

# UC Davis

## UC Davis Previously Published Works

### Title

Deep Learning Segmentation of Complex Features in Atomic-Resolution Phase-Contrast Transmission Electron Microscopy Images

### Permalink

<https://escholarship.org/uc/item/75k71756>

### Journal

Microscopy and Microanalysis, 27(4)

### ISSN

1431-9276

### Authors

Sadre, Robbie  
Ophus, Colin  
Butko, Anastasiia  
[et al.](#)

### Publication Date

2021-08-01

### DOI

10.1017/s1431927621000167

Peer reviewed

# Deep Learning Segmentation of Complex Features in Atomic-Resolution Phase Contrast Transmission Electron Microscopy Images

Robbie Sadre,<sup>1,\*</sup> Colin Ophus,<sup>2,†</sup> Anastasiia Butko,<sup>1</sup> and Gunther H Weber<sup>1</sup>

<sup>1</sup>*Computational Research Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720*

<sup>2</sup>*NCEM, Molecular Foundry, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720*

Phase contrast transmission electron microscopy (TEM) is a powerful tool for imaging the local atomic structure of materials. TEM has been used heavily in studies of defect structures of 2D materials such as monolayer graphene due to its high dose efficiency. However, phase contrast imaging can produce complex nonlinear contrast, even for weakly-scattering samples. It is therefore difficult to develop fully-automated analysis routines for phase contrast TEM studies using conventional image processing tools. For automated analysis of large sample regions of graphene, one of the key problems is segmentation between the structure of interest and unwanted structures such as surface contaminant layers. In this study, we compare the performance of a conventional Bragg filtering method to a deep learning routine based on the U-Net architecture. We show that the deep learning method is more general, simpler to apply in practice, and produces more accurate and robust results than the conventional algorithm. We provide easily-adaptable source code for all results in this paper, and discuss potential applications for deep learning in fully-automated TEM image analysis.

## I. Introduction

High resolution transmission electron microscopy (HRTEM) is a very powerful technique for imaging atomic structure, due to its extremely high spatial resolution. HRTEM has found wide application in studies of the local atomic structure of two-dimensional materials, such as graphene (Meyer et al. 2007, Warner et al. 2009, Meyer et al. 2011, Mas-Balleste et al. 2011, Robertson and Warner 2013, Rasool et al. 2013, 2015). Monolayer graphene is composed of a single 2D sheet of carbon atoms, with the same in-plane structure as the parent material graphite (Cooper et al. 2012). Most synthesis methods that can produce monolayer graphene will also produce defect structures, including point defects (Hashimoto et al. 2004, Jeong et al. 2008, Kotakoski et al. 2014), edges (Russo and Golovchenko 2012, Wang et al. 2014), and line defects, such as grain boundaries (Huang et al. 2011, Yu et al. 2011).

Grain boundaries in graphene are scientifically interesting because of their distinctive mechanical (Rasool et al. 2013, Grantab et al. 2010, Lee et al. 2013), electronic (Jauregui et al. 2011, Tapasztó et al. 2012, Fei et al. 2013), optical (Duong et al. 2012, Podila et al. 2012), and chemical properties (Yasaei et al. 2014, Kim et al. 2014). In a previous study, Ophus et al. (2015) used experimental HRTEM imaging and numerical simulations to map out the parameter space of single-layer graphene grain boundaries, as a function of misorientation and boundary tilt angle. This previous work used semi-automated analysis routines to map out the atomic positions of the boundaries. Once boundary regions were identified, the atomic position analysis was almost entirely automated. However, each of these boundaries had to be hand selected and individually masked, due to the presence of surface contaminants.

These contaminants are likely amorphous carbon (Zhang et al. 2019), which tends to be attracted by the charging induced by the electron beam to the boundary regions. This previous work did not utilize a reliable fully-automated computational method for segmenting between the desirable and undesirable atomic structures.

Recently however, new image analysis methods have been developed under the umbrella of deep learning (Garcia-Garcia et al. 2017). Deep learning as an approach to data processing problems has substantially grown in popularity over the last decade. This can be attributed to increasing availability of large labeled datasets, such as Image-Net (Deng et al. 2009), breakthrough research publications in the field (Krizhevsky et al. 2010), and availability of high performance deep learning frameworks, such as PyTorch (Paszke et al. 2019) and TensorFlow (Abadi et al. 2015). Convolutional Neural Networks (CNNs) have been used for various different image processing tasks, such as image classification (recognition of the object class within an image), object detection (classification and detection of objects in an image as well as generation of the bounding box around the object) and semantic segmentation (pixel-wise classification of an image).

Various works (Ziatdinov et al. 2018, Lee et al. 2019) have successfully applied deep learning methodologies to analyzing atomic defects in microscopy images of materials. In particular, Madsen et al. (2018) used a deep learning network trained on simulated TEM data to recognize local structures in graphene images.

Various studies have made use of neural networks for segmentation of images of cells, such as Akram et al. (2016), Al-Kofahi et al. (2018), as well as other biological datasets, such as vasculature stacks (Teikari et al. 2016), brain tumors (Dong et al. 2017), and neuron structures (Dahmen et al. 2019). Many works have introduced

application specific architectures for their studies, e.g., [Kassim et al. \(2017\)](#), [Roberts et al. \(2019\)](#).

