Automated Indexing Of TEM Diffraction Patterns Using Machine Learning

Nathaniel Tomczak, Sanmukh Kuppannagari
Department of Computer and Data Sciences, Case Western Reserve University
Contact: {nkt8, sxk1942}@case.edu

***Abstract*—Indexing Transmission Electron Microscopy (TEM) diffraction patterns is a critical step in materials characterization. Despite the manually intensive indexing process, work related to Machine Learning (ML) in its space is sparse. We present an evaluation of current state-of-the-art classification models and a Convolutional Neural Network (CNN), found through a Neural Architecture Search (NAS), in the TEM diffraction domain. Both convolution and transformer-based architectures were considered. Our NAS model achieved the greatest top-1 accuracy of 77.03% and F1 score of 0.751. The convolution-based architectures performed better, with EfficientNet-B3 achieving the highest average accuracy of 71.82% and tying the NAS model with the largest average F1 score of 0.686. These results can be used to guide further research into the better classification and creation of TEM diffraction data.**

***Index Terms*—Machine Learning, TEM, SAED, AI, CNN, GAN, Transfer Learning, Diffraction, Automation, Transformer**

I. Introduction

Transmission Electron Microscopy (TEM) is a fundamental tool in the field of materials science and engineering. High energy electrons are used to generate patterns containing orientation and structural information that is essential to the development and improvement of materials [1]. This information drives the development of stronger and lighter materials that underlie industries that range from aerospace and automotive to biological implants and electronics [2].

Unlike the biological fields where microscopy data analysis has transitioned to automated processing [3]–[5], the electron pattern identification process in materials science is predominantly manual, taking experienced researchers hundreds of hours to traverse, creating large backlogs of data. With the availability of High Performance Computing (HPC), Machine Learning (ML) models such as neural networks offer a promising solution for automated indexing of TEM diffraction patterns, as they have been shown to drastically reduce the time required to perform complicated tasks as evidenced by models like AlphaFold, which, while using a GPU, can predict protein structures in as little as minutes [6]. As the TEM diffraction field is relatively unexplored in regards to the employment of ML techniques, the application of neural networks could drastically expedite the indexing process and accelerate the rate of materials science research.

The goal of this paper is to foray into the research direction of high throughput automated indexing of TEM diffraction patterns by surveying the performance (accuracy) of current state-of-the-art ML models and drawing comparisons with a proposed Convolutional Neural Network (CNN) architecture discovered through a Neural Architecture Search (NAS). Four state-of-the-art ML architectures were assessed: VGG16 [7], ResNet-50 [8], EfficientNet-B3 [9], and ViT-L16 [10]. Each of these models were trained using two approaches: initialized from scratch (raw) where all layers are trained and pre-trained where only the model head was trained. This work made use of the High Performance Computing Resource in the Core Facility for Advanced Research Computing at Case Western Reserve University; specifically an AMD EPYC 7742 CPU and two Nvidia A100 GPUs (using distributed data parallelization). The results will be used to guide future research on developing better classification models as well as their parallelization strategies to achieve a high throughput indexing pipeline for TEM diffraction patterns, thereby exposing a new application domain to the power of HPC-driven Artificial Intelligence (AI)/ML.

II. Background

A. TEM Diffraction

Planes of atoms within a single crystal of a material coherently diffract a TEM's beam of high energy electrons [11]. These single crystals are composed of unique compositions and structures of elements. A sensor identifies the location of transmitted electrons and generates an image known as a Selected Area Electron Diffraction (SAED) pattern. These SAED patterns, also known as diffraction patterns, are key to the identification of a material and its structure. This information is fundamental to the process of materials characterization and discovery.

Figure 1 shows one of the 738 SAED patterns used in this work, provided by Professor Assel Aitkaliyeva of the University of Florida [12]. The individual spots in the pattern correspond to atomic planes; distinct arrangements of atoms vary the level and angle that electrons exit the single crystal [13]. on the unique combinations of diffraction angles and distances, researchers can identify a material's phase (also known as its structure). In the case of this work, phase is the unique classifier. One will notice a two-fold symmetry along the 8 to 2 o'clock and 11 to 5 o'clock axes, however there is no uniform grid of spots. This pattern for a Plutonium Zirconium (Pu-Zr) alloy was experimentally obtained from a TEM. There are irregularities

