



# A semi-supervised deep-learning approach for automatic crystal structure classification

Satvik Lolla,<sup>a,b</sup> Haotong Liang,<sup>c</sup> A. Gilad Kusne,<sup>b,c,d</sup> Ichiro Takeuchi<sup>c,e</sup> and William Ratcliff<sup>b,c,\*</sup>

Received 10 November 2021

Accepted 8 June 2022

Edited by T. J. Sato, Tohoku University, Japan

**Keywords:** machine learning; powder neutron diffraction; semi-supervised; indexing.


<sup>a</sup>Poolesville High School, Poolesville, MD 20837, USA, <sup>b</sup>NIST Center for Neutron Research, NIST, Gaithersburg, MD 20899, USA, <sup>c</sup>Department of Materials Science and Engineering, University of Maryland, College Park, MD 20742, USA, <sup>d</sup>Materials Measurement Laboratory, NIST, Gaithersburg, MD 20899, USA, and <sup>e</sup>Maryland Quantum Materials Center, College Park, MD 20742, USA. \*Correspondence e-mail: william.ratcliff@nist.gov

The structural solution problem can be a daunting and time-consuming task. Especially in the presence of impurity phases, current methods, such as indexing, become more unstable. In this work, the novel approach of semi-supervised learning is applied towards the problem of identifying the Bravais lattice and the space group of inorganic crystals. The reported semi-supervised generative deep-learning model can train on both labeled data, *i.e.* diffraction patterns with the associated crystal structure, and unlabeled data, *i.e.* diffraction patterns that lack this information. This approach allows the models to take advantage of the troves of unlabeled data that current supervised learning approaches cannot, which should result in models that can more accurately generalize to real data. In this work, powder diffraction patterns are classified into all 14 Bravais lattices and 144 space groups (the number is limited due to sparse coverage in crystal structure databases), which covers more crystal classes than other studies. The reported models also outperform current deep-learning approaches for both space group and Bravais lattice classification using fewer training data.

## 1. Introduction

The first step towards understanding the properties of a crystalline material at a microscopic level is identifying the crystal structure. However, this is nontrivial. The first part of crystal structure determination is indexing. There are several programs which can be used, such as *DICVOL06* (Boultif & Louër, 1991), *TOPAS* (Coelho, 2018), *GSAS-II* (Toby & Von Dreele, 2013) or *N-TREOR* (Werner *et al.*, 1985; Altomare *et al.*, 2000). These programs output a set of space groups and lattice parameters that could represent the crystal. Using Le Bail (Le Bail *et al.*, 1988) and Pawley (Pawley, 1981) refinements, the space group that fits the diffraction pattern the best can be identified. Rietveld (Rietveld, 1967, 1969) refinement can then be applied to profile the lattice parameters and check the space group. In the presence of impurity phases, this approach becomes more expensive as peaks must be selected manually or tolerance levels must be tuned to discard a certain number of peaks.

One of the approaches for identifying the positions of atoms in a crystal is the charge-flipping algorithm (CFA) (Oszlányi & Sütő, 2008; Palatinus, 2013; Baerlocher *et al.*, 2007). CFA is an iterative approach that relies on fast Fourier transforms to determine the crystal structure of a material (Nussbaumer, 1981; Palatinus & Chapuis, 2007). For CFA, the unit cell and Bravais lattice have to be known, that is, we must have already been successful with some degree of indexing. CFA also cannot handle impurity phases, which are prevalent in many real-world samples.



OPEN ACCESS

Published under a CC BY 4.0 licence

Data science methods are being used increasingly in materials development (Balachandran, 2020; Vandermause *et al.*, 2020; Reyes & Maruyama, 2019; Karigerasi *et al.*, 2018). An example of this is the use of supervised neural networks (NNs) to analyze diffraction patterns. Supervised learning is an approach that seeks to learn a functional mapping between data and their labels. The benefit of NNs is that they, unlike CFA, do not require additional parameters, such as the Bravais lattice or lattice parameters. Although some approaches use NNs to aid in the Rietveld refinement (Ozaki *et al.*, 2020; Chang *et al.*, 2020; Schmidt *et al.*, 2019; Schleder *et al.*, 2019), others use NNs to classify diffraction patterns on the basis of the crystal structure. These classifiers can be trained to identify impurity phases and can be tailored towards specific detectors or parameters. For example, Ryu *et al.* (2019) trained an NN to classify the diffraction patterns of crystals that had defects. Liu *et al.* (2019) used the pair-density function with powder neutron diffraction data for space-group classification. Ziletti *et al.* (2018) used a convolutional neural network to classify simulated single-crystal diffraction X-ray image data into eight space groups.

A number of studies represent powder diffraction patterns as 2D images. However, the information is inherently one dimensional. Previous groups probably used the image approach to take easy advantage of trained models developed by the machine-learning community. Unfortunately, this could introduce more complexities to the model. Garcia-Cardona *et al.* (2019) carried out one of the only studies to examine neutron scattering data and used a 1D approach with simulated powder diffraction data to both differentiate perovskites into five crystal systems and tune the lattice parameters using regression. This study only looked at a small subset of crystals.

