [75, 76], while JARVIS-2D (JV2D) provides high-throughput DFT data on thousands of 2D monolayers—including formation energies, band gaps and exfoliation energies—to support screening, benchmarking and discovery efforts [77, 78, 79, 80].

Similarly, the recently developed “Alexandria” database provides an extensive collection of DFT calculations for 1D, 2D, and 3D compounds, supporting accurate property predictions and robust universal interatomic potentials [81]. We extracted a subset of two-dimensional (2D) materials from the Alexandria database, specifically those computed using the Perdew–Burke–Ernzerhof (PBE) [82] exchange–correlation functional. This subset was then converted into simulated scanning transmission electron microscopy (STEM) images using the JARVIS-Tools [77, 78, 79] codebase. We refer to this collection of images as

the ‘‘Alex-PBE’’ dataset throughout this paper. To handle the large volume of images, we employed parallelized data processing techniques, ensuring efficient and consistent image generation for our experiments.

We apply a DenseNet121 model for both regression and classification tasks on three 2D materials datasets (JV2D, C2DB, and Alex-PBE). In regression, the model predicts lattice parameters with promising accuracy (low mean absolute errors) on all three datasets. For classification, we train on the Alex-PBE dataset and achieve high accuracy overall. This demonstrates the model’s effectiveness in classifying diverse lattice types and underscores the potential of deep learning approaches in identifying and characterizing 2D crystal structures. Our training and inference pipeline is shown in Figure 1. All datasets, code, and trained models are publicly available to facilitate reproducibility and future research.

2 Results and Discussion

2.1 Dataset generation and curation

The STEM images [44, 46, 47, 48, 49] were produced using the convolution approximation, which relies on fast Fourier transform-based convolutions, as described by the equation:

$$I(r) = R(r, Z) \otimes PSF(r) \tag{1}$$

where \mathbf{r} represents a 2D vector in the image plane, and $I(r)$ denotes the image intensity, and PSF represents the microscope’s point spread function. The transmission function, $R(r, Z)$, is given by:

$$R(r, Z) = \sum_i^N Z_i^{1.7} \delta(r - r_i) \tag{2}$$

Here, $R(r, Z)$ represents the transmission function corresponding to the N atoms located at positions r_i . This function incorporates information about the atomic potential of the system, which is determined by the atomic number Z_i . According to Rutherford scattering from the nuclear charge, the intensity follows a Z^2 dependence. However, due to core electron screening and variations in detection collection angles, this exponent is typically reduced. A value of 1.7 has been chosen as an approximate compromise among these factors. Previous studies have used power values ranging between 1.3 and 1.7 [83] to better align with experimental results. However, for consistency, we adopt a fixed exponent of 1.7 across all systems.

It is important to highlight that only crystallographic coordinates and atom-type information are required for generating these STEM images. In our experiments, the power factor is 1.7, the Gaussian width is set to be 0.5, the Lorentzian width is set to be 0.5, and the intensity ratio is set at 0.5. The optimized geometries of the 2D materials were obtained from density functional theory (DFT) [84] calculations. The generated STEM images have an output resolution of 256×256 pixels, covering an area of at least $2.5 \text{ nm} \times 2.5 \text{ nm}$. The microscope’s point spread function (PSF) is represented by a normalized Gaussian function with a width of 0.5 \AA .

The 2D materials used in this study were obtained from the JARVIS-DFT-2D (JV2D) [77, 78, 79], Computational 2D Materials Database (C2DB) [75], and Alexandria [73, 72, 82] databases. All image datasets (JARVIS-2D, C2DB, Alex-PBE) were regenerated directly from their underlying DFT databases [53]. In our current work, we generated a new image dataset based on Alex-PBE. From each Alex-PBE relaxed structure, we generated STEM images by creating a supercell of the primitive cell, cropping a central square region for clear visualization, and then normalizing the resulting image to ensure consistent intensity scaling. The simulated STEM images are sufficiently representative of experimentally measured images, which can be seen from Ref. [53], Figure 3, where the comparison for computational and previously reported experimental STM and STEM images are shown for a few examples (i.e., graphene).

For each dataset, we split the data into 80 % training and 20 % testing partitions while preserving a balanced representation of each lattice class. In the Atom-Vision [53] paper, the dataset was split into only train and test. To be consistent with previous work, we have split JV2D, C2DB, and Alex-PBE datasets into only train-test. Note that we did not combine the 3 datasets due to the potential for data contamination. These 3 DFT datasets were all computed with slightly different methodologies. Specifically, these DFT relaxations were performed with different codes, different exchange correlation functionals, different pseudopotentials, and slightly different simulation settings. This can cause inconsistencies in the training/testing, especially if we have duplicate structures that are slightly different each other. This is even more of an issue for structures where a and b might be switched, structures that are oriented different from one another, and if there are supercells versus primitive cells in one database versus another. To avoid this issue, we treated each dataset separately.

From Supplementary Information (SI) Figure S3, we see that the training loss curve flattens out at the end of 50 epochs for the JV2D and C2DB datasets and at the end of 500 epochs for the Alex-PBE dataset.

2.2 How We Chose Our Machine Learning Model

In this study, we systematically evaluate three distinct deep learning approaches for classification tasks across two datasets, as shown in SI Table S1. The approaches encompass feature extraction using pretrained Vision Transformer (ViT-L/14) architectures, as well as training Densenet121 and GoogleNet models from scratch. The goal is to investigate each method’s effectiveness, measure comparative performances, and highlight practical considerations.

In the first approach, we leverage the pre-trained CLIP model employing the ViT-L/14 architecture to extract meaningful, high-dimensional feature vectors (each vector having 768 coordinates) from images. These extracted feature vectors serve as input to a Multi-Layer Perceptron (MLP) classifier. This two-step approach, which separates feature extraction from the classifier training stage, is particularly valuable as it reduces computational overhead and enables efficient experimentation with different downstream classifiers.

Specifically, the ViT-L/14 architecture generates image representations as feature vectors of length 768. These vectors, along with corresponding 2D Bravais Lattice types (class labels), are subsequently utilized for training MLP classifiers, tailored individually based on dataset scale. To ensure optimal performance across datasets of varying sizes, we configured the MLP

with two hidden layers sized (512, 256) for the medium-to-large datasets, training the MLP model for 400 epochs. For the smallest dataset, the hidden layer dimensions were adjusted to (256, 128), with the number of epochs reduced to 200 to avoid potential overfitting. The Adam optimizer was selected for optimization, and the cross-entropy loss function guided the training process.

Our second approach involves the direct training of a Densenet121 model from scratch, without utilizing pretrained weights. The DenseNet architecture, characterized by its densely connected convolutional layers, offers strong gradient flow and efficient parameter utilization, making it highly suitable for various image-classification tasks. To match the number of target classes for each specific dataset, we modified the final fully connected classification layer accordingly. The training process involved careful hyperparameter tuning: the model was trained for 50 epochs, using a batch size of 16. Training employed the Adam optimizer, configured with a learning rate of 0.001, and utilized the cross-entropy loss function.

To ensure consistent and effective feature extraction by the convolutional layers, we standardized image inputs through a carefully structured preprocessing pipeline prior to training. This preprocessing included resizing each image to 224×224 pixels, converting image data to PyTorch [85] tensors, and normalizing tensors with mean and standard deviation values derived from the dataset. Such standardization facilitates smoother convergence during training and reduces model sensitivity to variance in raw input data distributions.

The image preprocessing pipeline for the Densenet Model is as follows. Images are uniformly resized to dimensions of 224×224 pixels to meet model input specifications. The resized images are then converted into tensor representations to facilitate efficient computations on GPUs. Each channel (Red, Green, Blue) is independently normalized utilizing established ImageNet [86] mean ($\mu = [0.485, 0.456, 0.406]$) and standard deviation ($\sigma = [0.229, 0.224, 0.225]$) values. This process standardizes the data distribution, enhancing training stability and generalization capability.

Our third approach mirrors the preprocessing steps of the Densenet121 method but employs the GoogleNet architecture trained from scratch, without leveraging pretrained weights. The GoogleNet architecture, known for its inception modules and efficient utilization of computational resources, provides another meaningful comparison point in evaluating our classification tasks.

Consistent with our Densenet121 configuration, GoogleNet training involved 50 epochs, a batch size of 16, and the Adam optimizer with a learning rate of 0.001. To further regularize training and prevent overfitting, we applied an $L2$ weight decay parameter (λ_2), which penalizes large weight magnitudes and encourages the model towards simpler solutions. The choice of GoogleNet extends our evaluation to include an architecture designed explicitly for computational efficiency.

Based on the classification results in SI Table S1, we picked the Densenet121 model to predict the lattice parameters for three datasets. We adjust the last layer so that we can switch from the classification to the regression problem. The Train-Test Distributions are reported in SI Table S3 for each dataset on which we conducted our experiments. We have used the other models, namely the GoogleNet and Clip+MLP on the Alex-PBE dataset for the regression experiments of lattice parameters. The results can be found in SI Figure S1 and SI Table S4.

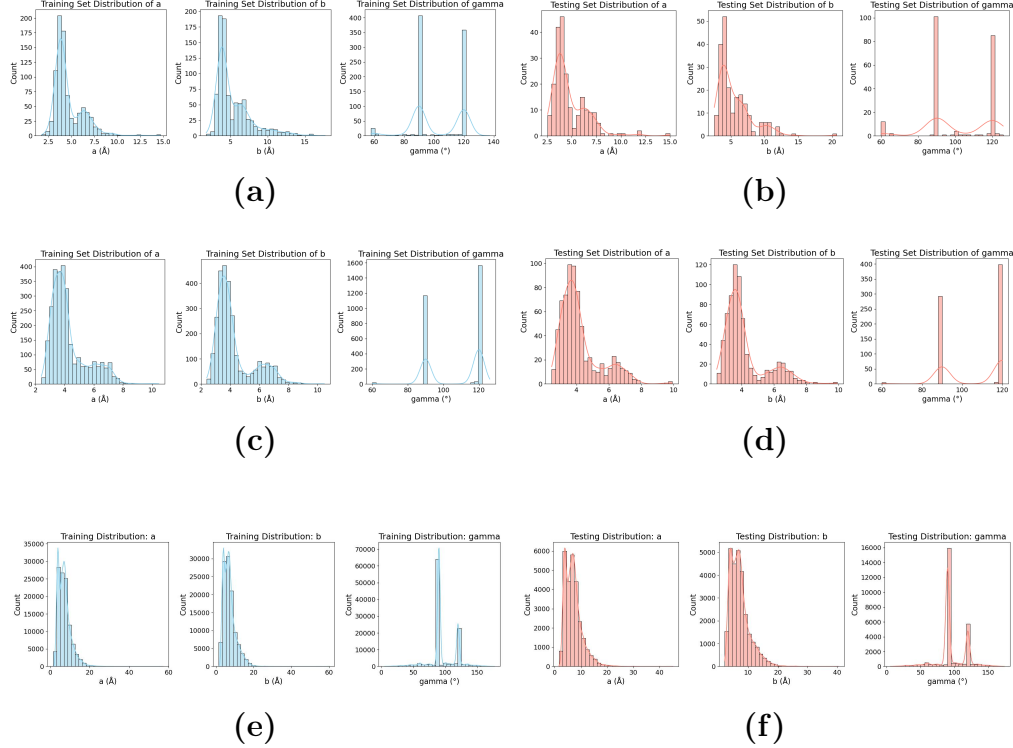


Figure 2: Train and test distribution plots of lattice parameters a , b , and γ for the datasets used in our experiments. The x-axis in each panel represents a , b (in Å), and γ (in °), while the y-axis shows the count. (a) JV2D Train dataset, (b) JV2D Test dataset, (c) C2DB Train dataset, (d) C2DB Test dataset, (e) Alex-PBE Train dataset, (f) Alex-PBE Test dataset.