For the segmentation task considered in this paper, we utilize the U-Net architecture as described in [\(Ronneberger et al. 2015\)](#), due to its proven ability to achieve high performance results on image segmentation tasks with limited training data. This aspect is crucial, as large databases of labeled data are typically not readily available for most scientific imaging applications. U-Net has been applied to various datasets, such as urine microscopy images ([Aziz et al. 2018](#)), ADF-STEM images ([Ge and Xin 2018](#)), corneal endothelial cell images ([Daniel et al. 2019](#)), and fluorescently labelled cell nuclei images ([Gudla et al. 2019](#)). Many other works performed similar microscopy segmentation tasks on the nanoscale using modified versions of the U-Net Architecture such as EM-Net ([Khadangi et al. 2020](#)), Fully Residual U-Net ([Gómez-de Mariscal et al. 2019](#)), Inception U-Net ([Punn and Agarwal 2020](#)), and the domain adaptive approach with two coupled U-Nets ([Bermúdez-Chacón et al. 2018](#)).

In this paper, we develop a deep learning-based image segmentation pipeline for detecting surface contaminants in HRTEM images of graphene and compare it to a conventional Bragg filtering approach ([Hýtch 1997](#), [Galindo et al. 2007](#)). Section II reviews materials and methods used in our study. First, we describe image acquisition and preprocessing methodologies as well as labeling training and test data for our modeling approach. We also review Bragg filtering as a classical image segmentation approach for detecting surface contaminants in graphene, which serves as a baseline model. We then introduce our new method that trains and evaluates a U-Net based neural network architecture using k-fold cross validation. We demonstrate that our neural network’s automated feature learning capabilities outperform Bragg filtering for detecting material properties and discuss two potential applications of this segmentation model (Section III). Furthermore, we show how it can be easily used to further automate software based scientific image analysis pipelines. Finally, we summarize results and suggest future extensions and uses (Section ??).

## II. Materials and Methods

We first describe the process through which we grow our graphene samples, how images are extracted, and the different classes of surface structures observed in the data. We then introduce the mathematical definitions for the preprocessing of this acquired data used in this study. We then describe the conventional Bragg filtering method for segmentation. Finally we introduce our deep learning approach to the segmentation task.

### *HRTEM Imaging of Graphene Structures*

The single-layer, polycrystalline graphene samples are grown on polycrystalline copper substrates at 135 °C by chemical vapor deposition. The copper substrate is first held under 150 mTorr of pressure in hydrogen for 1.5 hours, and then 400 mTorr pressure of methane is flowed at 5 standard cubic centimeters per minute (sccm) to form single-layer graphene. Further information of this sample preparation method are given by [Rasool et al. \(2013\)](#), [Li et al. \(2009\)](#), [Rasool et al. \(2011\)](#).

The majority of graphene HRTEM images utilized in the current study are published along with the measured atomic coordinates in a previous study ([Ophus et al. 2015](#)). Some additional HRTEM images, including those from time- and focal-series are from various studies of the structure of graphene grain boundaries ([Rasool et al. 2013, 2014](#), [Ophus et al. 2017](#)) were also included in the image dataset. All of our HRTEM images of graphene were recorded on the TEAM 0.5 microscope, a monochromated and aberration-corrected FEI / Thermo Fisher Titan microscope operated at 80 kV. The imaging conditions are optimized for fast data collection with a relatively low electron dose in order to record as many images as possible. The dose varied from approximately 1 000 to 10 000 electrons / Ångstrom<sup>2</sup> across all images.

Figs. 1a-f show examples of these HRTEM images. Note that due to the monochromation of the electron beam, the intensity of all images varies across the field of view. The graphene samples used in this study contain four primary structural classes. The first class is the graphene lattice itself, which consists of a periodically repeating honeycomb structure. In this structure, each carbon atom is bonded to 3 neighbors with 120° between each bond, and 6 carbon atoms form closed hexagonal rings, which are tiled in a close-packed triangular lattice. Figs. 1g-l show these regions marked in white. Depending on the microscope defocus, the contrast is either white-atom or black-atom, meaning either increased or decreased intensity at the location of each carbon atom respectively ([O’Keefe 2008](#)). Figs. 1b, c, d, and f show examples of white-atom contrast, while Figs. 1a and e show black-atom contrast. These images show varying degrees of residual imaging aberrations; this is a consequence of the low-dose measurement protocol used where the sample is exposed to as little electron fluence as possible.

We can define a simple order parameter calculated by using image convolution to measure the difference in signals between the atoms on a hexagonal ring (using the measured graphene lattice parameter) and the center of the ring. If we calculate this parameter for a range of hexagon orientations, the signal reaches a maximum when the measurement is oriented the same as the underlying lattice, giving an estimate for the lattice




FIG. 1. (a)-(f) HRTEM image examples of polycrystalline, single layer graphene, and (g)-(l) the corresponding labels.

orientation. The regions where two different orientations meet in a discontinuity define the second class, the graphene grain boundaries. These boundaries are shown as red lines in Figs. 1g-l, and were the focus of the previous study by [Ophus et al. \(2015\)](#).