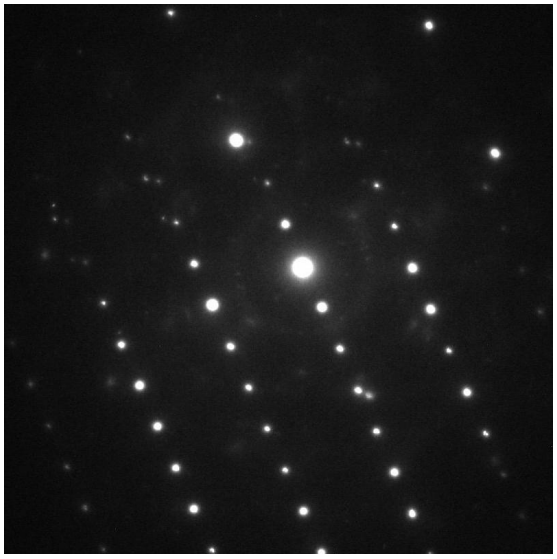

Fig. 1: An SAED pattern from a Plutonium Zirconium (Pu-Zr) alloy, specifically the κ-$PuZr_2$ phase. This is an actual pattern captured by a TEM. One can see an element of symmetry within the system; however there are noticeable irregularities in both spot size and intensity.

between patterns that come in the form of different spot sizes, varied levels of intensity throughout the pattern, and a soft haze around certain dots. An experienced researcher can analyze these patterns to identify the material and its structure.

B. Phase Indexing

SAED patterns are manually indexed by measuring the distances and angles between neighboring spots in relation to the highest intensity spot [14]. Physics-based algorithms show promise, but are complex, expensive, and slow to develop while still requiring human input and refinement [1], [15]. The manual method is used today in conjunction with physics-based software to help accelerate the process. However, even the most experienced researchers index at a rate of a few patterns per hour. A doctoral student could dedicate years to capturing and indexing the thousands of patterns required for their research.

ML has potential to expedite the process, though one potential barrier is the availability of training data; many of them are proprietary. In addition to issues with accessibility, SAED pattern datasets also tend to be relatively small, as evidenced by the 738 patterns available for this study. Despite these challenges, this work aims to distinguish between nine different phases of Pu-Zr alloys from a single SAED pattern using four state-of-the-art models and one discovered using NAS.

In conclusion, the problem statement of this work can be defined as follows: *given an SAED pattern as an input, an ML model will identify the corresponding phase of the image*.

The remaining portions of this paper are organized in the following manner: Section III covers adjacent work in automated diffraction pattern analysis, Section IV – Methods covers the dataset and models used in this research, Section V – Results & Discussion looks at the results observed while conducting the study, Section VI – Future Directions explores future avenues of research inspired by this work, and Section VII – Conclusion concludes as to the current state of automated indexing of TEM diffraction patterns using ML.

III. Related Works

ML research pertaining to the automated TEM indexing domain is relatively sparse. One group trained ResNet architectures of varying size to classify between 36 space groups (symmetry descriptions) for simulated SAED patterns at up 92.61% validation accuracy [16]. Another paper took a slightly different approach, using the 2D azimuthal integration of effective SAED patterns to feed peak positions into a 1D CNN. They were able to narrow 230 space groups to the top-2 with a confidence level ranging from 70% to 95% [17].

IV. Methods

A. Data

738 diffraction patterns were provided by Professor Aitkaliyeva [12]. Their distribution across nine different phases of Pu-Zr alloys are shown below in Table I.

TABLE I: The populations of each image's labelled material phase. The data set was split 80% training and 20% testing before rotations were applied to artificially increase the training size.

Image Phase	Training Population	Rotated Training Population	Testing Population
α-Zr	235	5,640	61
δ-(Pu, Zr)	194	4,656	47
β-Pu	36	864	6
Zr_3O-R-3c h	34	816	8
ZrO_2	26	624	3
PuO_2	25	600	10
PuO	19	456	4
κ-$PuZr_2$	17	408	7
Zr_3O-R32 h	4	96	2

In addition to the small number of images, the class divide is heavily skewed with α-Zr representing 40.1% of the dataset and Zr_3O-R32 h only 0.8%. The following pre-processing steps were performed to split the data, help alleviate the population size, and prepare it for model input:

- An 80% training and 20% testing split was used to proportionally divide the 738 patterns amongst each of the nine classes, leaving 590 and 148 patterns respectively in the training and testing datasets.