We selected the DenseNet121 model architecture for predicting lattice parameters across multiple datasets. Specifically, we adapted the final fully connected layer of the DenseNet121 network, originally designed for classification tasks, to accommodate regression outputs. The regression task involves predicting three lattice parameters: a , b , and γ . To rigorously evaluate model performance, we utilized standard train-test splits for each dataset. The datasets and corresponding sizes of the training and testing sets are detailed as follows:

- **STEM-JV2D:** 859 training samples and 215 test samples.
- **STEM-C2DB:** 2815 training samples and 705 test samples.
- **Alex-PBE:** 110,263 training samples and 27,569 test samples.

We evaluate the developed models using two metrics that are most common for regression tasks; they are Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). The error metrics computed are as follows:

- MAE = $\frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|$
- MSE = $\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2$ and RMSE = $\sqrt{\text{MSE}}$

Here $y_i = (a_i, b_i, \gamma_i)$ are the true values for sample i , $\hat{y}_i = (\hat{a}_i, \hat{b}_i, \hat{\gamma}_i)$ are the predicted values for sample i and N is the total number of test samples.

For evaluating classification tasks, we use the standard classification metrics.

- **Precision** = $\frac{\sum_{i=1}^N \mathbf{1}\{y_i = 1 \wedge \hat{y}_i = 1\}}{\sum_{i=1}^N \mathbf{1}\{\hat{y}_i = 1\}}$ where $\mathbf{1}\{\cdot\}$ is the indicator function (equals 1 if its argument is true, 0 otherwise), $y_i \in \{0, 1\}$ is the true label for sample i , and $\hat{y}_i \in \{0, 1\}$ is the predicted label.
- **Recall** = $\frac{\sum_{i=1}^N \mathbf{1}\{y_i = 1 \wedge \hat{y}_i = 1\}}{\sum_{i=1}^N \mathbf{1}\{y_i = 1\}}$
- **F1 Score** = $2 \cdot \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$

Baseline Model

To establish a reference point for evaluating the effectiveness of our DenseNet121-based model, we employed a simple yet robust baseline approach. The baseline model predicts the mean value of each parameter computed from the training set, denoted as:

$$\bar{y}_p = \frac{1}{N_{\text{train}}} \sum_{i=1}^{N_{\text{train}}} y_{i,p}, \quad p \in \{a, b, \gamma\}$$

where N_{train} represents the total number of training samples. Subsequently, these mean values were utilized as constant predictions for each parameter across all test samples:

$$\hat{y}_{i,p} = \bar{y}_p, \quad \forall i \in \{1, 2, \dots, N_{\text{test}}\}$$

Employing the mean-based baseline serves as a fundamental test to validate that the trained DenseNet121 model effectively captures meaningful patterns in the data. If the trained model's predictions fail to surpass the baseline performance, this suggests the model is insufficiently leveraging the available data for accurate regression.

2.3 Training Methodology and Experiments

A key consideration during our experimental design was whether to train a single DenseNet121 model to simultaneously predict all three lattice parameters (a, b, γ) , or to train three separate instances of the DenseNet121 model, each independently tasked with predicting only one of the parameters. Training separate models would eliminate inter-parameter correlations from influencing predictions, whereas training a single multi-output model might enable the model to leverage these correlations effectively.

To systematically evaluate both scenarios, we conducted preliminary experiments on the STEM-JV2D dataset. Specifically, we trained both the single multi-output DenseNet121 model and three separate single-output DenseNet121 models under identical training conditions: 50 epochs, a batch size of 16, a learning rate of 0.0005, and optimization via the Adam optimizer without pretrained weights. Performance across both configurations was compared using the previously defined regression metrics.

Based on comparative performance analysis in SI Table S2, we observed that the single multi-output DenseNet121 model consistently achieved better accuracy across all regression metrics. Consequently, we selected this approach (a single DenseNet121 model predicting all three parameters simultaneously) as our primary model architecture moving forward.

Following this methodological decision, comprehensive modelling experiments were carried out on the aforementioned three datasets (STEM-JV2D, STEM-C2DB, and Alex-PBE). We train our model on the full training set and then test on four subsets of the testing set, apart from testing on the full test set, for each parameter a , b , and γ . The 4 test subsets are:

- $a \leq 10 \text{ \AA}$, $b \leq 10 \text{ \AA}$, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$.
- a, b has no restrictions and $\gamma \notin \{60^\circ, 90^\circ, 120^\circ\}$
- $a > 10 \text{ \AA}$, b, γ have no restrictions
- $b > 10 \text{ \AA}$, a, γ have no restrictions

Out of the four subsets, the first subset is the most abundantly occurring and thus most relevant and interesting. Large databases of two-dimensional crystals indicate that *most 2D materials have in-plane lattice constants on the order of only a few angstroms*. For example, a high-throughput study identified approximately 1500 monolayers across more than 30 crystal prototypes, with the vast majority exhibiting lattice lengths well below 1 nm [75]. Common 2D materials such as graphene, h-BN, and monolayer MoS₂ have a (and b) values around 2.5 to 3.3 Å, while heavier or more complex 2D materials (e.g., antimonene, bismuthene) show in-plane constants on the order of 4 to 4.5 Å [87]. Extensive surveys of exfoliable layered compounds confirm that a and b rarely exceed 10 Å for stable monolayers [75]. Hence, nearly all known 2D crystals have relatively compact unit cells in-plane. Regarding unit-cell angles, 2D materials *predominantly adopt high-symmetry lattice angles of 60°, 90°, or 120°*. Crystallographic data show that *hexagonal/trigonal* (with $\gamma = 60^\circ$ or equivalently 120°) and *rectangular/square* ($\gamma = 90^\circ$) lattices are by far the most common in monolayers [75]. For example, graphene and MoS₂ have a hexagonal (120°) arrangement, while black phosphorus (phosphorene) is orthorhombic ($\gamma = 90^\circ$). Statistics from large databases confirm that over 70% of predicted 2D compounds fall into one of these high-symmetry categories [75], whereas lower-symmetry (oblique) lattices with arbitrary γ are relatively rare in stable 2D materials. This is why the first subset, where $a, b \leq 10 \text{ \AA}$, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$ should be the subset where we expect the best results. The same pattern is also observed from the data distribution of the Alex-PBE dataset as seen in Figure 2.

After conducting our modeling experiments as shown in Figure 1, we report our regression results across three datasets in Table 1. The scatter plots are shown in Figure 3 for each of the datasets when the model is run on the subset of the test set where $a, b \leq 10 \text{ \AA}$,

$\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$. It is important to note that in Fig. 3c, we observe two additional slopes that deviate from $y = x$ (perfect prediction, red dotted line) on the parity plot for a and b . This can be due in part to the large number of parallelogram/monoclinic structures in the Alex-PBE data (see SI Table S3), many of which represent supercells used for magnetic calculations. We report our classification experimental results across the 3 datasets in Table 2.

The SI depicts the histogram of errors for the test subsets across our three datasets. The histograms for the JV2D dataset are in SI Figure S4, the C2DB dataset in SI Figure S5, and the Alex-PBE dataset in SI Figure S6.

In Figure 3, it is clear that angle γ takes discrete values in the set $\{69, 90, 120\}$. To offer an alternative analysis of Figure 3 (γ), we generated confusion matrix plots (shown in SI Figure S2).

We analyzed how predictions vary by chemistry (chalcogenides vs halides) and by electronic properties (metal vs semiconductor/insulator). We examined subsets of the test set of each model with the following criteria: metals have $E_g \leq 0.05$ eV, semiconductors/insulators have $E_g > 0.05$ eV. In addition, we looked at chalcogenides (S/Se/Te), halides (F/Cl/Br/I), and transition metal dichalcogenides (TMDs) with MX_2 stoichiometry. For each subset of the test set, we report mean absolute error (MAE) and root mean squared error (RMSE) for lattice parameters a , b and angle γ . The results for this are reported in SI Table S5. We observe relatively consistent trends between each model. In particular, we see that the error is lower for metals than semiconductors/insulators for all three models. In addition, we observe the lowest error for TMD materials for each model. For more accurate models such as C2DB, we see less variation in predictions across material classes (errors ranging from 0.07 Å to 0.25 Å for lattice constants), while for models with higher overall error (JV2D), we see a larger variation with errors up to 1.20 Å for lattice constant of chalcogenides.