The third class is the vacuum regions, marked as green areas in Figs. 1g-l. No structure is present in these regions, and the electron beam passes straight through with no modulation. Each of these first three classes is easy to detect with simple algorithms. Measuring the location of the vacuum regions is trivial after illumination flat-field correction (described below) since these regions have unit intensity everywhere.

However the fourth class, which corresponds to surface contaminants such as amorphous carbon, is more difficult to accurately segment. These regions are shown as blue areas in Figs. 1g-l. The regions often show strong lattice contrast of ideal or near-ideal graphene structure, overlaid with random modulations. These modulations can be strong or weak, and consist of a complex mix of white-atom and black-atom contrast. Amorphous contaminants also tend to be attracted to the structures we would like to analyze, e.g. the grain boundaries. This is likely due to the surface topology induced by these boundaries ([Ni et al. 2019](#)). In the study of graphene grain boundaries by [Ophus et al. \(2015\)](#), most of the analysis steps were automated, but avoiding these surface contaminant regions was done manually. In this work, we tackle the segmentation of these regions, as it represents the most difficult step to automate.

#### Image Preprocessing

For both the U-Net and Bragg filtering image segmentation, we have applied the same pre-processing and normalization steps, based on the image processing described in [Ophus et al. \(2015\)](#). To normalize the intensity variation due to the monochromation, we have fit the average local intensity  $I_0(\mathbf{r})$  for each image with a 2x2 Bézier surface ([Farin 2001](#)) given by the Equation 1.

$$I_0(\mathbf{r}) = k_{i,j} \sum_{i=0}^m \binom{m}{i} u^i (1-u)^{m-i} \cdot \sum_{j=0}^n \binom{n}{j} v^j (1-v)^{n-j}, \quad (1)$$

where  $(u, v)$  are the image coordinates normalized to range from 0 to 1, and  $k_{i,j}$  are the Bézier surface coefficients. After fitting these coefficients, the normalized intensity  $I(\mathbf{r})$  is given by the Equation 2.

$$I(\mathbf{r}) = \frac{I_{\text{meas}}(\mathbf{r})}{I_0(\mathbf{r})}, \quad (2)$$



where  $\mathbf{r} = (x, y)$  are the real space coordinates, and  $I_{\text{meas}}(\mathbf{r})$  is the measured image intensity. After this step, the mean intensity is equal to one.

Next, we scale the intensity range by calculating the image standard deviation  $\sigma$ , equal to the root-mean-square of the intensity  $\sigma = \sqrt{\langle (I(\mathbf{r}) - 1)^2 \rangle}$ . We then normalize the image by subtracting the mean  $\mu$  and dividing by the standard deviation  $\sigma$  as described by Equation 3.

$$I_{\text{output}} = \frac{I(\mathbf{r}) - \mu}{\sigma} \quad (3)$$

The images in this dataset were originally 1024x1024 or 2048x2048 images. We resize these images down to the size of 256x256 pixels. Training/test labels were generated by hand using the Paint S software application for macOS.

### Segmentation by Fourier Filtering

The defining feature for crystalline samples is their high degree of ordering and long-range translation symmetry. When crystalline samples are imaged along a low index zone axis, the resulting images display local periodicity inside each crystalline grain. These periodic regions create sharply peaked maxima in the image’s 2D Fourier transform amplitude that are closely related to Bragg diffraction from a periodic crystal. These peaks are not strictly speaking due to true Bragg diffraction, but are nevertheless often referred to as “Bragg spots” (Hýtch 1997, Galindo et al. 2007). By applying numerical masks around a given Bragg spot, we measure the degree of local ordering over the image coordinates that corresponds to the associated crystal planes (Pan et al. 1998). We have developed a “Bragg Filtering” procedure to segment the images into two classes, corresponding to clean atomically resolved graphene, and the amorphous surface contaminants. Bragg filtering is a standard procedure in many image processing routines for atomic resolution micrographs such as lattice strain deformation mapping (Hýtch 1997). Our segmentation procedure is shown schematically in Fig. 2.

After preprocessing the initial image and padding the boundaries, we calculate a weighted Fourier transform  $G(\mathbf{q})$  of the image  $I(\mathbf{r})$  shown in Fig. 2a, given by the Equation 4.

$$G(\mathbf{q}) = |\mathbf{q}| |\mathcal{F}_{\mathbf{r} \rightarrow \mathbf{q}} \{I(\mathbf{r})W(\mathbf{r})\}|, \quad (4)$$

where  $\mathbf{r} = (x, y)$  and  $\mathbf{q} = (q_x, q_y)$  are the real space and Fourier space coordinates respectively,  $\mathcal{F}_{\mathbf{r} \rightarrow \mathbf{q}}$  is a 2D Fourier transform from real to Fourier space, and  $W(\mathbf{r})$  is a window function. Next, we find the local maxima

of this image that are above a threshold value  $G_{\text{thresh}}$ , shown in Fig. 2b.