- A series of 15 degree central rotations were applied to each image in the training set within the range of [0, 360] degrees, resulting in 14,160 patterns in the training dataset. The test images were left untransformed.
- Each pattern was center-cropped, resulting in 800 pixel by 800 pixel images.
- The ViT-L16 model had each image resized from 800 pixels by 800 pixels to 224 pixels by 224 pixels in order to allow use of the pre-trained model. The raw variant used the resized images as well.

B. Models

The VGG16, ResNet-50, EfficientNet-B3, and ViT-L16 models were chosen as bases of comparison due to their historic state-of-the-art capabilities on the ImageNet classification task. VGG16, ResNet-50, and EfficientNet-B3 are all convolution-based architectures while ViT-L16 is a vision transformer. Additionally, it should be noted that work conducted in [16] used ResNet architectures to classify simulated SAED patterns; [17] has not been compared due to the inability to access their methods.

The NAS model, known as 299-AN, was discovered through a brute-force search of a model space consisting of 41,472 possible architectures. The search parameters are shown in Table II, with 299-AN's specific values bolded and highlighted in gray.

TABLE II: The model variables that were adjusted during the NAS of which 41,472 models were trained. Those values that are bold and highlighted in gray were used by model 299-AN. The dataset adjustment indicates whether or not the rotation for increasing the dataset size was applied or not.

Adjusted Variables	Possible Values
Number of Convolution Layers	1, 2, 3, 4, 5, **6**
Number of Dense Layers	1, 2, **3**, 4, 5, 6
Number of Nodes in the First Dense Layer	8, 16, 32, **64**
Batch Normalization	**True**, False
2D Dropout	True, **False**
Dropout	True, **False**
Learning Rate	0.01, **0.001**, 0.0001
Dense Layer Architecture	**Constant**, Encoder
Dataset	Original, **Rotated**
Batch Size	8, **16**, 32

C. Training

Each of the models were trained ten times for 20 epochs using a batch size of 32 patterns and the Adam optimizer with a learning rate of 0.001. PyTorch was used both for training and testing every model [18]. All models, except for the NAS model (299-AN), were trained in two different ways. The first, called raw, took freshly instantiated versions of the model and trained them from scratch. The second, called pre-trained, only trained a newly instantiated head for each model; the model body used frozen weights from training on ImageNet. 299-AN was only trained on its raw state.

V. Results & Discussion

Individual model performances are outlined in Table III. All accuracies provided are top-1 while all F1 scores are weighted due to the imbalanced class distribution shown in Table I. The values both in bold and highlighted in gray correspond to the highest observed metric. Model 299-AN from the NAS is the smallest in terms of parameter count while also providing the highest accuracy at 77.03% and an F1 score of 0.751.

The raw variants of the other models, except for ViT-L16, also managed to achieve similar accuracies in the 70s. The inability of the vision transformer model to learn as well as as the convolution-based architectures could be explained by the fact that our version trained with a batch size of 32 patterns while its implementation study, [10], trained with a much larger batch size of 4,096. The batch size parameter should be further explored. However, it is unlikely that the raw ViT-L16 model did not perform well due to the image resizing operation as the pre-trained version classified at a rate on par with the other raw variants.

ResNet-50's pre-trained model was the only other model to use transfer learning that performed moderately well with a best accuracy of 64.19%. Our model did not match the same accuracy observed in [16]'s ResNet-50. Their work saw a performance of 90.339% accuracy while using a pre-trained variant of the model and fine-tuning its parameters with simulated SAED patterns. One potential explanation for the difference in performance is our use of real SAED patterns and their exclusive use of simulated data. Given the perfect nature of simulations, the noise present in real SAED patterns could explain the variance in model performance. The impact of noise present in TEM indexing is worth exploring in the future.

The EfficientNet-B3 and VGG16 pre-trained models exhibited similar behavior to the raw ViT-L16, with 31.76% and 43.92% accuracy respectively. These results indicate that both EfficientNet-B3 and VGG16 architectures were unable to apply what they learned from being pre-trained with the ImageNet dataset as well as the others. Additionally, the raw variants, were able to learn enough to predict the phase at above 70% accuracy in almost all cases, except for ViT-L16.