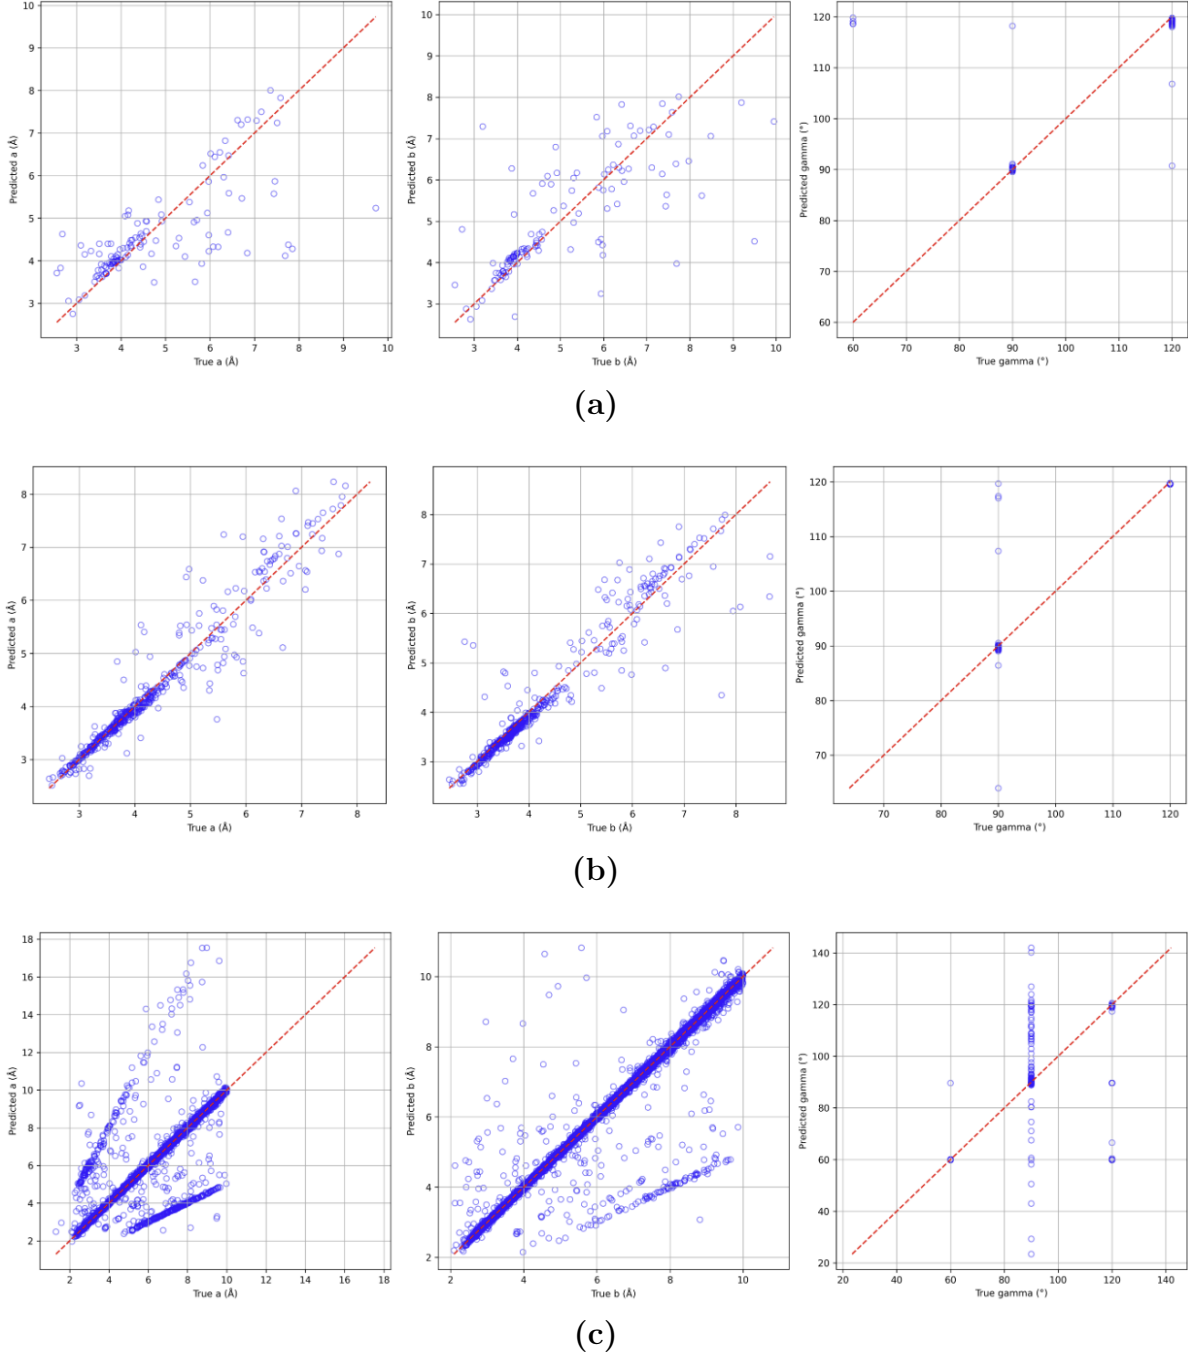


Figure 3: The Densenet model for prediction of the lattice parameters a, b, γ across the three datasets, (a) JV2D, (b) C2DB and (c) Alex-PBE. We show the scatter plots when the trained model is evaluated on only the subset of the test set where $a, b \leq 10 \text{ \AA}$, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$. In each scatter plot, we have three plots showing the scatter plot for a, b, γ where we have shown True vs Predicted values.

Table 1: Regression results for the lattice parameters a, b, γ using the Densenet model across JV2D, C2DB, and Alex-PBE datasets. Shown are the full test set, the (most abundant) subset where we have $a \leq 10 \text{ \AA}$, $b \leq 10 \text{ \AA}$, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$, and the baseline. n represents the number of samples in the set.

Dataset	Test Subset	Parameter	MAE	RMSE
JV2D	Full Test Set (n=215)	a (Å)	0.69	1.31
		b (Å)	0.99	1.91
		γ (°)	5.21	14.44
	$a \leq 10 \text{ \AA}$, $b \leq 10 \text{ \AA}$, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$ (n=126)	a (Å)	0.59	0.99
		b (Å)	0.64	1.07
		γ (°)	2.89	11.24
	Baseline (Mean)	a (Å)	1.29	1.74
		b (Å)	1.87	2.52
		γ (°)	16.02	17.74
C2DB	Full Test Set (n=705)	a (Å)	0.17	0.37
		b (Å)	0.19	0.42
		γ (°)	0.86	4.82
	$a \leq 10 \text{ \AA}$, $b \leq 10 \text{ \AA}$, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$ (n=481)	a (Å)	0.17	0.32
		b (Å)	0.19	0.40
		γ (°)	0.49	2.65
	Baseline	a (Å)	0.95	1.24
		b (Å)	1.00	1.29
		γ (°)	14.81	15.22
Alex-PBE	Full Test Set (n=27569)	a (Å)	0.35	1.16
		b (Å)	0.23	0.83
		γ (°)	0.90	4.34
	$a \leq 10 \text{ \AA}$, $b \leq 10 \text{ \AA}$, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$ (n=12139)	a (Å)	0.18	0.79
		b (Å)	0.09	0.43
		γ (°)	0.57	2.64
	Baseline	a (Å)	3.81	5.02
		b (Å)	4.23	5.35
		γ (°)	13.73	18.77

Table 2: Classification result of the Densenet model trained and tested separately on three different datasets: JV2D, C2DB, and Alex-PBE. We report the precision, recall, F1-score, and size for each class.

Dataset	Class	Precision	Recall	F1-score	Support
Alex-PBE	0	0.97	0.99	0.98	5715
	1	0.95	0.94	0.95	1792
	2	0.97	0.98	0.97	14579
	3	0.92	0.90	0.91	3513
	4	0.89	0.80	0.84	1970
	Accuracy		0.96		27569
	Macro avg	0.94	0.92	0.93	27569
	Weighted avg	0.96	0.96	0.96	27569
JV2D	0	0.97	0.90	0.93	92
	1	0.87	0.97	0.92	35
	2	0.78	0.88	0.83	65
	3	0.50	0.29	0.36	7
	4	0.77	0.62	0.69	16
	Accuracy		0.87		215
	Macro avg	0.78	0.73	0.75	215
	Weighted avg	0.86	0.87	0.86	215
C2DB	0	0.89	0.96	0.92	316
	1	0.99	0.98	0.98	148
	2	0.97	0.98	0.97	144
	3	0.57	0.46	0.51	63
	4	0.59	0.38	0.46	34
	Accuracy		0.89		705
	Macro avg	0.80	0.75	0.77	705
	Weighted avg	0.88	0.89	0.89	705

Through these experiments, we gain several important insights:

- *Restricted to STEM Data:* It is important to note that the current trained models can only be applied to STEM data. However, it is possible to design similar strategies to extend this work to other types of characterization data. We restricted our dataset to only STEM [44, 46, 47, 48, 49] images because we could easily generate a large quantity of STEM data from relaxed structures in DFT databases [53, 84].
- *Effect of Typical Lattice Ranges:* The best performance is consistently observed for lattices within typical ranges ($a, b \leq 10 \text{ \AA}$ and $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$). This aligns with the compact dimensions and high-symmetry angles common to naturally occurring 2D materials, enhancing prediction accuracy due to the higher representation of these cases in the training data. If we look at the results obtained in the estimation of the angle using the dataset Alex-PBE, it might seem that the model is not good enough in

sketching a linear trend between the predicted and actual angles. This problem has been also reported by other authors [88, 89] in the context of powder XRD.

- *Challenges with Out-of-Range and Low-Symmetry Parameters:* Model accuracy declines for lattices with large constants or non-standard angles. Some of these anomalies in experiment can include distorted phases (e.g., 1T'-MoS₂ [90]), asymmetric Janus materials (e.g., MoSSe [91]), and charge-density-wave distorted monolayers (e.g., 1T-TaS₂ [92]). This degradation likely arises from the limited presence of these atypical structures in the training dataset.
- *The Densenet Model is not robust to rotations.* Only translation equivariance is built-in for the Densenet model [93]. Convolution slides a fixed kernel over the image plane, so shifting an image shifts the feature map in the same way. Rotating the input, however, does not rotate the filter grid; pixels that used to fall under one weight now fall under a very different pattern. The standard Conv2d operator [94] is defined on a Cartesian grid. Rotation is not a symmetry of that grid. After a 90° turn, edge orientations, corners, etc., land in positions no filter has ever seen, so early activations collapse, and the prediction head gets arbitrary features. This can potentially become an issue for structures where a and b might be switched, or when the model looks at supercells versus primitive cells.
- We have included experimental 2D STEM images taken from literature (shown in SI Figure S7). The figures we have included are the following: a) tetragonal FeTe (modified from Ref. [95]), b) hexagonal FeTe (modified from Ref. [95]), c) graphene (modified from Ref. [96]) and d) graphene (modified from Ref. [97]). The images taken from literature were all larger than 2.5 nm x 2.5 nm. To match the simulated data on which these models were trained, we cropped out a region of the experimental image to be as close to 2.5 nm x 2.5 nm as possible and when necessary, we rotated to align its dimensions and crystallographic orientation with those of our simulated dataset. This preprocessing step is essential because the DenseNet architecture employed is not rotation-invariant and is sensitive to mismatched resolutions, which would otherwise degrade inference accuracy.

Subsequently, the three trained models were applied to these experimental images to infer lattice parameters. Ground-truth values were obtained from Ref. [98] for graphene, Ref. [99, 100] for tetragonal FeTe, and Ref. [95, 100] for hexagonal FeTe. The ground-truth and predicted parameters are reported in SI Table S6. The Alex-PBE model exhibited the largest errors, while the C2DB and JV2D models do an excellent job of predicting γ for both graphene images and both phases of FeTe. The C2DB model performs exceptionally well for the prediction of the graphene lattice constants and angle, while the JV2D model has more difficulty with the simultaneous prediction of a and b for graphene. However, prediction results for both phases of FeTe are better with the JV2D model. It is important to note that for the tetragonal and hexagonal phases of FeTe, there are more discrepancies in what the experimentally measured lattice constants are. For example, values ranging from 3.7 Å to 3.9 Å have been experimentally reported for tetragonal FeTe thin films [99, 100] and values ranging from 3.8 Å to 4.1 Å

have been reported for hexagonal FeTe thin films [95, 100]. It has also been suggested that magnetic ordering can impact these lattice constant values [95]. For comparison, the DFT computed lattice constant values from C2DB (which assumes a ferromagnetic ground state) are 3.62 Å for tetragonal FeTe (C2DB ID: Fe2Te2-b49eca087529) and 3.99 Å for hexagonal FeTe (C2DB ID: Fe2Te2-18809f787f88).

- *Benefits of Multi-Output Learning:* Training a single DenseNet121 model for simultaneous multi-parameter prediction improves overall accuracy compared to separate single-output models. Leveraging the correlations among lattice parameters (a , b , and γ) through multi-task learning is particularly advantageous for crystalline property prediction tasks.

3 Conclusion

In this work, we developed an effective and scalable deep-learning pipeline for automatically extracting lattice parameters, predicting lattice constants using regression techniques, and classifying lattice types from STEM images of 2D materials. Using the DenseNet121 model, we achieved particularly strong performance on a test subset where $a \leq 10$ Å, $b \leq 10$ Å, and $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$. For the JV2D dataset, we obtained MAE score of around 0.69 Å for a , and 0.99 Å for b on the full test set. On the C2DB dataset, the model yielded MAE of around 0.17 Å for a , and 0.19 Å for b for the full test set. On the Alex-PBE dataset (our largest dataset contribution and publicly available), MAE scores were around 0.35 Å for a , and 0.23 Å for b for the full test set. For classification (trained on the full Alex-PBE train dataset), we achieved up to 96 % accuracy on this largest dataset. Finally, we acknowledge that our model is not suitable when faced with challenges with out-of-range and low-symmetry parameters. Future work could integrate rotation and scale-equivariant architectures or physics-informed losses to extend robustness to low-symmetry lattices and larger lattice-parameter ranges, as well as to accommodate out-of-distribution experimental images. The trained models and processing scripts will be publicly available, facilitating reproducibility and fostering further research into automated lattice characterization.