Next, we perform Bragg filtering by applying a 2D Gaussian distribution aperture to  $N$  Bragg peaks at positions  $\mathbf{q}_n$ , given by the Equation 5, where  $\mathcal{F}_{\mathbf{q} \rightarrow \mathbf{r}}$  is the inverse Fourier transform, and  $\sigma$  is the aperture size of the Bragg filter.

$$I_{\text{Bragg}}(\mathbf{r}) = \mathcal{F}_{\mathbf{q} \rightarrow \mathbf{r}} \left\{ \mathcal{F}_{\mathbf{r} \rightarrow \mathbf{q}} \{I(\mathbf{r})\} \sum_{n=1}^N e^{-|\mathbf{q} - \mathbf{q}_n|^2 / 2\sigma^2} \right\}, \quad (5)$$

Note that if symmetric pairs of Bragg diffraction peaks are used, the output image will be real-valued for all pixels. The resulting image is shown in Fig. 2c. By subtracting the Bragg filtered image and the mean intensity from the original image, we generate an image consisting of the non-periodic components, shown in Fig. 2d. Since we consider both negative and positive deviations to be signals from the surface contaminants, we take the absolute value and then low pass filter  $F_{\text{LP}}(\dots)$  the output, giving an image like Fig. 2e. In this figure, we see weak signals in the aperiodic boundary between the two graphene grains, and a strong signal from the contaminants.

Finally, by choosing an appropriate mask threshold  $M_{\text{thresh}}$ , we compute the desired segmentation output  $I_{\text{seg}}(\mathbf{r})$  by the Equation 3. The output is shown in Fig. 2f.

$$I_{\text{seg}}(\mathbf{r}) = F_{\text{LP}} \left[ |I(\mathbf{r}) - I_{\text{Bragg}}(\mathbf{r})| \right] > M_{\text{thresh}} \quad (6)$$

This image compares favorably to the training dataset in Fig. 2g. We accurately mask the contaminant region, while not producing a “false positive” signal at the grain boundary. There are some false positives at the image boundary due to the breakdown of the lattice periodicity at the image boundaries. We use padding and normalization of the filter output to reduce the magnitude of these effects, but they are still present in some images. In practice however, these edge artifacts are insignificant, since we cannot perform accurate measurements of the local atomic neighborhood at the image boundaries.

In this study, we have optimized the parameters  $\sigma$ ,  $G_{\text{thresh}}$ , and  $M_{\text{thresh}}$  by using gradient descent to minimize the error between the segmentation maps generated and the training data. The Bragg filter parameters and performance was measured using five-fold cross validation, giving values of  $\sigma = 0.0156 \pm 0.0005$  1/pixels,  $G_{\text{thresh}} = 12.5 \pm 0.1$ , and  $M_{\text{thresh}} = 0.054 \pm 0.000$ . The maximum number of Bragg peaks included was coarsely tuned, but does not strongly affect the results (as long as it is high enough) and is therefore fixed at 24. The degree of low pass filtering both of the initial Fourier transform for Bragg peak detection,




FIG. 2. Segmenting surface contaminants and graphene lattices using Bragg filtering. (a) Input image data and (b) its Fourier transform amplitude. (c) Inverse Fourier transform after applying Bragg mask, and (d) difference from input image data. (e) Absolute difference smoothed and then (f) threshold segmentation image. (g) Corresponding training data.

and of the output difference image does not strongly affect the results. These two steps were performed by convolution with a 2D Gaussian function with 2 and 5 pixels standard deviation respectively. The cross-validation ensures that the segmentation map accuracy was measured only on the validation subset of the data, using parameters optimized from the other 80% of the training images. This procedure prevents over-fitting of the parameters to the training data. The accuracy of the resulting segmentation images is summarized in Table I.

### Deep Learning Segmentation

The U-Net architecture (Ronneberger et al. 2015) is a CNN architecture based on a fully convolutional neural network, modified and extended to improve segmentation performance for medical and scientific (microscopy) segmentation tasks with limited training data. A fully convolutional neural network (FCN) is a common baseline deep learning architecture for semantic segmentation. It is formed using convolutional layers, pooling layers, nonlinear activation layers, and transposed convolution or upconvolution layers. It generates a network output equal in dimensions to the input, offering a classification for each input pixel. U-Net, which is a network based on the FCN, consists of a down-sampling path and an up-sampling path. Features from the down-sampling path are combined with features generated by the up-sampling path.

Each U-Net block in the down-sampling path consists

of two  $3 \times 3$  kernel convolution layers, a Rectified Linear Unit (ReLU) layer and a  $2 \times 2$  Maximum Pooling operation (MaxPool) layer. Each block in this path doubles the number of features in the previous block. Each block in the up-sampling path starts with an up-convolution operation that doubles the size of the feature maps but also reduces the number of feature by a factor of two. The features from this up-convolution operation are concatenated with the feature maps of the corresponding layer of the down-sampling path, bridged across the network. This is followed by two convolution layers, and a single ReLU layer. The final layer in each U-Net block is a single  $1 \times 1$  convolution layer that translates the final feature maps to two separate output classes, foreground and background.