When comparing the models' average accuracies, the raw VGG16 model stands out. All models experienced a drop in reported accuracy due to averaging, however VGG16 experiences a drop of 28.58% between the best and average accuracies. The explanation can be seen in Figure 2a: the test accuracy violin plot shows a huge variance in model performance. Nine out of 10 raw VGG16 models performed

TABLE III: Indicates whether or not a model was pre-trained and their respective parameter counts. Additionally, it provides the best and average accuracies (top-1) and F1 scores from 10 unique trainings of the model. It should be noted that only the classification layers were trained for the pre-trained models. The bolded items highlighted in gray correspond to the highest metric values observed in this study.

Model	Pre-Trained	Parameter Count (Millions)	Best Accuracy (%)	Avg. Accuracy (%)	Best F1 Score	Average F1 Score
NAS (299-AN)	F	0.265	**77.03**	71.62	**0.751**	**0.686**
EfficientNet-B3	F	10.7	75.68	**71.82**	0.71	**0.686**
EfficientNet-B3	T	10.7	31.76	28.11	0.281	0.261
VGG16	F	1,300	72.97	44.39	0.708	0.287
VGG16	T	1,300	43.92	43.04	0.285	0.274
ResNet-50	F	23.5	75.0	68.85	0.745	0.678
ResNet-50	T	23.5	64.19	61.01	0.622	0.57
ViT-L16	F	302	45.95	44.93	0.388	0.374
ViT-L16	T	302	72.97	69.8	0.676	0.633

at 41.22% testing accuracy. They only learned to predict 0, the class label pertaining to α-Zr, which encompassed 39.83% of the training dataset as shown in Table I. The pre-trained VGG16 models behaved similarly, but also learned to pick the Zr_3O-R-3c h phase and, rarely, the β-Pu phase.

Raw ViT-L16 has a similarly tight distribution of model performances, with it being able to predict both the α-Zr and δ-(Pu, Zr) phases. With these observations, it becomes apparent that test accuracy may not be the best metric for a model's ability to learn as one can approach 50% accuracy by guessing one or two of the majority phases. The weighted F1 score helps to solve this problem by applying weights equal to a class's probability within the population.

Figure 2b shows the F1 score distributions for each model. The pre-trained EfficientNet-B3, raw and pre-trained VGG16, and raw ViT-L16 models all exhibit a leftward shift in relation to the other models. Although F1 score is not a direct relation to accuracy, it does show that these models are making more incorrect predictions than the others, penalizing gaming of the system by choosing only a couple of classes.

The 299-AN, raw EfficientNet-B3, raw and pre-trained ResNet-50, and pre-trained ViT-L16 models predict more classes than the other four while also being more accurate in regard to said classes. Each of the five higher-performing models (in terms of F1 score) makes predictions that encompass between five and nine classes. However, it is uncommon for a model to make predictions inclusive of all nine classes. The only two phases that are present amongst all of the five best architectures were α-Zr and δ-(Pu, Zr). Logically, this makes sense as these phases make up 72.7% of the training dataset meaning that the model had the most exposure to each of these classes.

These results indicate that a larger model does not always yield a better result. 299-AN tied with raw EfficientNet-B3's average F1 score despite having 2.48% of the amount of parameters as seen in Table III. However, some of the larger models may not have been trained with optimal parameters, meaning that they could offer better performance given fine-tuning. These results will help create future directives to push greater levels of performance in the indexing of SAED patterns.

VI. Future Directions

The models covered in Section V are promising, showcasing the ability to detect differences between diffraction patterns of varied phase. However, they do not represent the culmination of this work; said models serve as a baseline for future methods considered below.

A. Physics-Informed

This work showcases models that rely solely on the capacity of the network to extract features from a two-dimensional image of an SAED pattern. Researchers are capable of doing the indexing from an entirely visual standpoint, however, they extract key angles, vector magnitudes, and often compare these values to other patterns from the same sample. The model currently lacks the capability to compare its representation to other SAED patterns, and, as it stands, one cannot say with absolute certainty what features the model chooses to extract and focus on. There may be key aspects of the dataset used within the manual methodology that are not being directly input to the model.