4 Supporting Information

- Comparison of 2D Bravais lattice classification performance of three model architectures
- Predictive performance of a single multi-task DenseNet versus three independent DenseNet models trained separately for each parameter.
- The number of train and test samples for each of the three datasets
- Scatter plots created when the GoogleNet model and the Clip + MLP model are used for prediction of the lattice parameters a, b, γ on the Alex-PBE dataset.
- Table representing the comparison of regression results for lattice parameters a, b, γ on the Alex-PBE dataset using GoogLeNet and ViT+MLP

- Scatter plots and confusion matrices when the Densenet model is used for prediction of the lattice parameters a, b, γ across the three datasets, where γ predictions are represented as confusion matrix plots.
- Training loss curves for various models
- Table showing prediction of the lattice parameters a, b, γ when the datasets are broken into subsets based on chemical compositions of materials.
- Histogram of prediction errors for lattice parameters after training the DenseNet model on the JV2D dataset
- Histogram of prediction errors for lattice parameters after training the DenseNet model on the C2DB dataset
- Histogram of prediction errors for lattice parameters after training the DenseNet model on the Alex-PBE dataset
- Experimental 2D STEM images taken from literature and used for model inference.
- Comparison of predicted lattice parameters (a, b, γ) for selected materials from experiment (see Fig. S7).

Acknowledgments

This work was performed under the following financial assistance award 70NANB24H136 from U.S. Department of Commerce, National Institute of Standards and Technology as part of the Center for Hierarchical Materials Design (CHiMaD). Partial support is also acknowledged from NSF awards CMMI-2053929, OAC-2331329, and the Northwestern Center for Nanocombinatorics.

This work was performed with funding from the CHIPS Metrology Program, part of CHIPS for America, National Institute of Standards and Technology, U.S. Department of Commerce. Certain commercial equipment, instruments, software, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identifications are not intended to imply recommendation or endorsement by NIST, nor are they intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

Code and Data Availability

The code and data supporting the findings of this study will be available upon publication.

Conflicts of interest

There are no conflicts to declare.

References

- [1] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, and A. A. Firsov. Electric field effect in atomically thin carbon films. *Science*, 306(5696):666–669, 2004.
- [2] Andre K. Geim and Konstantin S. Novoselov. The rise of graphene. *Nature Materials*, 6(3):183–191, 2007.
- [3] Kin Fai Mak, Changgu Lee, James Hone, Jie Shan, and Tony F. Heinz. Atomically thin mos_2 : a new direct-gap semiconductor. *Physical Review Letters*, 105(13):136805, 2010.
- [4] Likai Li, Feng Yang, Gonglan Ye, Zhi Zhang, Zengwei Zhu, Wenjing Lou, Xiaoyan Zhou, Liang Li, Kenji Watanabe, Takashi Taniguchi, et al. Direct observation of the layer-dependent electronic structure in phosphorene. *Nature Nanotechnology*, 11(7):593–597, 2016.
- [5] Yujie Zhang, Ziliang Ye, Yuki Matsushashi, and Yoshihiro Iwasa. Direct observation of the transition from indirect to direct bandgap in atomically thin epitaxial mose_2 . *Nano Letters*, 14(5):2419–2425, 2010.
- [6] Peter J Goodhew, John Humphreys, and Richard Beanland. *Electron microscopy and analysis*. CRC press, 2000.
- [7] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014.
- [8] Elizabeth A Holm, Ryan Cohn, Nan Gao, Andrew R Kitahara, Thomas P Matson, Bo Lei, and Srujana Rao Yarasi. Overview: Computer vision and machine learning for microstructural characterization and analysis. *Metallurgical and Materials Transactions A*, 51(12):5985–5999, 2020.
- [9] Kamal Choudhary, Brian DeCost, Chi Chen, Anubhav Jain, Francesca Tavazza, Ryan Cohn, Cheol Woo Park, Alok Choudhary, Ankit Agrawal, Simon JL Billinge, et al. Recent advances and applications of deep learning methods in materials science. *npj Computational Materials*, 8(1):1–26, 2022.
- [10] Rama K Vasudevan, Kamal Choudhary, Apurva Mehta, Ryan Smith, Gilad Kusne, Francesca Tavazza, Lukas Vlcek, Maxim Ziatdinov, Sergei V Kalinin, and Jason Hattrick-Simpers. Materials science in the artificial intelligence age: high-throughput library generation, machine learning, and a pathway from correlations to the underpinning physics. *MRS communications*, 9(3):821–838, 2019.
- [11] Mohammad Hadi Modarres, Rossella Aversa, Stefano Cozzini, Regina Ciancio, Angelo Leto, and Giuseppe Piero Brandino. Neural network for nanoscience scanning electron microscope image recognition. *Scientific reports*, 7(1):1–12, 2017.

- [12] Mengshu Ge, Fei Su, Zhicheng Zhao, and Dong Su. Deep learning analysis on microscopic imaging in materials science. *Materials Today Nano*, 11:100087, 2020.
- [13] Michiel Larmuseau, Michael Sluydts, Koenraad Theuwissen, Lode Duprez, Tom Dhaene, and Stefaan Cottenier. Compact representations of microstructure images using triplet networks. *npj Computational Materials*, 6(1):1–11, 2020.
- [14] Zijiang Yang, Xiaolin Li, L Catherine Brinson, Alok N Choudhary, Wei Chen, and Ankit Agrawal. Microstructural materials design via deep adversarial learning methodology. *Journal of Mechanical Design*, 140(11), 2018.
- [15] Tim Hsu, William K Epting, Hokon Kim, Harry W Abernathy, Gregory A Hackett, Anthony D Rollett, Paul A Salvador, and Elizabeth A Holm. Microstructure generation via generative adversarial network for heterogeneous, topologically complex 3d materials. *JOM*, 73(1):90–102, 2021.
- [16] Sehyun Chun, Sidhartha Roy, Yen Thi Nguyen, Joseph B Choi, HS Udaykumar, and Stephen S Baek. Deep learning for synthetic microstructure generation in a materials-by-design framework for heterogeneous energetic materials. *Scientific reports*, 10(1):1–15, 2020.
- [17] Minyi Dai, Mehmet F Demirel, Yingyu Liang, and Jia-Mian Hu. Graph neural networks for an accurate and interpretable prediction of the properties of polycrystalline materials. *npj Computational Materials*, 7(1):1–9, 2021.
- [18] Tiberiu Stan, Zachary T Thompson, and Peter W Voorhees. Optimizing convolutional neural networks to perform semantic segmentation on large materials imaging datasets: X-ray tomography and serial sectioning. *Materials Characterization*, 160:110119, 2020.
- [19] Joshua A Taillon, Thomas F Bina, Raymond L Plante, Marcus W Newrock, Gretchen R Greene, and June W Lau. Nexuslims: A laboratory information management system for shared-use electron microscopy facilities. *Microscopy and Microanalysis*, 27(3):511–527, 2021.
- [20] S Jesse, M Chi, A Belianinov, C Beekman, SV Kalinin, AY Borisevich, and AR Lupini. Big data analytics for scanning transmission electron microscopy ptychography. *Scientific reports*, 6(1):1–8, 2016.
- [21] Florian Schroff, Antonio Criminisi, and Andrew Zisserman. Harvesting image databases from the web. *IEEE transactions on pattern analysis and machine intelligence*, 33(4):754–766, 2010.
- [22] Xian-Sheng Hua and Jin Li. Prajna: Towards recognizing whatever you want from images without image labeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- [23] Alexandra L Day, Carolin B Wahl, Vishu Gupta, Roberto Dos Reis, Wei-keng Liao, Chad A Mirkin, Vinayak P Dravid, Alok Choudhary, and Ankit Agrawal. Machine

- learning-enabled image classification for automated electron microscopy. *Microscopy and Microanalysis*, 30(3):456–465, 2024.
- [24] Angelo Ziletti, Devinder Kumar, Matthias Scheffler, and Luca M Ghiringhelli. Insightful classification of crystal structures using deep learning. *Nature communications*, 9(1):1–10, 2018.
- [25] Kamal Choudhary, Kevin F Garrity, Charles Camp, Sergei V Kalinin, Rama Vasudevan, Maxim Ziatdinov, and Francesca Tavazza. Computational scanning tunneling microscope image database. *Scientific data*, 8(1):1–9, 2021.
- [26] Eric Schwenker, Weixin Jiang, Trevor Spreadbury, Nicola Ferrier, Oliver Cossairt, and Maria KY Chan. Exsclaim!—an automated pipeline for the construction of labeled materials imaging datasets from literature. *arXiv preprint arXiv:2103.10631*, 2021. <https://arxiv.org/abs/2103.10631>. Accessed 2025-05-01.
- [27] Brian L DeCost, Matthew D Hecht, Toby Francis, Bryan A Webler, Yoosuf N Picard, and Elizabeth A Holm. Uhcsdb: ultrahigh carbon steel micrograph database. *Integrating Materials and Manufacturing Innovation*, 6(2):197–205, 2017.
- [28] Jacob Madsen, Pei Liu, Jens Kling, Jakob Birkedal Wagner, Thomas Willum Hansen, Ole Winther, and Jakob Schiøtz. A deep learning approach to identify local structures in atomic-resolution transmission electron microscopy images. *Advanced Theory and Simulations*, 1(8):1800037, 2018.
- [29] Artem Maksov, Ondrej Dyck, Kai Wang, Kai Xiao, David B Geohegan, Bobby G Sumpter, Rama K Vasudevan, Stephen Jesse, Sergei V Kalinin, and Maxim Ziatdinov. Deep learning analysis of defect and phase evolution during electron beam-induced transformations in ws₂. *npj Computational Materials*, 5(1):1–8, 2019.
- [30] Sang-Hyeok Yang, Wooseon Choi, Byeong Wook Cho, Frederick Osei-Tutu Agyapong-Fordjour, Sehwan Park, Seok Joon Yun, Hyung-Jin Kim, Young-Kyu Han, Young Hee Lee, Ki Kang Kim, et al. Deep learning-assisted quantification of atomic dopants and defects in 2d materials. *Advanced Science*, 8(16):2101099, 2021.
- [31] Graham Roberts, Simon Y Haile, Rajat Sainju, Danny J Edwards, Brian Hutchinson, and Yuanyuan Zhu. Deep learning for semantic segmentation of defects in advanced stem images of steels. *Scientific reports*, 9(1):1–12, 2019.
- [32] Lukas Vlcek, Maxim Ziatdinov, Artem Maksov, Alexander Tselev, Arthur P Baddorf, Sergei V Kalinin, and Rama K Vasudevan. Learning from imperfections: predicting structure and thermodynamics from atomic imaging of fluctuations. *ACS nano*, 13(1):718–727, 2019.
- [33] Maxim Ziatdinov, Artem Maksov, and Sergei V Kalinin. Learning surface molecular structures via machine vision. *npj Computational Materials*, 3(1):1–9, 2017.