A significant challenge with NNs is that they struggle to generalize to new data sets. Most models that predict the space group of a material use less than 100 space groups in their training data set, which limits their application to new diffraction patterns. However, large labeled diffraction data sets are often rare, as labeling them is an expensive task. For this reason, we use a semi-supervised model, which takes advantage of both labeled and unlabeled data during training (Odena, 2016; Zhu & Goldberg, 2009; Kingma *et al.*, 2014; Kipf & Welling, 2016). We employ a generative network that can extract features from the unlabeled data distribution and match these features with the corresponding crystal structure. This allows semi-supervised learning to be used on more data sets, especially ones where labels are not available.

In this study, we propose a 1D semi-supervised model for Bravais lattice and space-group classification using powder neutron diffraction data. Our NNs are trained with data spanning 144 space groups and 14 Bravais lattices. The models used in this study are freely available and can be downloaded (Lolla & Liang, 2021).

## 2. Methods

### 2.1. Data

To test our approach under conditions where we know the correct answer, we worked with simulated data sets. Our data

were taken from the Inorganic Crystal Structure Database (ICSD; <https://www.pds.ac.uk/icsd>), which contains structural information about more than 210 000 crystals (Bergerhoff *et al.*, 1983). A total of 138 362 diffraction patterns were simulated using *TOPAS* (Coelho, 2018). For the Bravais lattices, we combine the rhombohedral and tetragonal classes for a total of 14 classes with ‘F’, ‘I’, ‘P’ and ‘C’ representing the face-centered, body-centered, primitive and base-centered lattices, respectively. We note that there is an inherent class imbalance in the ICSD, as shown in Fig. 1. The most prevalent classes in this data set were the primitive hexagonal, the face-centered cubic and the primitive orthorhombic lattices. The least represented lattices were the face-centered orthorhombic and body-centered orthorhombic lattices.

For the space groups, we used 136 454 of the simulated diffraction patterns. We used only the space groups that had more than 50 diffraction patterns, leaving us with 144 out of the 225 space groups present in the ICSD. The most frequent space group is No. 62 (*Pnma*), which is orthorhombic, and accounts for the disproportionately large number of orthorhombic (P) diffraction patterns in the ICSD data set. A complete list of the 144 space groups used is shown in Table 1.