The original U-Net architecture consists of a total of 9 U-Net blocks. The modified U-Net architecture used in our work is illustrated in Fig. 3. First, we use padded convolution, allowing us to produce a segmentation map with the same size as the input image. In order to build a model optimized to our data, we need to select the best hyperparameters. In the context of deep learning, hyperparameters are configurations set by the developer for a given data modeling problem. These are typically factors such as the learning rate, number of features, and depth of the network. The optimal hyperparameters for a given data modeling problem are not known from the start and are usually optimized through trial and error. In this study, we used the grid search method for hyperparameter tuning. Grid search involves exhaustive trial and error of all combinations

of possible hyperparameters in a search space defined by the programmer. Once each version of the network was trained for 150 epochs, the model with the best performance based on the accuracy score was selected. The selected model was then retrained for an additional 300 epochs for two trials to generate accuracy metrics in order to verify full convergence. We optimized across 3 hyperparameters: number of features in the first block (with each subsequent block having an adjusted number of features to maintain network symmetry), number of blocks, and learning rate. We found that a U-Net with 7 blocks and 32 features in the first block with a learning rate of .0001 achieved the highest accuracy of .9792 and a Jaccard score of .8125. For reference, the original U-Net model achieved an accuracy score of .9775 and a Jaccard score of .8015 using a .0001 learning rate, which yielded the best results for this network architecture as well. An additional advantage of a smaller network is achieving higher throughput which is essential for speed dependant applications such as compression of real time tracking.




FIG. 3. Our Neural Network Architecture based on a scaled down version of U-Net. Each color arrow corresponds to the different operation: 2D Convolution (Conv) with the kernel size of 3x3, Rectified Linear Unit (ReLU), Maximum Pooling (MaxPool) with the window size of 2x2, feature map Copying (Copy), 2x2 Up (Up-conv). Each dark grey box represents a multi-channel feature map resulting from the previous convolution operation, where light grey boxes represent features copied from the down path to the up path. The number of feature maps between each convolution layer is labeled at that top of the box.

We train this network using an Adam optimizer [Kingma and Ba \(2014\)](#) to estimate parameters and pixel-wise cross entropy as a loss function. We train for 300 epochs/approximately 41 minutes on an Nvidia an GeForce GTX TITAN Black GPU, using a learning rate of 0.0001. We use 5 fold cross validation for training. In this scheme, the dataset is split into 5 different groups. For each unique group, the data in that group is excluded, and the model is fit (from scratch) to the remaining 4 groups. Then this model is evaluated for accuracy metrics on the excluded fold in order to have a separate

test set from the training set. at the end, the accuracy metrics from the 5 folds are averaged together. We repeat this experiment twice over the different folds, randomizing the dataset each trial, and then average the performance results for evaluation. We train our model using the PyTorch Deep Learning Framework ([Paszke et al. 2019](#)).

### III. Results and Discussion

In this section, we describe the results of our experiments. We first qualitatively compare the two methods in terms of the types of errors perceived via two separate cases. We then define several key performance metrics that are commonly used for evaluating segmentation models, and report the performance of each model across these different metrics. We also describe how these networks can be used in atomic position workflows and discuss the different results of using each model. Finally, we offer an additional usage of the segmentation model for data reduction and sorting.

#### Qualitative Comparison

Fig. 4 shows the performance of each method on two different example input images. Fig. 4a shows an example of a critical failure by the Bragg filter to detect amorphous carbon. The Bragg filter produces incomplete segmentation results ignoring the bottom section of amorphous carbon, possibly due to its low contrast with respect to the other region of amorphous carbon at the top of the image. On the other hand, U-Net successfully identifies both of these regions as amorphous with precision. The resolution of our segmented boundaries (positional uncertainty of the edges) is roughly given by the feature size of the structures being imaged, approximately equal to the graphene lattice parameter.

Fig. 4b shows an example where both the U-Net and Bragg Filter methods produce precise results; however, the results of the Bragg filter produces a substantial amount of false negatives, where as the U-Net method produces visibly superior results. Furthermore, the bright spot at the right side of the image is incorrectly identified as deformed graphene. U-Net, in contrast, has a much higher coverage of the main defect present in the image, and does not incorrectly classify the bright spot as amorphous. Moreover, the U-Net segmentation detects false positives of amorphous graphene in the upper left hand corner. However, if one inspects the raw image, one will observe that there is in fact a small amount of amorphous graphene present in this part of the image that is overlooked during the labeling process. This demonstrates that the capabilities of deep learning




FIG. 4. Comparison of the performance of the Bragg filter method and the U-Net method for segmentation. Detected errors are visualized on the left. True Positives are shown in white, True Negatives in black, false positives in red and false negatives in blue. The sectioning generated by U-Net is shown on the original image in the middle column, and the sectioning generated by the Bragg filter is shown in the right column.

surpass human performance on some tasks and effectively avoid over-fitting to human labels in certain cases.

#### Quantitative Comparison

Performance Metric	U-Net	Bragg	Difference [%]
Accuracy Score	<b>0.979</b>	0.956	2.4%
Balanced Accuracy Score	<b>0.950</b>	0.897	5.3%
Jaccard Score	<b>0.812</b>	0.688	12.4%
F1 Score	<b>0.884</b>	0.791	9.2%
Precision Score	<b>0.886</b>	0.829	5.6%
Recall Score	<b>0.889</b>	0.792	9.6%

TABLE I. Table of Performance Metrics

Multiple metrics for evaluating segmentations and comparing models exist (Taha and Hanbury 2015, Udupa et al. 2002), where the choice of a particular metric

	$Y' = 1$	$Y' = 0$
$Y = 1$	TP	FN
$Y = 0$	FP	TN

TABLE II. Definitions of True Positive, False Positive, True Negative, and False Negative, where  $Y$  is the ground truth and  $Y'$  is the value predicted by a model.

depends on the problem scope and goals. In the following we consider and define several metrics, discuss their typical usage as well as limitations, and use them to compare our U-Net-based approach to Bragg filtering. Table I summarizes all results, which were computed using SciKit-Learn (Pedregosa et al. 2011). All metrics are defined in terms of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN), see Table II.