- [34] Oleg S Ovchinnikov, Andrew O'Hara, Stephen Jesse, Bethany M Hudak, Shi-Ze Yang, Andrew R Lupini, Matthew F Chisholm, Wu Zhou, Sergei V Kalinin, Albina Y Borisevich, et al. Detection of defects in atomic-resolution images of materials using cycle analysis. *Advanced Structural and Chemical Imaging*, 6(1):1–9, 2020.
- [35] Kevin de Haan, Zachary S Ballard, Yair Rivenson, Yichen Wu, and Aydogan Ozcan. Resolution enhancement in scanning electron microscopy using deep learning. *Scientific reports*, 9(1):1–7, 2019.
- [36] Jeffrey M Ede and Richard Beanland. Partial scanning transmission electron microscopy with deep learning. *Scientific reports*, 10(1):1–10, 2020.
- [37] Mohammad Rashidi and Robert A Wolkow. Autonomous scanning probe microscopy in situ tip conditioning through machine learning. *Acs Nano*, 12(6):5185–5189, 2018.
- [38] Sagi Eppel, Haoping Xu, Mor Bismuth, and Alan Aspuru-Guzik. Computer vision for recognition of materials and vessels in chemistry lab settings and the vector-labpics data set. *ACS central science*, 6(10):1743–1752, 2020.
- [39] Luke Scime, Derek Siddel, Seth Baird, and Vincent Paquit. Layer-wise anomaly detection and classification for powder bed additive manufacturing processes: A machine-agnostic algorithm for real-time pixel-wise semantic segmentation. *Additive Manufacturing*, 36:101453, 2020.
- [40] Ankit Agrawal and Alok Choudhary. Deep materials informatics: Applications of deep learning in materials science. *Mrs Communications*, 9(3):779–792, 2019.
- [41] John Bardeen. Tunnelling from a many-particle point of view. *Physical review letters*, 6(2):57, 1961.
- [42] Jerry Tersoff and Donald R Hamann. Theory and application for the scanning tunneling microscope. *Physical review letters*, 50(25):1998, 1983.
- [43] C Julian Chen and Walter F Smith. Introduction to scanning tunneling microscopy. *American Journal of Physics*, 62(6):573–574, 1994.
- [44] Aidan H Combs, Jason J Maldonis, Jie Feng, Zhongnan Xu, Paul M Voyles, and Dane Morgan. Fast approximate stem image simulations from a machine learning model. *Advanced Structural and Chemical Imaging*, 5(1):1–10, 2019.
- [45] Earl J Kirkland. *Advanced computing in electron microscopy*, volume 12. Springer, 1998.
- [46] Earl J Kirkland. Computation in electron microscopy. *Acta Crystallographica Section A: Foundations and Advances*, 72(1):1–27, 2016.
- [47] LJ Allen, SD Findlay, MP Oxley, and CJ Rossouw. Lattice-resolution contrast from a focused coherent electron probe. part i. *Ultramicroscopy*, 96(1):47–63, 2003.

- [48] JM Cowley and AF Moodie. Fourier images: I-the point source. *Proceedings of the Physical Society. Section B*, 70(5):486, 1957.
- [49] Jacob Madsen and Toma Susi. The abtem code: transmission electron microscopy from first principles. *Open Research Europe*, 1(24):24, 2021.
- [50] Maxim Ziatdinov, Ayana Ghosh, Tommy Wong, and Sergei V Kalinin. Atomai: A deep learning framework for analysis of image and spectroscopy data in (scanning) transmission electron microscopy and beyond. *Nature Mach. Intell.*, 4(arXiv: 2105.07485):1101–1112, 2022.
- [51] Colin Ophus. A fast image simulation algorithm for scanning transmission electron microscopy. *Advanced structural and chemical imaging*, 3(1):1–11, 2017.
- [52] Christoph Tobias Koch. *Determination of core structure periodicity and point defect density along dislocations*. Arizona State University, 2002.
- [53] Kamal Choudhary, Ramya Gurunathan, Brian DeCost, and Adam Biacchi. Atomvision: A machine vision library for atomistic images. *Journal of Chemical Information and Modeling*, 63(6):1708–1722, 2023.
- [54] Anwai Archit, Luca Freckmann, Sushmita Nair, Nabeel Khalid, Paul Hilt, Vikas Rajashekar, Marei Freitag, Carolin Teuber, Melanie Spitzner, Constanza Tapia Contreras, et al. Segment anything for microscopy. *Nature Methods*, 22(3):579–591, 2025.
- [55] Zhuo Diao, Keiichi Ueda, Linfeng Hou, Fengxuan Li, Hayato Yamashita, and Masayuki Abe. Ai-equipped scanning probe microscopy for autonomous site-specific atomic-level characterization at room temperature. *Small Methods*, 2024.
- [56] Ryan Jacobs, Mingren Shen, Yuhan Liu, Wei Hao, Xiaoshan Li, Ruoyu He, Jacob RC Greaves, Donglin Wang, Zeming Xie, Zitong Huang, et al. Performance and limitations of deep learning semantic segmentation of multiple defects in transmission electron micrographs. *Cell Reports Physical Science*, 3(5), 2022.
- [57] I. Lobato, T. Friedrich, and S. Van Aert. Deep convolutional neural networks to restore single-shot electron microscopy images. *npj Computational Materials*, 10:10, 2024.
- [58] Shiro Ihara, Hikaru Saito, Mizumo Yoshinaga, Lavakumar Avala, and Mitsuhiro Murayama. Deep learning-based noise filtering toward millisecond-order imaging by using scanning transmission electron microscopy. *Scientific Reports*, 12:13462, 2022.
- [59] Rajat Sainju, Wei-Ying Chen, Samuel Schaefer, Qian Yang, Caiwen Ding, Meimei Li, and Yuanyuan Zhu. Defecttrack: a deep learning-based multi-object tracking algorithm for quantitative defect analysis of in-situ tem videos in real-time. *Scientific Reports*, 12:15705, 2022.
- [60] Thomas Bilyk, Alexandra M. Goryaeva, Mihai-Cosmin Marinica, Camille Flament, Catherine Sabathier, Eric Leroy, Marie Loyer-Prost, and Estelle Meslin. Accurate quantification of dislocation loops in complex functional alloys enabled by deep learning image analysis. *Scientific Reports*, 14:25168, 2024.

- [61] Ang’eline Lafleur and Soo-hyon Phark. Automated classification of individual atoms on surfaces using machine learning. *arXiv preprint arXiv:2410.13711*, 2024. <https://arxiv.org/abs/2410.13711>. Accessed 2025-05-01.
- [62] Prateek Verma, Minh-Hao Van, and Xintao Wu. Beyond human vision: The role of large vision language models in microscope image analysis. In *2024 IEEE International Conference on Big Data (BigData)*, pages 1700–1705. IEEE, 2024.
- [63] Yaotian Yang, Hao Xiong, Zirong Wu, Zhiyao Luo, Xiao Chen, Xiaonan Wang, and Fei Wei. Deep learning-enabled stem imaging for precise single-molecule identification in zeolite structures. *Advanced Science*, page e2408629, 2024.
- [64] Alejandro Lozano, Jeffrey Nirschl, James Burgess, Sanket R. Gupte, Yuhui Zhang, Alyssa Unell, and Serena Yeung-Levy. Micro-Bench: A vision-language benchmark for microscopy understanding. In *Proc. 38th Conf. on Neural Information Processing Systems (NeurIPS 2024) Datasets and Benchmarks Track*, 2024.
- [65] Shanghang Zhang, Gaole Dai, Tiejun Huang, and Jianxu Chen. Multimodal large language models for bioimage analysis. *Nature Methods*, 21(8):1390–1393, 2024.
- [66] Yuwei Mao, Mahmudul Hasan, Md Maruf Billah, Youjia Li, Sayak Chakrabarty, Claire Songhyun Lee, Kewei Wang, Muhammed Nur Talha Kilic, Vishu Gupta, Weikeng Liao, et al. An ai framework for time series microstructure prediction from processing parameters. *Scientific Reports*, 15(1):24074, 2025.
- [67] Ge Lei, Ronan Docherty, and Samuel J. Cooper. Materials science in the era of large language models: a perspective. *Digital Discovery*, 3:1257–1272, 2024.
- [68] Indrajeet Mandal, Jitendra Soni, Mohd Zaki, Morten M. Smedskjaer, Katrin Wondraczek, Lothar Wondraczek, Nitya Nand Gosvami, and N. M. Anoop Krishnan. Autonomous microscopy experiments through large language model agents. *arXiv preprint arXiv:2501.10385*, 2024. <https://arxiv.org/abs/2501.10385>. Accessed 2025-05-01.
- [69] Sams Jarin, Yufan Yuan, Mingxing Zhang, Mingwei Hu, Masud Rana, Sen Wang, and Ruth Knibbe. Predicting the crystal structure and lattice parameters of the perovskite materials via different machine learning models based on basic atom properties. *Crystals*, 12(11):1570, 2022.
- [70] Logan Williams, Arpan Mukherjee, and Krishna Rajan. Deep learning based prediction of perovskite lattice parameters from hirshfeld surface fingerprints. *The Journal of Physical Chemistry Letters*, 11(17):7462–7468, 2020.
- [71] Yuxin Li, Wenhui Yang, Rongzhi Dong, and Jianjun Hu. Mlatticeabc: generic lattice constant prediction of crystal materials using machine learning. *ACS omega*, 6(17):11585–11594, 2021.
- [72] Jonathan Schmidt, Tiago FT Cerqueira, Aldo H Romero, Antoine Loew, Fabian Jäger, Hai-Chen Wang, Silvana Botti, and Miguel AL Marques. Improving machine-learning

- models in materials science through large datasets. *Materials Today Physics*, 48:101560, 2024.
- [73] Jonathan Schmidt, Noah Hoffmann, Hai-Chen Wang, Pedro Borlido, Pedro JMA Carriço, Tiago FT Cerqueira, Silvana Botti, and Miguel AL Marques. Machine-learning-assisted determination of the global zero-temperature phase diagram of materials. *Advanced Materials*, 35(22):2210788, 2023.
- [74] Hai-Chen Wang, Jonathan Schmidt, Miguel AL Marques, Ludger Wirtz, and Aldo H Romero. Symmetry-based computational search for novel binary and ternary 2d materials. *2D Materials*, 10(3):035007, 2023.
- [75] Sten Hastrup, Mikkel Strange, Mohnish Pandey, Thorsten Deilmann, Per S Schmidt, Nicki F Hinsche, Morten N Gjerding, Daniele Torelli, Peter M Larsen, Anders C Riis-Jensen, et al. The computational 2d materials database: high-throughput modeling and discovery of atomically thin crystals. *2D Materials*, 5(4):042002, 2018.
- [76] Morten Niklas Gjerding, Alireza Taghizadeh, Asbjørn Rasmussen, Sajid Ali, Fabio Bertoldo, Thorsten Deilmann, Nikolaj Rørbæk Knøsgaard, Mads Kruse, Ask Hjorth Larsen, Shidrokh Manti, et al. Recent progress of the computational 2d materials database (c2db). *2D Materials*, 8(4):044002, 2021.
- [77] Kamal Choudhary, Kevin F Garrity, Andrew CE Reid, Brian DeCost, Adam J Biacchi, Angela R Hight Walker, Zachary Trautt, Jason Hattrick-Simpers, A Gilad Kusne, Andrea Centrone, et al. The joint automated repository for various integrated simulations (jarvis) for data-driven materials design. *npj Computational Materials*, 6(1):1–13, 2020.
- [78] Kamal Choudhary, Irina Kalish, Ryan Beams, and Francesca Tavazza. High-throughput identification and characterization of two-dimensional materials using density functional theory. *Scientific Reports*, 7(1):1–16, 2017.
- [79] Daniel Wines, Ramya Gurunathan, Kevin F Garrity, Brian DeCost, Adam J Biacchi, Francesca Tavazza, and Kamal Choudhary. Recent progress in the jarvis infrastructure for next-generation data-driven materials design. *Applied Physics Reviews*, 10(4), 2023.
- [80] Kamal Choudhary, Kevin F. Garrity, Andrew CE Reid, Brian DeCost, Adam J. Biacchi, Angela R. Hight Walker, Zachary Trautt, Jason Hattrick-Simpers, Aaron G. Kusne, Andrea Centrone, et al. The joint automated repository for various integrated simulations (jarvis) for data-driven materials design. *npj Computational Materials*, 6(1):173, 2020.
- [81] Jonathan Schmidt, Love Pettersson, Claudio Verdozzi, Silvana Botti, and Miguel AL Marques. Crystal graph attention networks for the prediction of stable materials. *Science advances*, 7(49):eabi7948, 2021.
- [82] John P Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Physical review letters*, 77(18):3865, 1996.