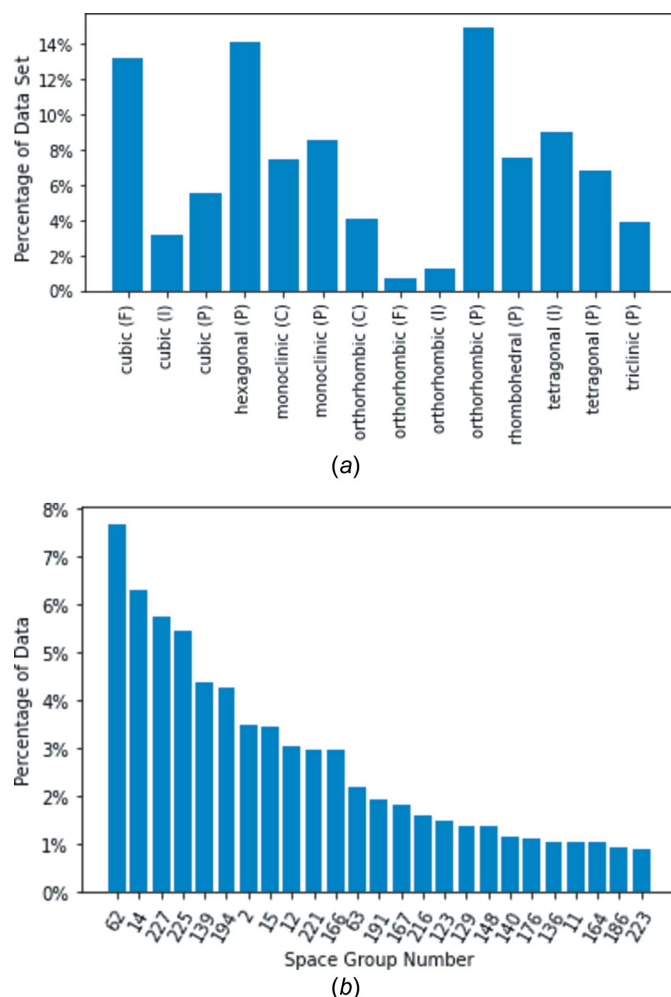


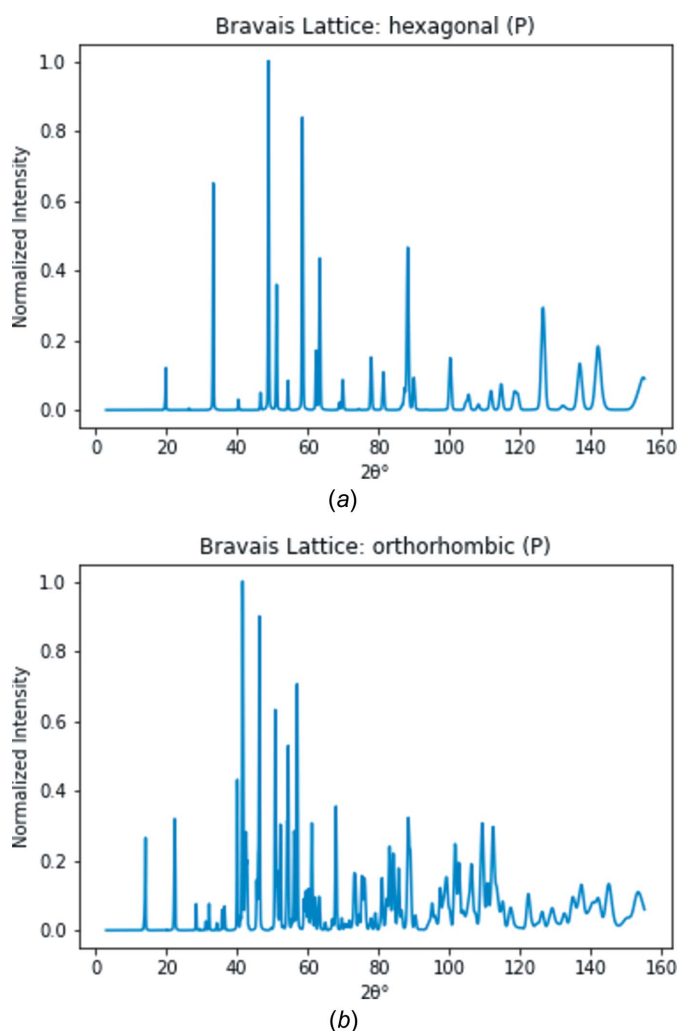
Figure 1 The (a) Bravais lattice and (b) space-group class distribution of the simulated data.

**Table 1**  
List of space groups sorted by crystal system.

Crystal system	Space groups
Triclinic	1, 2
Monoclinic	4–15
Orthorhombic	18–20, 26, 29, 31, 33, 34, 36, 38, 40–47, 51–53, 55–67, 69–74
Tetragonal	82, 85–88, 92, 96, 99, 100, 107, 109, 113, 114, 119, 121, 122, 123, 125–131, 135–137, 139–142
Trigonal	143, 144, 146–148, 150, 152, 154–157, 159–167
Hexagonal	173–176, 180, 182, 185–187, 189–194
Cubic	197–201, 203–206, 212, 213, 215–218, 220, 221, 223–227, 229, 230

A complete list of the ICSD IDs used in this study can be found in the GitHub repository at the URL <https://github.com/usnistgov/semi-supervised-neutron> (Lolla & Liang, 2021).

In this study, we use a 1D approach rather than the traditional 2D image approach. Our data set consists of diffraction patterns of powders. Examples of the 1D diffraction patterns are shown in Fig. 2. To normalize these diffraction patterns, we



**Figure 2**  
Example 1D diffraction patterns for (a) hexagonal (P) crystals and (b) orthorhombic (P) crystals.

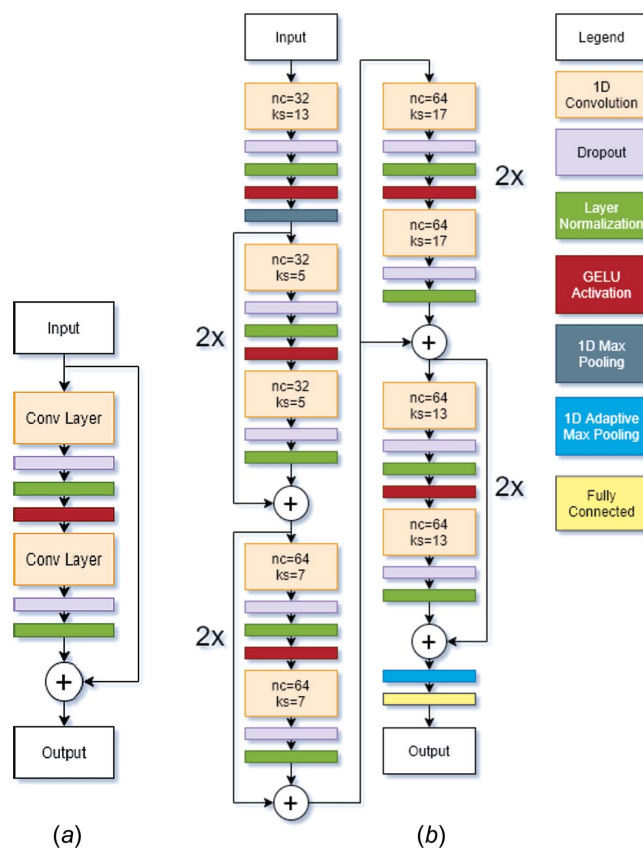
divided all intensities in each diffraction pattern by the maximum intensity. This ensures that the new maximum intensity is equal to 1 and the minimum is equal to 0.

**2.2. Models**

We use two approaches to classify the diffraction patterns: a supervised approach using convolutional neural networks (CNNs) and a semi-supervised approach using a semi-supervised generative adversarial network (SGAN).

**2.3. Supervised model**

We used a 1D ResNet-18, a residual network (He *et al.*, 2016), model to identify the crystal structure of the diffraction patterns. ResNets are examples of CNNs which are commonly used for image classification network algorithms. CNNs consist of convolutional layers, which are responsible for extracting high-level features, such as edges and colors, from images. These layers are used to create a feature map consisting of the most relevant characteristics of the image. To



**Figure 3**  
(a) The ResNet block used in this study and (b) the ResNet-18 architecture. In both parts, the orange, purple, green, red, gray, blue and yellow layers represent 1D Convolutional layers, Dropout layers, Layer Normalization layers, Gaussian error linear unit (GELU) activation functions, 1D Max Pooling layers, 1D Adaptive Max Pooling layers and Fully Connected layers, respectively. The white circles with '+' signs in them represent the addition of two layers. In part (b), each ResNet block was repeated twice, shown by the '2x' next to each block. 'nc' and 'ks' represent the number of channels and the kernel stride for each convolutional layer.