One typical score for evaluating a segmentation is *pixel wise accuracy*, calculated as the total number of correctly



classified pixels (TP, TF) divided by the number of predictions (i.e., total number of pixels)

$$\text{PixelwiseAccuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad .$$

Both segmentation methods achieve a pixel wise accuracy above 95%, with U-Net performing roughly 2% better. While this metric is easy to interpret, it fails to properly take large class imbalances into account. In our case, segmentations are dominated by the background class, making the results of this metric potentially skewed. To avoid performance metric inflation due to class imbalance, *balanced accuracy* weighs samples by the inverse prevalence of their true classes. For the binary case, balanced accuracy is computed as

$$\text{BalancedAccuracy} = \frac{1}{2} \left( \frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right) \quad .$$

Using this metric, we start to see the performance of our models to diverge numerically, with a difference of around 5% in favor of U-Net.

The *Jaccard score* and the *F1 score* are two of the most common metrics used for evaluating semantic segmentation models or in the presence of class imbalances. Mathematically, both metrics evaluate the same aspects of model performance (Taha and Hanbury 2015). The Jaccard score is computed as

$$\text{JaccardScore} = \frac{TP}{TP + FP + FN} \quad ,$$

and the F1 score is computed as

$$\text{F1Score} = \frac{2TP}{2TP + FP + FN} \quad .$$

Our U-Net model shows Jaccard and F1 improvements over the Bragg filtering method of approximately 12% and 9% respectively. This further shows that for our class imbalanced dataset, U-Net holds significant advantages over Bragg filtering that are not as clear when using other metrics like accuracy.

To analyze performance differences more specifically in terms of in terms of false positives and false negatives, we compare the *precision* and *recall scores*. *Precision* represents the ratio of correctly predicted positives to total predicted positives, i.e.,

$$\text{Precision} = \frac{TP}{TP + FP} \quad ,$$

while *recall* is the ratio of correctly predicted positives to the total number of positives in the ground truth, i.e.,

$$\text{Recall} = \frac{TP}{TP + FN} \quad .$$

High precision corresponds to a low false positive rate, where as high recall indicates to a low false negative rate. Our study shows that U-Net achieves around 4% and 10% improvement over the Bragg filtering method in precision and recall scores respectively.

## Memory Usage and Speed

Using a workstation running an Intel Xeon 6138 CPU, the Bragg filtering routine requires 0.0203 seconds to process each 256x256 image. The RAM requirements for this routine are approximately 10 times those required for a single image, equal to 5 MB for a 256x256 pixel image. The U-Net implementation requires 0.0294 seconds per image, and uses 710MiB GPU RAM using an NVIDIA Titan Black GPU, for 256x256 size images.

## Segmentation in Atomic Position Workflows

First, we demonstrate the use of Bragg and U-Net segmentation as part of an automated atom and bond finding routine (Fig. 5). After an initial illumination correction step, we determine candidate atom positions using the method described by Ophus et al. (2015). Subsequently, we discard all atom positions that fall inside contamination regions, determined either by using the mask output of the Bragg filtering routine or the U-Net segmentation routines, and remove them from the list of candidates. Finally, we connect neighboring atoms within a given distance threshold by candidate atomic bonds. Fig. 5 shows candidate atom positions as white dots and candidate bonds colored by the bond angle modulo 60°, overlaid on the masks generated by the two segmentation filters. For this paper, we do not perform further any refinement of the atom positions and bonds and focus on segmentation quality.

Fig. 5a shows scenario where both the Bragg filter and U-Net produce suitable results. In both cases, two surface contamination regions are correctly identified and masked, in agreement with the training data. In contrast, Fig. 5b shows an example where the Bragg filter produces a series of false positive regions (type I errors) at the boundary between the two graphene grains. This error can be addressed by decreasing the masking threshold  $M_{\text{thresh}}$ , but this change would overall increase the error across the full dataset. Fig. 5c shows an example of the kind of errors introduced into the Bragg filter segmentation when  $M_{\text{thresh}}$  is set too low. In this example, the Bragg filter failed to mask off a region of surface contamination, which in turn leads to many erroneous atom positions and bonds. These two examples illustrate one weakness of the conventional Bragg filtering routine, which is that it relies on a small number of hyperparameters that cannot be set to values that will successfully perform the segmentation across the full dataset.