- [83] Shunsuke Yamashita, Jun Kikkawa, Keiichi Yanagisawa, Takuro Nagai, Kazuo Ishizuka, and Koji Kimoto. Atomic number dependence of z contrast in scanning transmission electron microscopy. *Scientific reports*, 8(1):1–7, 2018.
- [84] Pierre Hohenberg and Walter Kohn. Inhomogeneous electron gas. *Physical review*, 136(3B):B864, 1964.
- [85] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [86] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [87] Nicolas Mounet, Marco Gibertini, Philippe Schwaller, Davide Campi, Andrius Merkys, Antimo Marrazzo, Thibault Sohier, Ivano Eligio Castelli, Andrea Cepellotti, Giovanni Pizzi, et al. Two-dimensional materials from high-throughput computational exfoliation of experimentally known compounds. *Nature nanotechnology*, 13(3):246–252, 2018.
- [88] Juan Ivan Gomez-Peralta, Xim Bokhimi, and Patricia Quintana. Convolutional neural networks to assist the assessment of lattice parameters from x-ray powder diffraction. *The Journal of Physical Chemistry A*, 127(36):7655–7664, 2023.
- [89] Sathya R Chitturi, Daniel Ratner, Richard C Walroth, Vivek Thampy, Evan J Reed, Mike Dunne, Christopher J Tassone, and Kevin H Stone. Automated prediction of lattice parameters from x-ray powder diffraction patterns. *Applied Crystallography*, 54(6):1799–1810, 2021.
- [90] Xiaofeng Qian, Junwei Liu, Liang Fu, and Ju Li. Quantum spin hall effect in two-dimensional transition metal dichalcogenides. *Nano Letters*, 14(9):4959–4965, 2014.
- [91] Anlian Lu, Han Zhang, Rui Zhang, Bing Liu, and Hualing Zeng. Janus monolayers of transition metal dichalcogenides. *ACS Nano*, 12(8):7445–7452, 2018.
- [92] Yizhong Chen, Wenjin Wang, Shu Guo, Qingsong Wang, and Tao Wu. Strong electron correlations in monolayer 1t-tas₂. *Nature Communications*, 11(1):2453, 2020.
- [93] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016.
- [94] Valerio Biscione and Jeffrey S Bowers. Convolutional neural networks are not invariant to translation, but they can learn to be. *Journal of Machine Learning Research*, 22(229):1–28, 2021.
- [95] Lixing Kang, Chen Ye, Xiaoxu Zhao, Xieyu Zhou, Junxiong Hu, Qiao Li, Dan Liu, Chandreyee Manas Das, Jiefu Yang, Dianyi Hu, Jieqiong Chen, Xun Cao, Yong Zhang, Manzhang Xu, Jun Di, Dan Tian, Pin Song, Govindan Kutty, Qingsheng

- Zeng, Qundong Fu, Ya Deng, Jiadong Zhou, Ariando Ariando, Feng Miao, Guo Hong, Yizhong Huang, Stephen J. Pennycook, Ken-Tye Yong, Wei Ji, Xiao Renshaw Wang, and Zheng Liu. Phase-controllable growth of ultrathin 2d magnetic fete crystals. *Nature Communications*, 11(1):3729, 2020.
- [96] Science and Industry Museum. Graphene: A new way of thinking about materials. <https://www.scienceandindustrymuseum.org.uk/objects-and-stories/graphene>, November 2021. Accessed: 2025-10-20.
- [97] Dongha Shin, Jong Bo Park, Yong-Jin Kim, Sang Jin Kim, Jin Hyoun Kang, Bora Lee, Sung-Pyo Cho, Byung Hee Hong, and Konstantin S Novoselov. Growth dynamics and gas transport mechanism of nanobubbles in graphene liquid cells. *Nature Communications*, 6(1):6068, 2015.
- [98] Gao Yang, Lihua Li, Wing Bun Lee, and Man Cheung Ng. Structure of graphene and its disorders: a review. *Science and Technology of Advanced Materials*, 19(1):613–648, 2018. PMID: 30181789.
- [99] Zhimo Zhang, Min Cai, Rui Li, Fanqi Meng, Qinghua Zhang, Lin Gu, Zijin Ye, Gang Xu, Ying-Shuang Fu, and Wenhao Zhang. Controllable synthesis and electronic structure characterization of multiple phases of iron telluride thin films. *Phys. Rev. Mater.*, 4:125003, Dec 2020.
- [100] Yen-Hui Lin, Yu-Hsiang Chung, Guan-Yu Chen, Hung-Chin Lee, Nitin Kumar, Chia-Ju Chen, Jia-An Lin, Chun-Yu Mou, Allan H. MacDonald, Chih-Kang Shih, Dah-An Luh, Jung-Jung Su, and Pin-Jui Hsu. Fabricating a hexagonal fete monolayer with a moiré superlattice on topological insulator Bi_2Te_3 . *Nanoscale Adv.*, 7:3088–3095, 2025.

Predicting Lattice Parameters from Atomic-Scale Images of Two Dimensional Materials Using Deep Learning

Sayak Chakrabarty¹, Kamal Choudhary^{2,3,4}, Daniel Wines², Youjia Li⁵, Vishu Gupta^{5,6,7},
Muhammed Nur Talha Kilic⁵, Alok Choudhary⁵, Ankit Agrawal⁵

¹ *Department of Computer Science, Northwestern University, Evanston, IL 60208, USA*

² *Material Measurement Laboratory, National Institute of Standards and Technology, Gaithersburg, MD 20899, USA*

³ *Department of Materials Science and Engineering, Whiting School of Engineering, The Johns Hopkins University, Baltimore, Maryland 21218, USA*

⁴ *Department of Electrical and Computer Engineering, Whiting School of Engineering, The Johns Hopkins University, Baltimore, Maryland 21218, USA*

⁵ *Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL 60208, USA*

⁶ *Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, NJ 08544, USA*

⁷ *Cancer Research Institute and Institute of Integrative Genomics, Princeton University, NJ 08544, USA*

Table S1: Comparison of 2D Bravais lattice classification performance (accuracy % and F1 score) of three model architectures: Clip+MLP, Densenet121, and GoogleNet on the STEM-JV2D and STEM-C2DB datasets. Densenet121 achieves the highest overall accuracy and F1 score and is selected as the best-performing model.

Dataset	Clip + MLP		Densenet121		GoogleNet	
	Accuracy (%)	F1 Score	Accuracy (%)	F1 Score	Accuracy (%)	F1 Score
Stem_jv2d	83.41	0.71	80.65	0.81	79.72	0.78
Stem_c2db	87.09	0.73	89.22	0.88	84.11	0.84

Table S2: Predictive performance of a single multi-task DenseNet (jointly regressing a , b , and γ) versus three independent DenseNet models trained separately for each parameter. Mean absolute error (MAE) and root-mean-square error (RMSE) are reported on the JV2D test set; lower values indicate better accuracy.

Parameter	Single Multi-Target DenseNet		Three Single-Target DenseNet	
	MAE	RMSE	MAE	RMSE
a (Å)	0.60	1.22	1.28	1.75
b (Å)	1.00	1.83	1.87	2.52
γ (°)	6.17	16.13	16.10	17.77

Table S3: The number of train and test samples for each of the three datasets, across each lattice class (namely 0, 1, 2, 3, 4), is given in this table.

Dataset	Data Type	0	1	2	3	4
STEM-JV2D	Training	367	137	261	27	67
	Testing	92	35	65	7	16
STEM-C2DB	Training	1263	592	574	252	134
	Testing	316	148	144	63	34
Alex-PBE	Training	22856	7164	58315	14052	7876
	Testing	5715	1792	14579	3513	1970

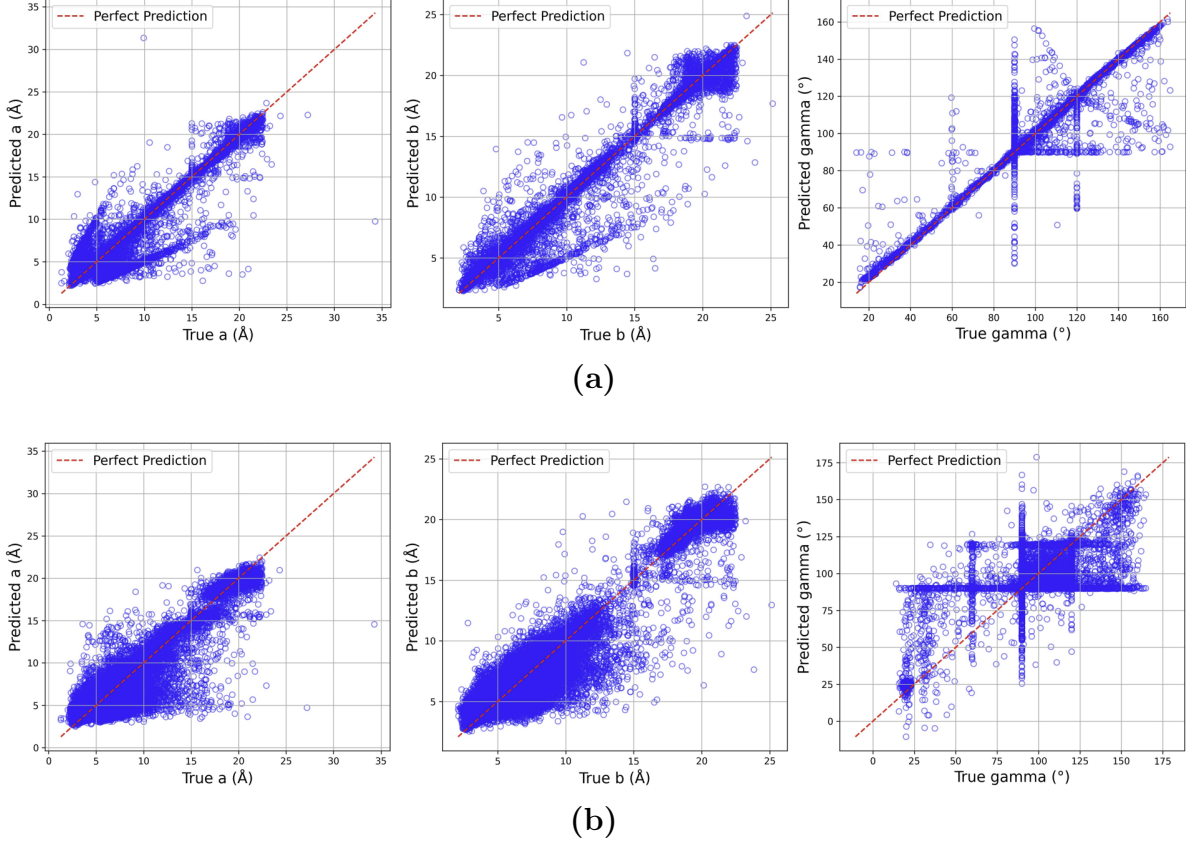


Figure S1: The (a) GoogleNet model and the (b) Clip + MLP model for prediction of the lattice parameters a, b, γ on the Alex-PBE dataset. We show the scatter plots when the trained model is evaluated on the full test subset of the Alex-PBE dataset. In each scatter plot, we have three subplots showing the scatter plot for a, b, γ where we have shown True vs Predicted values.

Table S4: Comparison of regression results for lattice parameters a, b, γ on the Alex-PBE dataset using GoogLeNet and ViT+MLP. n represents the number of samples.