There is no guarantee that giving the model this specific information will help it better classify between phases, but a similar approach has been shown to work well in regards to material property predictions. The field often represents molecules as graphs, embedding physical information into its structure, and feeds them into Graph Neural Networks (GNN) with good results [19], [20]. These models are not necessarily capable of representing local (atomistic) interactions, which can be important in the prediction of electrical properties, after multiple graph convolutions. The Atomistic LIne Graph Neural Network (ALIGNN) model added an alternative, line graph neural networks, that incorporates bond information in a separate graph. This provided performance better than some of the previous GNN models in regards to atomistic

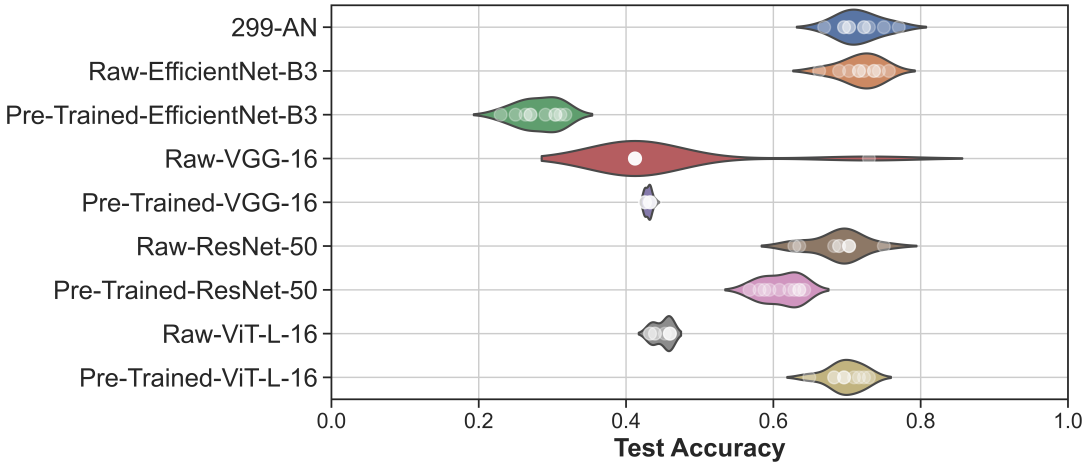

(a)

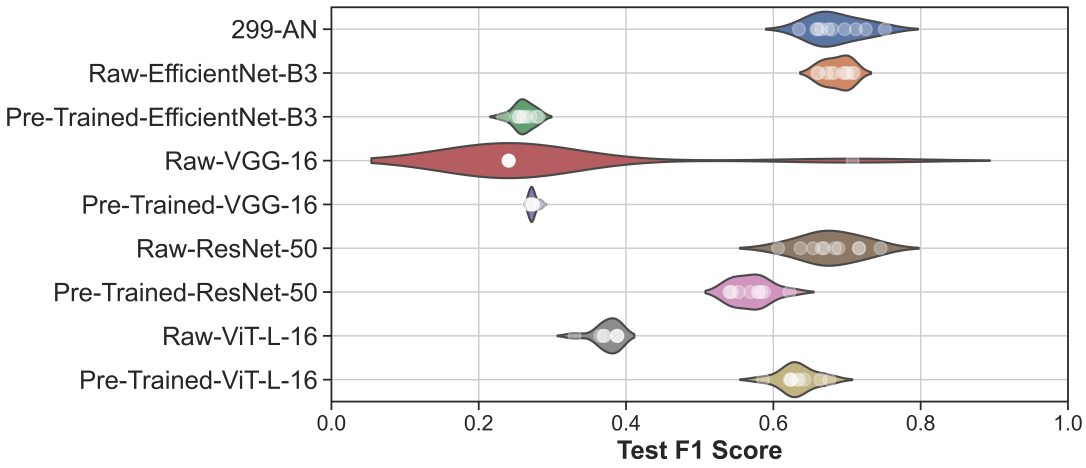

(b)

Fig. 2: Violin plots detailing the performances of 10 separately trained instances of each model architecture across two metrics. Each white point within the violin represents one trained model. All models were trained and tested on the same datasets outlined in Table I. **(a)** shows the model accuracies and **(b)** the weighted F1 scores. The weighted F1 score is useful for adjusting the model performances relative to the unbalanced dataset.

prediction [21].

GNNs are one potential area of research, but the takeaway is that the inclusion of alternative data representing physics-based knowledge could benefit the model.