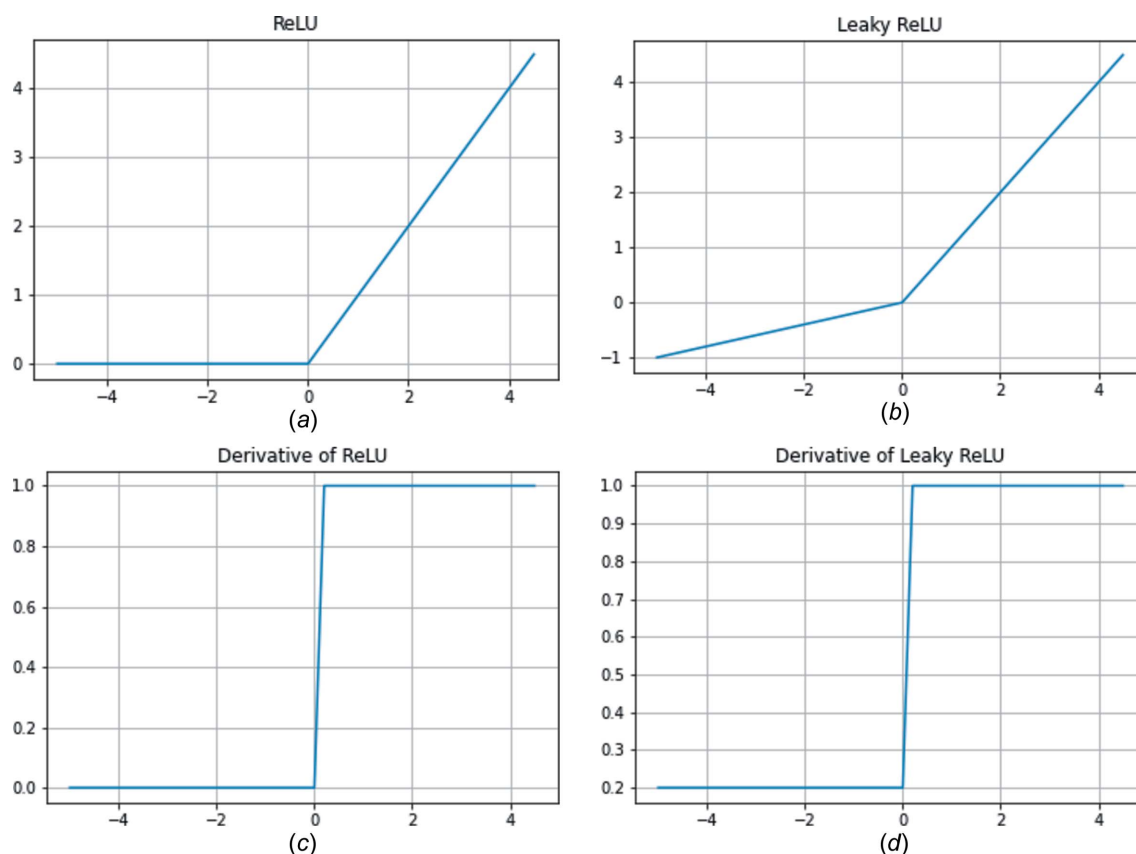
create this map using 1D data, a filter of size  $n$  is applied to a larger sequence with size  $m$ , and the dot product of every  $n$  consecutive values and the filter is computed. This generates a smaller matrix that only includes the relevant features (LeCun & Bengio, 1995).

A ResNet was used in this study to overcome the degradation problem, which occurs when neural networks are too dense so that the accuracy saturates and then quickly degrades (He *et al.*, 2016). ResNets are characterized by their residual blocks, which contain convolutional layers with an identity function. Fig. 3 shows the model architecture used for the ResNet-18, and includes an example of a ResNet block used in this model. During training, we randomly selected 90% of the data to use as the training set and the remaining 10% of the data were used to test the model. This testing data set was distinct from the training one, so the model did not learn from the testing data. These models and the associated training scripts are available on GitHub (Lolla & Liang, 2021).


#### 2.4. Semi-supervised model

We also used an SGAN (Odena, 2016; Goodfellow *et al.*, 2020; Salimans *et al.*, 2016). The SGAN consists of two models: a Generator and a Discriminator. The Generator tries to fool the Discriminator with fake diffraction patterns, while the Discriminator aims to differentiate between real and fake diffraction patterns. The Discriminator also classifies the real labeled data into the corresponding crystal structure class.