Test Subset	Parameter	GoogLeNet		Clip+MLP	
		MAE	RMSE	MAE	RMSE
Full Test Set (n=27568)	a (Å)	0.40	1.15	1.05	1.68
	b (Å)	0.29	0.85	0.92	1.43
	γ (°)	1.00	4.62	4.08	11.03
$a \leq 10$ Å, $b \leq 10$ Å, $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$	a (Å)	0.22	0.74	0.84	1.26
	b (Å)	0.14	0.50	0.77	1.11
	γ (°)	0.59	3.12	1.94	6.29

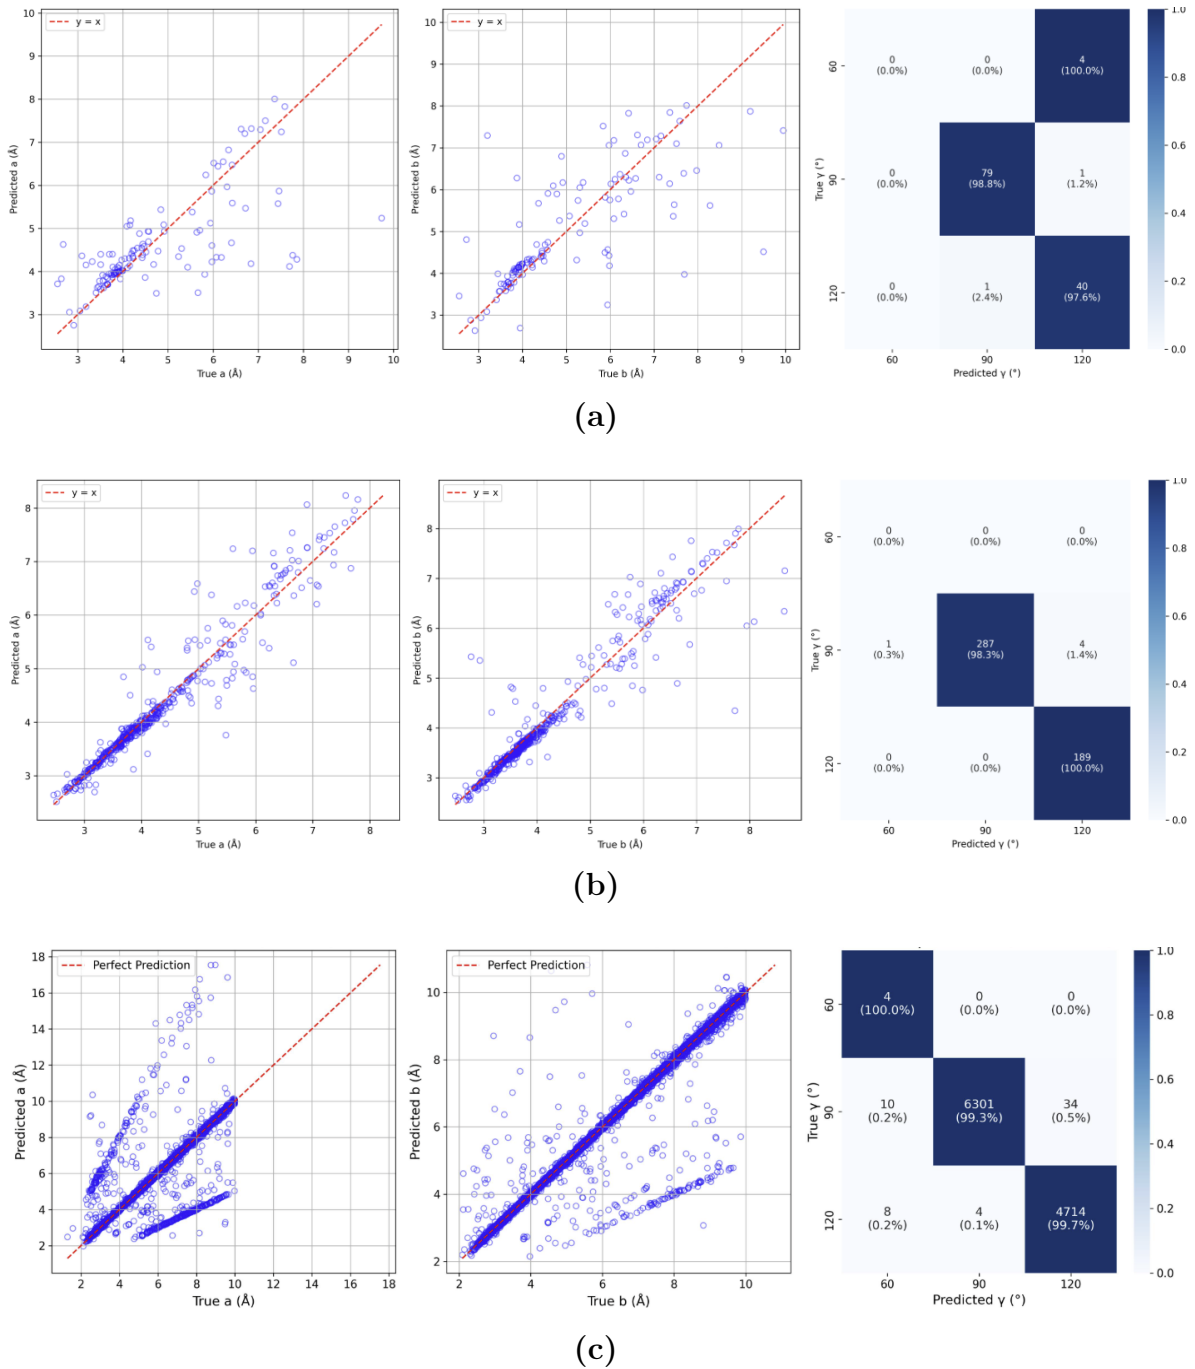

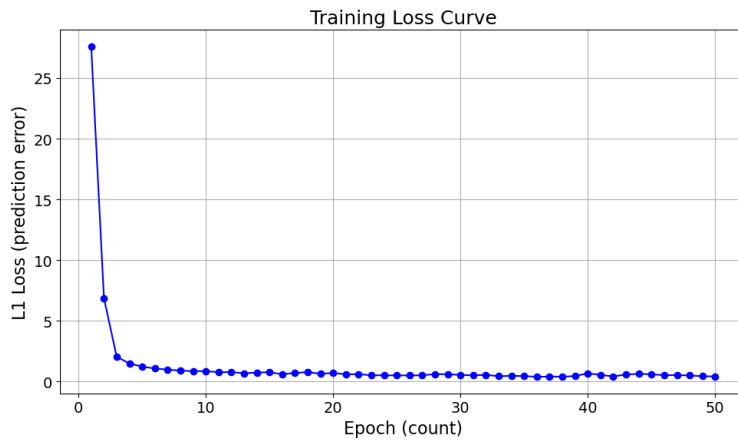


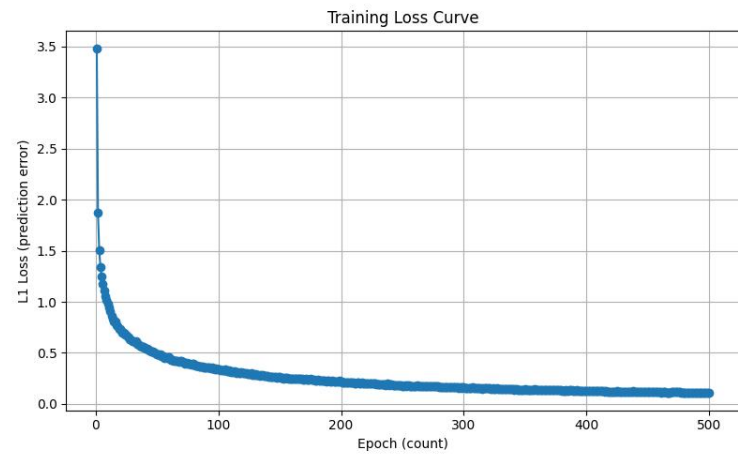
Figure S2: The Densenet model for prediction of the lattice parameters a , b , γ across the three datasets, (a) JV2D, (b) C2DB and (c) Alex-PBE. We show the scatter plots when the trained model is evaluated on only the subset of the test set where $a, b \leq 10 \text{ \AA}$, and we show confusion matrix plots when $\gamma \in \{60^\circ, 90^\circ, 120^\circ\}$.



(a)



(b)



(c)

Figure S3: Training loss vs epochs when we trained our Densenet model on the (a) JV2D dataset (b) C2DB dataset (c) Alex-PBE dataset

Table S5: Subset-wise performance of dataset-specific DenseNet models predicting lattice parameters a , b (Å) and γ (°). Results are computed on test datasets only for JARVIS-2D, C2DB and Alex-PBE. Groups are defined as: metals ($E_g \leq 0.05$ eV), semiconductors/insulators ($E_g > 0.05$ eV), chalcogenides (S/Se/Te), halides (F/Cl/Br/I), and TMDs (MX_2).

(a) JARVIS-2D

Group	Samples	a (MAE / RMSE)	b (MAE / RMSE)	γ (MAE / RMSE)
Semiconductors/Insulators	137	0.79 / 1.50	1.15 / 2.15	5.75 / 15.20
Metals	78	0.53 / 0.91	0.71 / 1.42	4.29 / 13.02
Chalcogenides (S/Se/Te)	90	0.89 / 1.61	1.20 / 2.10	5.83 / 14.28
Halides (F/Cl/Br/I)	94	0.66 / 1.27	0.80 / 1.57	4.00 / 12.84
TMD (MX_2)	13	0.15 / 0.20	0.11 / 0.14	0.66 / 0.71

(b) C2DB

Group	Samples	a (MAE / RMSE)	b (MAE / RMSE)	γ (MAE / RMSE)
Semiconductors/Insulators	203	0.22 / 0.40	0.25 / 0.48	0.76 / 3.22
Metals	502	0.15 / 0.37	0.16 / 0.40	0.91 / 5.34
Chalcogenides (S/Se/Te)	308	0.21 / 0.38	0.23 / 0.44	0.57 / 2.63
Halides (F/Cl/Br/I)	303	0.13 / 0.25	0.15 / 0.35	0.41 / 2.19
TMD (MX_2)	52	0.07 / 0.12	0.11 / 0.18	0.41 / 1.29

(c) Alex-PBE

Group	Samples	a (MAE / RMSE)	b (MAE / RMSE)	γ (MAE / RMSE)
Semiconductors/Insulators	10895	0.41 / 1.33	0.29 / 0.96	1.03 / 4.78
Metals	16674	0.32 / 1.05	0.20 / 0.74	0.82 / 4.03
Chalcogenides (S/Se/Te)	9022	0.35 / 1.13	0.24 / 0.86	0.88 / 4.39
Halides (F/Cl/Br/I)	10821	0.37 / 1.24	0.25 / 0.88	0.85 / 4.09
TMD (MX_2)	307	0.26 / 0.84	0.10 / 0.44	0.49 / 2.05