B. Data Simulation

Although physics informed neural networks can be trained using small datasets [22], having more data is expected to improve accuracy. As evidenced by [16], [17] in Section III, diffraction patterns can be simulated. Figure 3 shows a simulated diffraction pattern [12], [23]. Unlike a real SAED pattern, as seen in Figure 1, simulated patterns have no noise. To our knowledge, there is no work that explores the generalizability of models trained on simulated SAED patterns being applied to real ones. This, as well as the more realistic simulation of SAED patterns, should be explored in the future. Work has just begun into using Generative Adversarial Networks (GAN) in order to simulate atomistic [24] and nanoparticle [25] images. Applying GANs to TEM diffraction could help solve the dataset size issue while simultaneously driving the development of a better classifier model, potentially helping to refine which physical properties and concepts are required in order to solve the problem.

C. High Performance Computing

The models showcased in this research are extremely resource intensive. Moreover, with the employment of Physics informed models such as GNNs similar to the one proposed in [21], we anticipate that complexity of models will increase due to irregular graph based computations. Moreover, as the models are fine-tuned using the vast

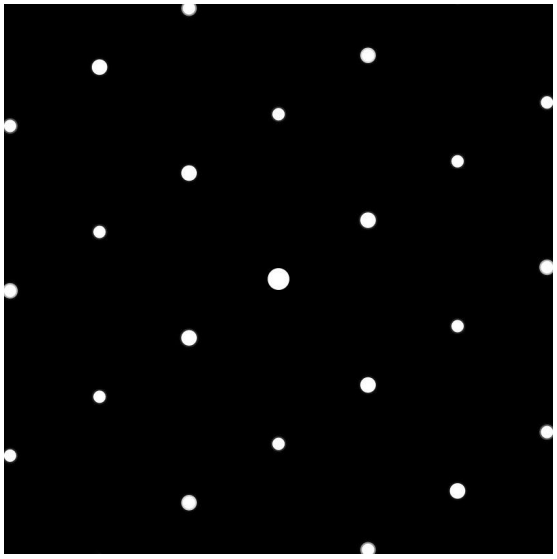

Fig. 3: An SAED pattern from the κ-$PuZr_2$ phase. This is a simulated pattern [23]. The diffraction is perfect, with no noise present in the system.

amounts of "unindexed" data available in the research community, these computational costs will continue to grow. Thus, one of our future directions will be to work on streamlining the training as well as indexing (inference) processes through exploration of parallelization techniques that are customized to the models developed for TEM diffraction classification.

VII. Conclusion

The automated classification of SAED pattern phases is an important issue as it serves as a major backbone for many materials science research projects. Despite a dataset of 738 images disproportionately distributed amongst nine classes, CNN architectures were capable of identifying SAED patterns at up to 77.03% accuracy, or a weighted F1 score of 0.751. Based on this work, we can draw the conclusion that the TEM diffraction domain is a great candidate for neural network-based methods. Additionally, there is potential for both further model and data development through the integration of the physics of TEM diffraction.

All model architectures and prediction results are available at: https://github.com/KLab-AI3/auto-saed

VIII. Acknowledgements

This work would not have been possible without Professor Assel Aitkaliyeva's generous donation of her dataset. We would like to also extend thanks to Professor Jennifer Carter and her oversight during the development of the NAS model, which is based upon work supported by the National Science Foundation (NSF) Award 1552716. This research was supported in part by NSF Award 2117439. Additionally, this work made use of the High Performance Computing Resource in the Core Facility for Advanced Research Computing at Case Western Reserve University.

References

[1] S. Nebaba, D. Zavyalov, and A. Pak, "Patterns detection in saed images of transmission electron microscopy," *Proceedings of the 30th International Conference on Computer Graphics and Machine Vision (GraphiCon 2020). Part 2*, pp. paper63–1, 12 2020.

[2] Z. Fan, L. Zhang, D. Baumann, L. Mei, Y. Yao, X. Duan, Y. Shi, J. Huang, Y. Huang, and X. Duan, "In situ transmission electron microscopy for energy materials and devices," *Advanced Materials*, vol. 31, no. 33, p. 1900608, 2019. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.201900608

[3] M. Balluet, F. Sizaire, Y. El Habouz, T. Walter, J. Pont, B. Giroux, O. Bouchareb, M. Tramier, and J. Pecreaux, "Neural network fast-classifies biological images through features selecting to power automated microscopy," *Journal of Microscopy*, vol. 285, no. 1, pp. 3–19, 2022. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/jmi.13062