**2.4.1. Generator.** The purpose of the Generator is to sample the latent space, a high-dimensional feature space, to generate realistic diffraction patterns. The inputs to the Generator were sampled from a random normal distribution with a mean of 0 and a standard deviation of 1. The Generator consists of 1D Convolutional Transpose layers, 1D Batch Normalization layers and a Leaky rectified linear unit (ReLU) activation function. The Convolutional Transpose layers are used to upsample the data (Radford *et al.*, 2015; Dumoulin & Visin, 2016). The Batch Normalization layers standardize the output of each layer, which reduces error when the model tries to generalize to new inputs (Ioffe & Szegedy, 2015) and has also been shown to reduce mode collapse, a major problem in GANs (Radford *et al.*, 2015). Mode collapse occurs when the Generator only produces a few distinct diffraction patterns despite the latent space input. The Leaky ReLU (Xu *et al.*, 2015) activation with  $\alpha = 0.2$  is used rather than ReLU to reduce the vanishing gradients problem (Radford *et al.*, 2015). Graphs of the ReLU and the Leaky ReLU activation functions are shown in Fig. 4. For negative values, the derivative for the Leaky ReLU function is equal to  $\alpha$ , but for the ReLU function, it is equal to 0. By having a nonzero derivative for all values, the Leaky ReLU is used to combat the sparse gradient problem that occurs while training GANs. Due to our normalization method, which was dividing all values in a diffraction pattern by the maximum intensity, the Discriminator's inputs were in the range from 0 to 1. For this reason, a sigmoid activation was applied to the last layer of our



**Figure 4**  
(a) ReLU, (b) Leaky ReLU with  $\alpha = 0.2$ , (c) derivative of ReLU and (d) derivative of Leaky ReLU with  $\alpha = 0.2$ .



**Figure 5**  
The architecture for the Generator. The blue, green, orange and red boxes represent 1D Convolutional Transpose layers, Leaky ReLU activations, Batch Normalization layers and sigmoid activations, respectively. ‘nc’ and ‘ks’ in the blue boxes represent the number of channels and the kernel size, respectively, for each 1D Convolutional Transpose layer.

Generator, rather than the hyperbolic tangent function recommended by Salimans *et al.* (2016). Fig. 5 shows the model architecture of the Generator.

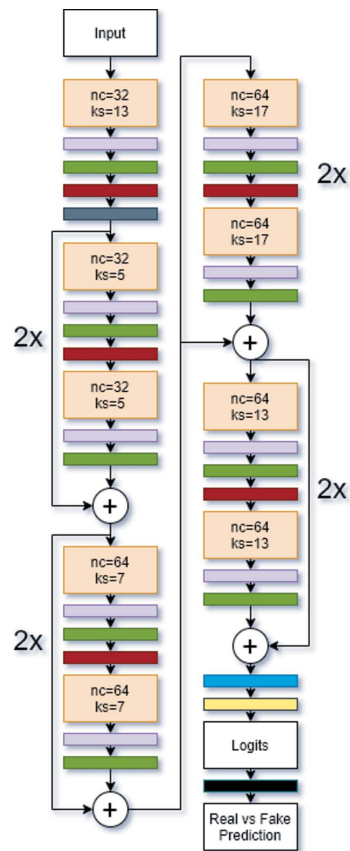
**2.4.2. Discriminator.** The Discriminator has two objectives: to differentiate between real and generated data, and to classify the real data into the correct class. To do this, we used the same 1D ResNet-18 model described in Section 2.3, but an activation function to the last fully connected layer, as proposed by Salimans *et al.* (2016). This activation function is shown in equation (1) and is a version of the softmax activation:

$$D(x) = \frac{Z(x)}{Z(x) + 1}, \tag{1}$$

$$Z(x) = \sum_{k=1}^K \exp[l_k(x)].$$

In this equation,  $l_k(x)$  represents the logit for class  $k$  with data  $x$ . By doing this, we eliminate the need for a second output layer and instead use only the logits from the classification layer. By applying this activation function, diffraction patterns with larger logits, which signify more confident predictions, will be classified as ‘real’, whereas diffraction patterns with smaller logits will be classified as ‘fake’. This encourages the Discriminator to be more confident in its predictions, which sharpens the decision boundary between classes. The Discriminator’s architecture is shown in Fig. 6.

While training the Discriminator, there are two modes: supervised and unsupervised. During unsupervised training, the Discriminator acts the same way it would in a regular GAN as it tries to determine that the generated diffraction patterns are fake and the data drawn from the unsupervised set is real. In the supervised mode, the Discriminator is trained



**Figure 6**  
The model architecture for the ResNet-18 Discriminator. The orange, purple, green, red, gray, blue and yellow layers represent 1D Convolutional layers, Dropout layers, Layer Normalization layers, GELU activation functions, 1D Max Pooling layers, 1D Adaptive Max Pooling layers and Fully Connected layers, respectively. The white circles with ‘+’ signs in them represent the addition of two layers. In part (b), each ResNet block was repeated twice. ‘nc’ and ‘ks’ represent the number of channels and the kernel stride for each convolutional layer, respectively.

to predict the class label for real samples. Training in the unsupervised mode can help the Discriminator extract features from the data, and training on the supervised data will allow the Discriminator to use those extracted features for classification.

**2.4.3. Loss functions and objective functions.** The modified min–max loss proposed by Goodfellow *et al.* (2020) was used for the adversarial loss between the networks. The objective function that the Generator tries to maximize is shown in equation (2):

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m -\log\{D[G(z^{(i)})]\}. \tag{2}$$

$\theta_g$  represents the parameters in the Generator and  $z^{(i)}$  represents the random values in the latent space.  $G(z^{(i)})$  is the generated diffraction pattern from the Generator and  $-\log\{D[G(z^{(i)})]\}$  is the probability that the Discriminator predicts that the generated pattern is real.