In comparison, the U-Net segmentation shown in Fig. 5b and c outperforms the Bragg filter and produces successful results. These examples show that the deep learning approaches can be both more accurate and more robust than simpler, conventional imaging filters,




FIG. 5. Three examples comparing Bragg and U-net filtering. (a) Output masks where both filters performed well, (b) masks where Bragg filter produced false positive regions, and (c) masks where Bragg filter produced false negative regions. In both (b) and (c), the U-Net filter produced a more accurate result.

i.e. Bragg filter. Both segmentation methods produce a false positive at the left edge of Fig. 5c, due to the correction of the low initial intensity value in that region causing boosting of the noise. However, edge pixels are significantly less valuable in analysis of the atomic structure because the neighboring atomic environment is not visible at image boundaries.

#### *Data Reduction and sorting*

Lastly, we demonstrate how our segmentation methods can be applied to a data reduction and image sorting problem. Scientific experiments can yield tens thousands of images which can often have a high variance in terms of information that is useful for the particular experiment being practiced. To be able to sort images by content ratio could be very useful to scientists. In the case of graphene images, a scientist may want to sort images by the ratio of graphene to amorphous carbon. We demonstrate the capability of our segmentation models to perform this task in Fig. 6. Images are sorted based on what percentage is dominated by graphene rather than amorphous carbon. In some cases where frame rates of the data acquisition instrument are too high for

the transference of the data to be done in a reasonable amount of time, it might be useful to compress certain parts of the image data. One way approach to this is the implementation of a run length compression on the content class that may be less important to that specific experiment being carried out.

#### IV. Conclusion

In this study, we have compared two methods to perform segmentation of complex features in phase contrast HRTEM images of monolayer graphene. The two methods we used were a conventional Bragg filtering algorithm, and a deep learning method utilizing the U-Net architecture. The U-Net filter outperformed the conventional method in every performance metric tested, and was very robust against incorrect determination of structurally important regions. The U-Net method has the additional advantage of being adaptable to many different pixel-wise classification problems, and only requires a labeled dataset with a sufficient number of images that contain the desired segmentation features to perform the training. In the future, it may be possible to randomly generate structures and perform image






FIG. 6. A subset of the images from our dataset sorted by the percentage of pixels that contain potentially useful information as computed by U-Net in decreasing order from left to right. Regions classified as amorphous carbon are highlighted in red.

simulation to automatically generate labeled training datasets, removing even this relatively minor barrier. Because of their generality and robustness, deep learning methods such as U-Net segmentation are extremely valuable for fully-automated image processing in TEM.

#### Source Code Availability

The adapted U-Net source code, HRTEM images and the amorphous region labeled images for training are all available at our [DL\\_Segmentation\\_HRTEM Github repo](#)

#### Acknowledgements

This work was supported by the Office of Advanced Scientific Computing Research for the Computational Research Division, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Work at the Molecular Foundry was supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. RS, AB, and GHW acknowledge additional support of the Laboratory-Directed Research and Development project “Network Computing for Experimental and Observational Data Management.” CO acknowledges additional support from a Department of Energy Early Career Research Program award.

\*[rssadre@lbl.gov](mailto:rssadre@lbl.gov)

†[cophus@gmail.com](mailto:cophus@gmail.com)

- J. C. Meyer, A. K. Geim, M. I. Katsnelson, K. S. Novoselov, T. J. Booth, and S. Roth, *Nature* **446**, 60 (2007).
- J. H. Warner, M. H. Rummeli, L. Ge, T. Gemming, B. Montanari, N. M. Harrison, B. Büchner, and G. A. D. Briggs, *Nature nanotechnology* **4**, 500 (2009).
- J. C. Meyer, S. Kurasch, H. J. Park, V. Skakalova, D. Künzel, A. Groß, A. Chuvilin, G. Algara-Siller, S. Roth, T. Iwasaki, *et al.*, *Nature materials* **10**, 209 (2011).
- R. Mas-Balleste, C. Gomez-Navarro, J. Gomez-Herrero, and F. Zamora, *Nanoscale* **3**, 20 (2011).
- A. W. Robertson and J. H. Warner, *Nanoscale* **5**, 4079 (2013).
- H. I. Rasool, C. Ophus, W. S. Klug, A. Zettl, and J. K. Gimzewski, *Nature communications* **4**, 1 (2013).
- H. I. Rasool, C. Ophus, and A. Zettl, *Advanced Materials* **27**, 5771 (2015).
- D. R. Cooper, B. D’Anjou, N. Ghattamaneni, B. Harack, M. Hilke, A. Horth, N. Majlis, M. Massicotte, L. Vandsburger, E. Whiteway, *et al.*, *ISRN Condensed Matter Physics* **2012** (2012).
- A. Hashimoto, K. Suenaga, A. Gloter, K. Urita, and S. Iijima, *Nature* **430**, 870 (2004).
- B. W. Jeong, J. Ihm, and G.-D. Lee, *Physical Review B* **78**, 165403 (2008).