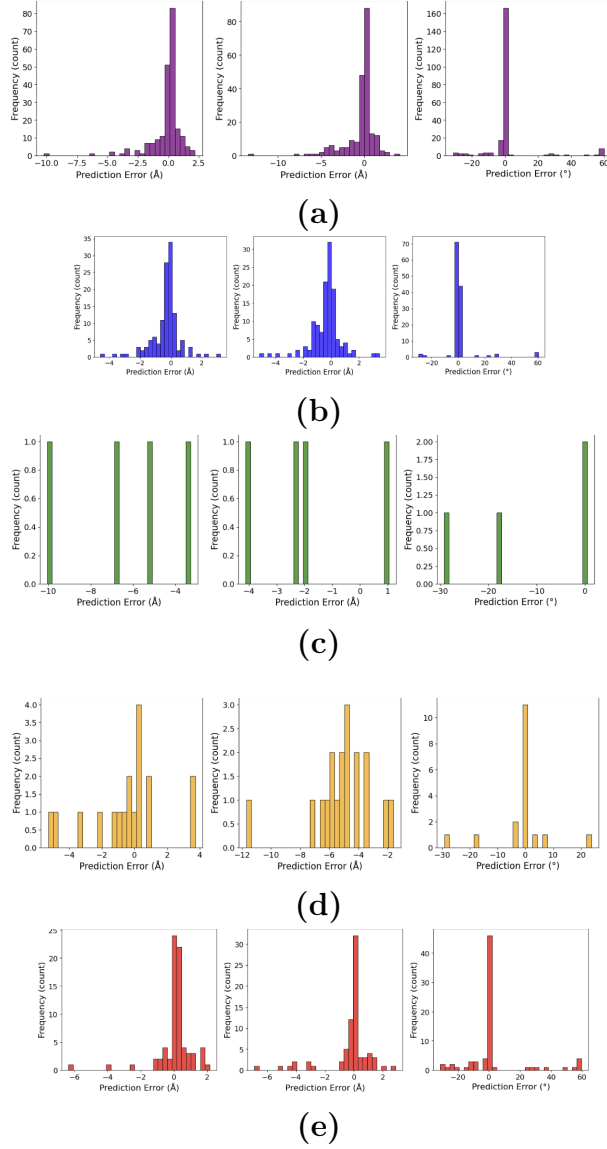


Figure S4: Histogram of prediction errors for lattice parameters a , b , and γ after training the DenseNet121 model on the JV2D dataset. The x-axis shows the prediction error (predicted - actual), and the y-axis shows the count. Panels show: (a) full test set, (b) subset with $a \leq 10$, $b \leq 10$, and $\gamma = 60^\circ, 90^\circ, 120^\circ$, (c) subset with $a > 10$, (d) subset with $b > 10$, (e) subset with non-standard angles $\gamma \neq 60^\circ, 90^\circ, 120^\circ$

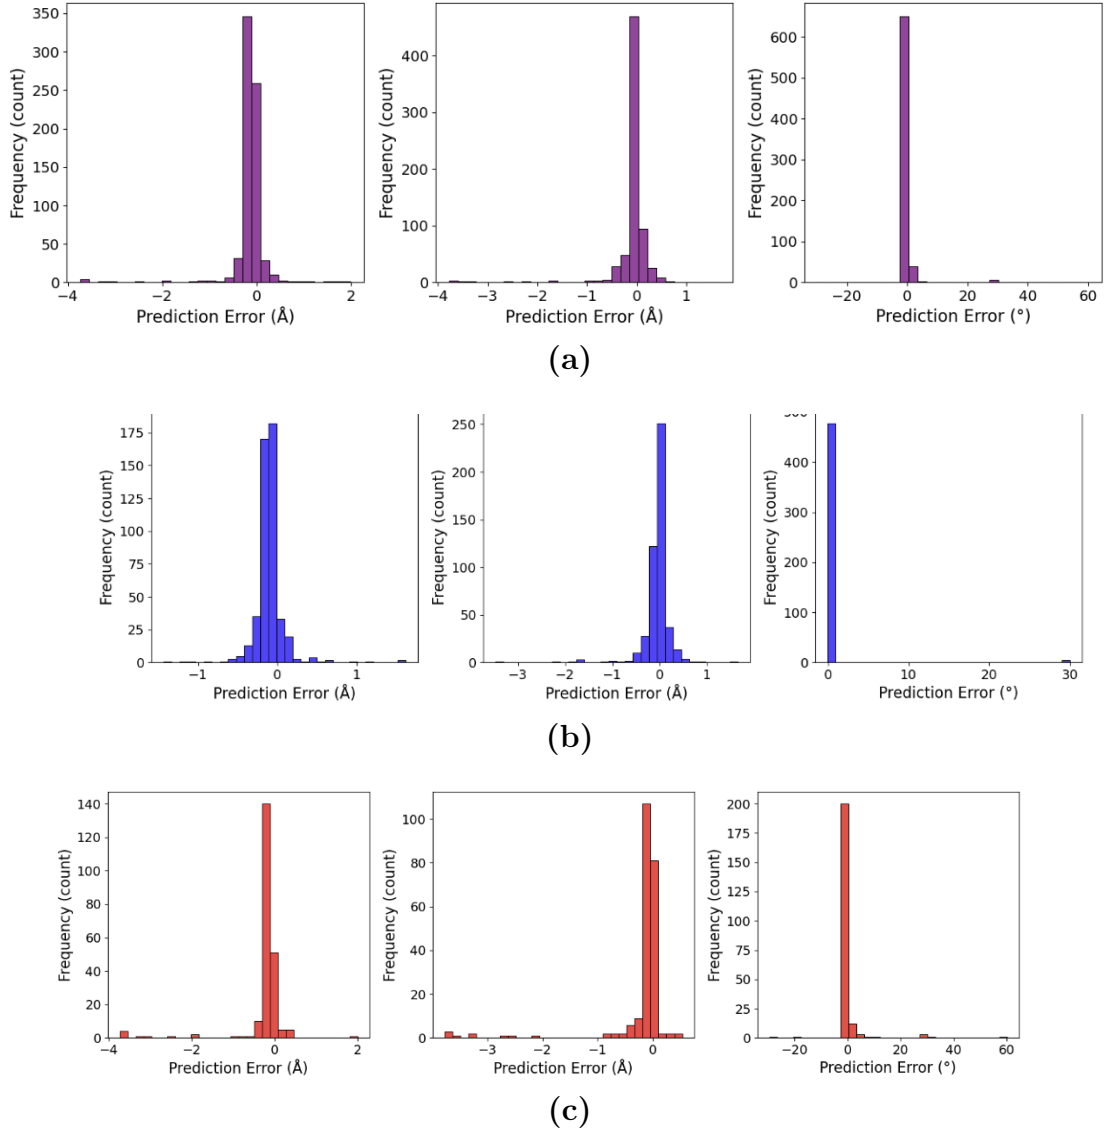


Figure S5: Histogram of prediction errors for lattice parameters a , b , and γ after training the DenseNet121 model on the C2DB dataset. The x-axis shows the prediction error (predicted - actual), and the y-axis shows the count. Panels show: (a) full test set, (b) subset with $a \leq 10$, $b \leq 10$, and $\gamma = 60^\circ, 90^\circ, 120^\circ$, (c) subset with non-standard angles $\gamma \notin \{60^\circ, 90^\circ, 120^\circ\}$.

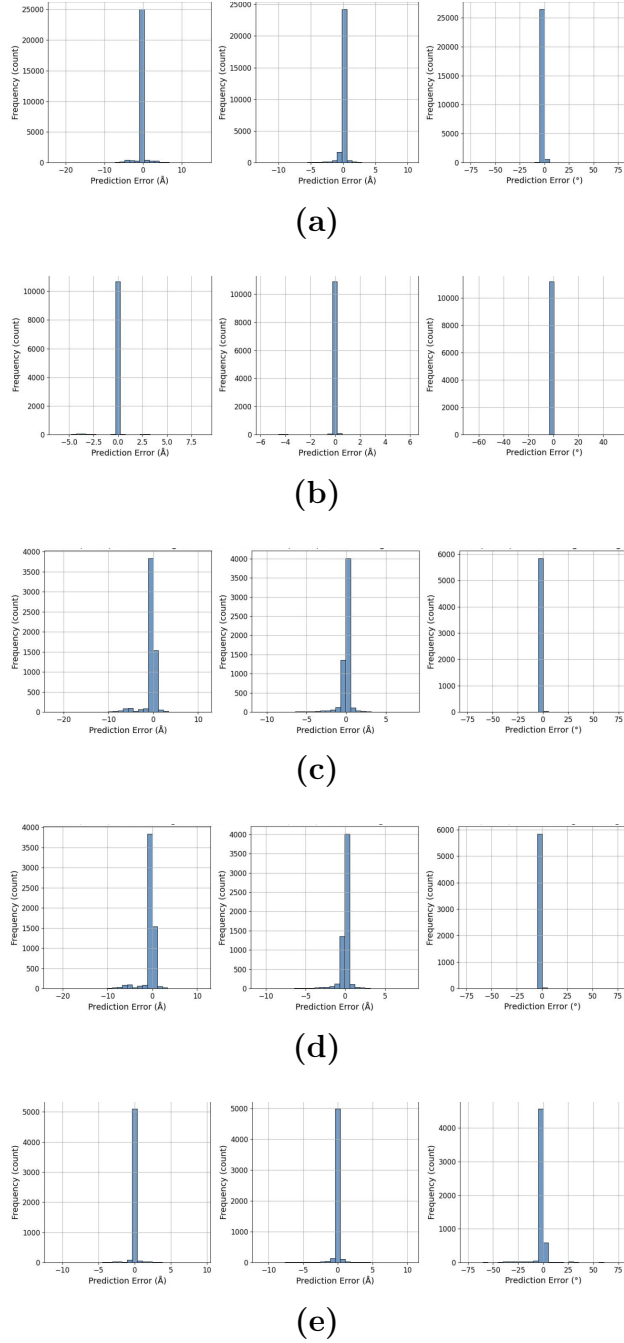


Figure S6: Histogram of prediction errors for lattice parameters a , b , and γ after training the DenseNet121 model on the Alex-PBE dataset. The x-axis shows the prediction error (predicted - actual), and the y-axis shows the count. Panels show: (a) full test set, (b) subset with $a \leq 10$, $b \leq 10$, and $\gamma = 60^\circ, 90^\circ, 120^\circ$, (c) subset with $a > 10$, (d) subset with $b > 10$, (e) subset with non-standard angles $\gamma \notin \{60^\circ, 90^\circ, 120^\circ\}$.




Figure S7: Experimental 2D STEM images taken from literature and used for model inference a) graphene (adapted from Ref. [96]) and b) graphene (adapted from Ref. [97]), c) tetragonal FeTe (adapted from Ref. [95]), d) hexagonal (adapted from Ref. [95]). a): Image credit: Sarah Haigh, University of Manchester and Quentin Ramasse, EPSRC SuperSTEM Laboratory, Daresbury. b): Adapted with permission from Springer Nature: Shin et al., Nature Communications 6, 6068 (2015). Copyright 2015 Springer Nature Ltd., c) and d): Adapted with permission from Springer Nature: Kang et al., Nature Communications 11, 3729 (2020). Copyright 2020 Author(s), licensed under a Creative Commons Attribution (CC BY) license.

Table S6: Comparison of predicted lattice parameters (a , b , γ) for selected materials from experiment (see Fig. S7).

Graphene (Ref. [96])			
Source	a (Å)	b (Å)	γ (°)
Experiment (Ref. [98])	2.46	2.46	120.00
JV2D	3.60	1.85	114.44
C2DB	2.62	2.13	119.31
Alex-PBE	1.82	6.22	89.78
Graphene (Ref. [97])			
Experiment (Ref. [98])	2.46	2.46	120.00
JV2D	3.30	1.99	114.21
C2DB	2.73	2.32	119.36
Alex-PBE	2.95	4.47	119.76
Tetragonal FeTe (Ref. [95])			
Experiment (Ref. [99, 100])	3.7–3.9	3.7–3.9	90.00
JV2D	3.87	3.77	90.32
C2DB	3.60	3.86	89.90
Alex-PBE	4.76	8.81	89.72
Hexagonal FeTe (Ref. [95])			
Experiment (Ref. [95, 100])	3.8–4.1	3.8–4.1	120.00
JV2D	4.14	3.91	118.57
C2DB	4.53	4.41	114.25
Alex-PBE	9.55	8.61	89.67