For the Discriminator, the objective function to be maximized is shown in equation (3):

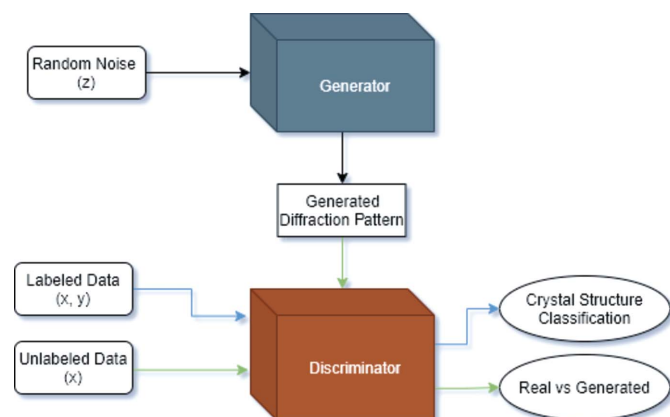


$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m (\log[D(x^{(i)})] + \log\{1 - D[G(z^{(i)})]\}) + \sum_i^C t_i \log(s_i). \quad (3)$$

$\theta_d$  represents the parameters in the Discriminator,  $x^{(i)}$  is the unsupervised real data, and maximizing  $D(x^{(i)})$  implies that the model can identify real data. Like the Generator’s objective function,  $z^{(i)}$  represents the random values in the latent space and  $G(z^{(i)})$  is the generated diffraction pattern from the latent space. Increasing the value of  $1 - \log\{D[G(z^{(i)})]\}$  shows that the Discriminator can determine that the generated patterns are fake. Equation (3) also includes the categorical cross entropy loss, which is shown in the term  $\sum_i^C t_i \log(s_i)$ . Here,  $C$  represents the number of classes,  $t_i$  shows whether the  $i$ th class is the label of the diffraction pattern and  $s_i$  is the Discriminator’s prediction.

**2.4.4. Training details.** Fig. 7 shows the training pipeline used in the SGAN. During SGAN training, the Discriminator has three inputs: a generated sequence from the Generator, a powder diffraction pattern that is labeled with either the space group or the Bravais lattice, and an unlabeled powder diffraction pattern that does not include the crystal structure. We train our SGAN using four different amounts of labeled training data. In all scenarios, we randomly select 10% of the data as testing data, which is distinct from the labeled training data and the unlabeled training data. In the first scenario, we use 5% of the data as labeled training data and 85% as unlabeled training data. In the second, we use 10% of the data as labeled training data and 80% as unlabeled training data. In the third, we use 25% of the data as labeled training data and 65% of the data as unlabeled training data, and finally we use 50% of the data as labeled training data with 40% of the data as unlabeled training data. We also train our supervised ResNet with the same 5, 10, 25 and 50% of the data to compare the accuracy of the SGAN with that of the purely supervised approach.

To train a supervised classifier, we use only the percentage of labeled training data. The model uses a powder diffraction



**Figure 7**  
The architecture for the semi-supervised GAN.

**Table 2**  
Hyperparameters used in the supervised ResNet, Generator and Discriminator.

Hyperparameter	Supervised	Generator	Discriminator
Optimizer	Adam	Adam	Adam
Learning rate	$1 \times 10^{-4}$	$1 \times 10^{-6}$	$5 \times 10^{-6}$
Dropout rate	0.1	None	0.1
Batch size	64	32	32
Nonlinear activations	GELU	Leaky ReLU & sigmoid	GELU

**Table 3**  
Comparison with other supervised space-group classifiers.

Study	Number of space groups	Type of data	Accuracy (%)
Our model	144	Powder diffraction	80
Liu <i>et al.</i> (2019)	45	Pair distribution function	71
Ziletti <i>et al.</i> (2018)	8	2D X-ray diffraction	96.3
Tiong <i>et al.</i> (2020)	72	2D X-ray diffraction	80.2

pattern as input and aims to differentiate between the various crystal structure classes.

Table 2 shows the hyperparameters used in the ResNet and the SGAN.

We used *PYTORCH* (Paszke *et al.*, 2017) as a deep-learning framework. To accelerate training, each model was trained on eight NVIDIA Tesla V100 Tensor Cores.

### 3. Results and discussion

#### 3.1. Supervised model

Our supervised ResNet trained on 90% of the data set had an accuracy of 88%. The confusion matrix for the Bravais lattice model is shown in Fig. 8. By plotting the predicted Bravais lattice against the actual Bravais lattice, the confusion matrix provides more information about the sets of classes that the network misclassified. If the model had a perfect testing accuracy, the values along the principal diagonal would sum to 100% as the network would have classified every diffraction pattern correctly. Again, there is a clear imbalance in the sampled ICSD data set, with orthorhombic (F) and orthorhombic (I) having the least samples. From the confusion matrix, we can see that, despite the fact that orthorhombic (P) is the most prevalent class, the model misclassifies some of these as monoclinic (P) crystals. The network also has trouble differentiating between triclinic (P) and monoclinic (P) diffraction patterns, as both of these classes have low symmetries, agreeing with previous studies (Garcia-Cardona *et al.*, 2019). Similarly to Suzuki *et al.* (2020), we believe that this result was caused by undersampling the triclinic crystals.