[4] E. A. Hay and R. Parthasarathy, "Performance of convolutional neural networks for identification of bacteria in 3d microscopy datasets," *PLOS Computational Biology*, vol. 14, no. 12, pp. 1–17, 12 2018. [Online]. Available: https://doi.org/10.1371/journal.pcbi.1006628

[5] L. Möckl, A. R. Roy, P. N. Petrov, and W. E. Moerner, "Accurate and rapid background estimation in single-molecule localization microscopy using the deep neural network BGnet," *Proceedings of the National Academy of Sciences*, vol. 117, no. 1, pp. 60–67, dec 2019. [Online]. Available: https://doi.org/10.1073%2Fpnas.1916219117

[6] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, and et al., "Highly accurate protein structure prediction with alphafold," *Nature*, vol. 596, no. 7873, p. 583–589, 2021.

[7] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations*, 2015.

[8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: http://arxiv.org/abs/1512.03385

[9] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," *CoRR*, vol. abs/1905.11946, 2019. [Online]. Available: http://arxiv.org/abs/1905.11946

[10] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," *CoRR*, vol. abs/2010.11929, 2020. [Online]. Available: https://arxiv.org/abs/2010.11929

[11] D. Geelen, A. Thete, O. Schaff, A. Kaiser, S. J. van der Molen, and R. Tromp, "ev-tem: Transmission electron microscopy in a low energy cathode lens instrument," *Ultramicroscopy*, vol. 159, pp. 482–487, 2015, special Issue: LEEM-PEEM 9. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0304399115300024

[12] A. Aitkaliyeva, "Materials for nuclear advancement and technology in extreme environments," 2020, university of Florida.

[13] C. J. Humphreys, "The significance of Bragg's law in electron diffraction and microscopy, and Bragg's second law," *Acta Crystallographica Section A*, vol. 69, no. 1, pp. 45–50, Jan 2013. [Online]. Available: https://doi.org/10.1107/S0108767312047587

[14] Y. Liao, "Practical electron microscopy and database," 2006. [Online]. Available: www.globalsino.com/EM/

[15] J. Simbrunner, J. Domke, F. Sojka, D. Knez, R. Resel, T. Fritz, and R. Forker, "Automatic indexing of two-dimensional patterns in reciprocal space," *Phys. Rev. B*, vol. 104, p. 195402, Nov 2021. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevB.104.195402

[16] M. Ra, Y. Boo, J. M. Jeong, J. Batts-Etseg, J. Jeong, and W. Lee, "Classification of crystal structures using electron diffraction patterns with a deep convolutional neural network," *RSC Adv.*, vol. 11, pp. 38 307–38 315, 2021. [Online]. Available: http://dx.doi.org/10.1039/D1RA07156D

[17] J. A. Aguiar, M. L. Gong, R. R. Unocic, T. Tasdizen, and B. D. Miller, "Decoding crystallography from high-resolution electron imaging and diffraction datasets with deep learning," *Science Advances*, vol. 5, no. 10, p. eaaw1949, 2019. [Online]. Available: https://www.science.org/doi/abs/10.1126/sciadv.aaw1949

[18] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[19] C. W. Park and C. Wolverton, "Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery," *Phys. Rev. Mater.*, vol. 4, p. 063801, Jun 2020. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevMaterials.4.063801

[20] T. Xie and J. C. Grossman, "Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties," *Phys. Rev. Lett.*, vol. 120, p. 145301, Apr 2018. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.120.145301

[21] K. Choudhary and B. DeCost, "Atomistic line graph neural network for improved materials property predictions," *npj Computational Materials*, vol. 7, no. 1, nov 2021. [Online]. Available: https://doi.org/10.1038%2Fs41524-021-00650-1

[22] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations," 2017.

[23] E. CrystalMaker Software Ltd, Oxford, "Crystaldiffract." [Online]. Available: http://crystalmaker.com/crystaldiffract/index.html

[24] A. Khan, C.-H. Lee, P. Y. Huang, and B. K. Clark, "Leveraging generative adversarial networks to create realistic scanning transmission electron microscopy images," *npj Computational Materials*, vol. 9, no. 1, p. 85, 2023.

[25] J. Bals and M. Epple, "Artificial scanning electron microscopy images created by generative adversarial networks from simulated particle assemblies," *Advanced Intelligent Systems*, p. 2300004, 2023.