For the space-group identification, our model had a top-1 accuracy of 80.6% and a top-5 accuracy of 90.27% across all 144 space groups. We trained our model on all 230 space groups and found that the model had a top-1 accuracy of 74% and a top-5 accuracy of 85%. We decided to investigate further the model on only the 144 most prevalent space groups within




Figure 8 The confusion matrix for the supervised ResNet-18.

the data set, due to a major class imbalance. Some space groups had less than 50 diffraction patterns, less than 0.03% of our data set. Accuracy was measured by dividing the number of correctly classified diffraction patterns in the testing set by the total number of patterns in the testing set. Top-5 accuracy is the percentage of samples for which the actual space group was one of the model’s top five predictions. This outperforms most current models of which we are aware. Liu *et al.* (2019) used machine learning with a pairwise distribution function with a top-1 accuracy of 71% and a top-5 accuracy of 90% across 45 space groups. Tiong *et al.* (2020) classified X-ray diffraction data into 8, 20, 49 and 72 space groups (Table 3). Their accuracy decreased from 99 to 80% for 8 and 72 space groups, respectively, implying that this accuracy would decrease further if their model was trained on more space groups. Aguiar *et al.* (2019) had a top-2 accuracy greater than 80% across all space groups, but used a data set consisting of 650 000 diffraction patterns, more than five times the size of the data set used in this study. However, they used a 1D network, suggesting that a 1D approach can lead to more accurate predictions. We note that we did not take advantage of data augmentation.

### 3.2. Semi-supervised model

We compare the accuracy of the SGAN with the accuracy of the supervised model in Table 4. The SGAN consistently outperforms the purely supervised model, showing that the

semi-supervised approach has the potential to be more applicable in the real world. A graph comparing the accuracy of the supervised and semi-supervised models is shown in Fig. 9. This graph shows that although the accuracy of the SGAN is impacted by a lack of data, the difference between the accuracy of the SGAN and the accuracy of the supervised model is greatest when only 5% of the data are used.

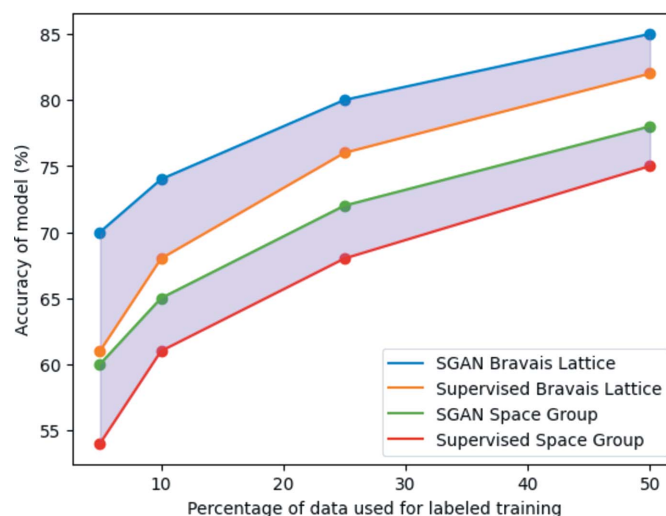


Figure 9 Graph comparing the accuracies of the SGAN and supervised ResNet for Bravais lattice and space-group identification.

Table 4

Comparing the accuracy (%) of the SGAN to the purely supervised approach with different quantities of labeled training data for both Bravais lattice and space-group classification.

Percentage of data	Supervised Bravais lattice	SGAN Bravais lattice	Supervised space group	SGAN space group
5	61	70	54	60
10	68	74	61	65
25	76	80	68	72
50	82	85	75	78

#### 4. Conclusion

In this study, we use both CNNs and a semi-supervised GAN to investigate supervised and semi-supervised approaches for crystal structure classification. We demonstrate that SGANs can prove to be more accurate with limited quantities of labeled data for Bravais lattice and space-group classification. Further, we explore a 1D approach rather than a traditional 2D one. Our 1D model is more accurate than 2D image models, which agrees with previous results in the literature. Our semi-supervised model is also more applicable to real data sets which will lack large quantities of labeled data.

In the future, we would like to train the SGAN to identify impurity phases and to test the method on real data sets.

#### Acknowledgements

The authors are grateful to Austin McDannald and Hui Wu from NIST, and Brian Toby from Argonne National Laboratory. We acknowledge Stephan Rühl for the ICSD. The work at the University of Maryland was supported by NIST. Commercial disclaimer: Any mention of commercial products within this article is for information only; it does not imply recommendation or endorsement by NIST.

#### Funding information

Funding for this research was provided by the National Institute of Standards and Technology (grant No. 60NANB19D027). We also acknowledge support from the Center for High Resolution Neutron Scattering (CHRNS), a national user facility jointly funded by the NCNR and the NSF under agreement No. DMR-2010792.

#### References

- Aguiar, J., Gong, M. L., Unocic, R., Tasdizen, T. & Miller, B. (2019). *Sci. Adv.* **5**, eaaw1949.
- Altomare, A., Giacovazzo, C., Guagliardi, A., Moliterni, A. G. G., Rizzi, R. & Werner, P.-E. (2000). *J. Appl. Cryst.* **33**, 1180–1186.
- Baerlocher, C., McCusker, L. B. & Palatinus, L. (2007). *Z. Kristallogr.* **222**, 47–53.
- Balachandran, P. V. (2020). *MRS Bull.* **45**, 579–586.
- Bergerhoff, G., Hundt, R., Sievers, R. & Brown, I. (1983). *J. Chem. Inf. Comput. Sci.* **23**, 66–69.
- Boultif, A. & Louër, D. (1991). *J. Appl. Cryst.* **24**, 987–993.
- Chang, M.-C., Wei, Y., Chen, W.-R. & Do, C. (2020). *MRS Commun.* **10**, 11–17.
- Coelho, A. A. (2018). *J. Appl. Cryst.* **51**, 210–218.
- Dumoulin, V. & Visin, F. (2016). *arXiv:1603.07285*.

- Garcia-Cardona, C., Kannan, R., Johnston, T., Proffen, T., Page, K. & Seal, S. K. (2019). *IEEE International Conference on Big Data (Big Data)*, pp. 4490–4497. New York: IEEE.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. & Bengio, Y. (2020). *Commun. ACM*, **63**, 139–144.
- He, K., Zhang, X., Ren, S. & Sun, J. (2016). *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778. Los Alamitos: IEEE Computer Society.
- Ioffe, S. & Szegedy, C. (2015). *Proceedings of the 32nd International Conference on Machine Learning*, pp. 448–456. PMLR.
- Karigerasi, M. H., Wagner, L. K. & Shoemaker, D. P. (2018). *Phys. Rev. Mater.* **2**, 094403.
- Kingma, D. P., Mohamed, S., Rezende, D. J. & Welling, M. (2014). *Advances in Neural Information Processing Systems*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence & K. Q. Weinberger, pp. 3581–3589. Curran Associates Inc.
- Kipf, T. N. & Welling, M. (2016). *arXiv:1609.02907*.
- Le Bail, A., Duroy, H. & Fourquet, J. (1988). *Mater. Res. Bull.* **23**, 447–452.
- LeCun, Y. & Bengio, Y. (1995). *The Handbook of Brain Theory and Neural Networks*. Cambridge: MIT Press.
- Liu, C.-H., Tao, Y., Hsu, D., Du, Q. & Billinge, S. J. L. (2019). *Acta Cryst.* **A75**, 633–643.
- Lolla, S. & Liang, H. (2021). *Semi-supervised Neutron*, <https://github.com/usnistgov/semi-supervised-neutron>.
- Nussbaumer, H. J. (1981). *Fast Fourier Transform and Convolution Algorithms*, pp. 80–111. Berlin: Springer.
- Odena, A. (2016). *arXiv:1606.01583*.
- Oszlányi, G. & Sütő, A. (2008). *Acta Cryst.* **A64**, 123–134.
- Ozaki, Y., Suzuki, Y., Hawaii, T., Saito, K., Onishi, M. & Ono, K. (2020). *NPJ Comput. Mater.* **6**, 75.
- Palatinus, L. (2013). *Acta Cryst.* **B69**, 1–16.
- Palatinus, L. & Chapuis, G. (2007). *J. Appl. Cryst.* **40**, 786–790.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L. & Lerer, A. (2017). The Future of Gradient-Based Machine Learning Software and Techniques, NIPS 2017 Autodiff Workshop, Long Beach, California, USA, 9 December 2017. <https://openreview.net/forum?id=BJJsrmfCZ>.
- Pawley, G. S. (1981). *J. Appl. Cryst.* **14**, 357–361.
- Radford, A., Metz, L. & Chintala, S. (2015). *arXiv:1511.06434*.
- Reyes, K. G. & Maruyama, B. (2019). *MRS Bull.* **44**, 530–537.
- Rietveld, H. M. (1967). *Acta Cryst.* **22**, 151–152.
- Rietveld, H. M. (1969). *J. Appl. Cryst.* **2**, 65–71.
- Ryu, D., Jo, Y., Yoo, J., Chang, T., Ahn, D., Kim, Y. S., Kim, G., Min, H.-S. & Park, Y. (2019). *Sci. Rep.* **9**, 15239.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A. & Chen, X. (2016). *Adv. Neural Inf. Process. Syst.* **29**, 2234–2242.
- Schleder, G. R., Padilha, A. C., Acosta, C. M., Costa, M. & Fazzio, A. (2019). *J. Phys. Mater.* **2**, 032001.
- Schmidt, J., Marques, M. R., Botti, S. & Marques, M. A. (2019). *NPJ Comput. Mater.* **5**, 83.
- Suzuki, Y., Hino, H., Hawaii, T., Saito, K., Kotsugi, M. & Ono, K. (2020). *Sci. Rep.* **10**, 21790.
- Tiong, L. C. O., Kim, J., Han, S. S. & Kim, D. (2020). *NPJ Comput. Mater.* **6**, 196.
- Toby, B. H. & Von Dreele, R. B. (2013). *J. Appl. Cryst.* **46**, 544–549.
- Vandermause, J., Torrisi, S. B., Batzner, S., Xie, Y., Sun, L., Kolpak, A. M. & Kozinsky, B. (2020). *NPJ Comput. Mater.* **6**, 20.
- Werner, P.-E., Eriksson, L. & Westdahl, M. (1985). *J. Appl. Cryst.* **18**, 367–370.
- Xu, B., Wang, N., Chen, T. & Li, M. (2015). *arXiv:1505.00853*.
- Zhu, X. & Goldberg, A. B. (2009). *Introduction to Semi-Supervised Learning*, Synthesis Lectures on Artificial Intelligence and Machine Learning 6. Morgan & Claypool.
- Ziletti, A., Kumar, D., Scheffler, M. & Ghiringhelli, L. M. (2018). *Nat. Commun.* **9**, 2775.