- J. Kotakoski, C. Mangler, and J. C. Meyer, *Nature communications* **5**, 3991 (2014).
- C. J. Russo and J. A. Golovchenko, *Proceedings of the National Academy of Sciences* **109**, 5953 (2012).
- W. L. Wang, E. J. Santos, B. Jiang, E. D. Cubuk, C. Ophus, A. Centeno, A. Pesquera, A. Zurutuza, J. Ciston, R. Westervelt, et al., *Nano letters* **14**, 450 (2014).
- P. Y. Huang, C. S. Ruiz-Vargas, A. M. Van Der Zande, W. S. Whitney, M. P. Levendorf, J. W. Kevek, S. Garg, J. S. Alden, C. J. Hustedt, Y. Zhu, et al., *Nature* **469**, 389 (2011).
- Q. Yu, L. A. Jauregui, W. Wu, R. Colby, J. Tian, Z. Su, H. Cao, Z. Liu, D. Pandey, D. Wei, et al., *Nature materials* **10**, 443 (2011).
- R. Grantab, V. B. Shenoy, and R. S. Ruoff, *Science* **330**, 946 (2010).
- G.-H. Lee, R. C. Cooper, S. J. An, S. Lee, A. Van Der Zande, N. Petrone, A. G. Hammerberg, C. Lee, B. Crawford, W. Oliver, et al., *Science* **340**, 1073 (2013).
- L. A. Jauregui, H. Cao, W. Wu, Q. Yu, and Y. P. Chen, *Solid State Communications* **151**, 1100 (2011).
- L. Tapasztó, P. Nemes-Incze, G. Dobrik, K. Jae Yoo, C. Hwang, and L. P. Biró, *Applied Physics Letters* **100**, 053114 (2012).
- Z. Fei, A. Rodin, W. Gannett, S. Dai, W. Regan, M. Wagner, M. Liu, A. McLeod, G. Dominguez, M. Thiemens, et al., *Nature nanotechnology* **8**, 821 (2013).
- D. L. Duong, G. H. Han, S. M. Lee, F. Gunes, E. S. Kim, S. T. Kim, H. Kim, Q. H. Ta, K. P. So, S. J. Yoon, et al., *Nature* **490**, 235 (2012).
- R. Podila, B. Anand, J. T. Spear, P. Puneet, R. Philip, S. S. S. Sai, and A. M. Rao, *Nanoscale* **4**, 1770 (2012).
- P. Yasaei, B. Kumar, R. Hantehzadeh, M. Kayyalha, A. Baskin, N. Repnin, C. Wang, R. F. Klie, Y. P. Chen, P. Král, et al., *Nature communications* **5**, 1 (2014).
- K. Kim, R. W. Johnson, J. T. Tanskanen, N. Liu, M.-G. Kim, C. Pang, C. Ahn, S. F. Bent, Z. Bao, et al., *Nature communications* **5**, 4781 (2014).
- C. Ophus, A. Shekhawat, H. Rasool, and A. Zettl, *Physical Review B* **92**, 205402 (2015).
- J. Zhang, K. Jia, L. Lin, W. Zhao, H. T. Quang, L. Sun, T. Li, Z. Li, X. Liu, L. Zheng, et al., *Angewandte Chemie International Edition* **58**, 14446 (2019).
- A. Garcia-Garcia, S. Orts-Escolano, S. Oprea, V. Villena-Martinez, and J. Garcia-Rodriguez, *arXiv preprint arXiv:1704.06857* (2017).
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, in *CVPR09* (2009).
- A. Krizhevsky, I. Sutskever, and G. Geoffrey, *Neural Networks* (2010).
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., in *Advances in Neural Information Processing Systems* (2019) pp. 8024–8035.
- M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “**TensorFlow: Large-scale machine learning on heterogeneous systems,**” (2015), software available from tensorflow.org.
- M. Ziatdinov, O. Dyck, S. Jesse, and S. V. Kalinin, *Microscopy and Microanalysis* **24**, 60–61 (2018).
- C.-H. Lee, C. Shi, D. Luo, A. Khan, B. E. Janicek, S. Kang, W. Zhu, B. K. Clark, and P. Y. Huang, *Microscopy and Microanalysis* **25**, 172–173 (2019).
- J. Madsen, P. Liu, J. Kling, J. B. Wagner, T. W. Hansen, O. Winther, and J. Schiøtz, *Advanced Theory and Simulations* **1**, 1800037 (2018).
- S. U. Akram, J. Kannala, L. Eklund, and J. Heikkilä, in *Deep Learning and Data Labeling for Medical Applications* (Springer, 2016) pp. 21–29.
- Y. Al-Kofahi, A. Zaltsman, R. Graves, W. Marshall, and M. Rusu, *BMC bioinformatics* **19**, 1 (2018).
- P. Teikari, M. Santos, C. Poon, and K. Hynynen, *arXiv preprint arXiv:1606.02382* (2016).
- H. Dong, G. Yang, F. Liu, Y. Mo, and Y. Guo, in *annual conference on medical image understanding and analysis* (Springer, 2017) pp. 506–517.
- T. Dahmen, P. Potocek, P. Trampert, M. Peemen, and R. Schoenmakers, *Microscopy and Microanalysis* **25**, 196 (2019).
- Y. M. Kassim, O. V. Glinskii, V. V. Glinsky, V. H. Huxley, and K. Palaniappan, *Microscopy and Microanalysis* **23**, 140 (2017).
- G. Roberts, R. Sainju, B. Hutchinson, M. B. Toloczko, D. J. Edwards, and Y. Zhu, *Microscopy and Microanalysis* **25**, 164 (2019).
- O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” (2015), [arXiv:1505.04597 \[cs.CV\]](https://arxiv.org/abs/1505.04597).
- A. Aziz, H. Pande, B. Cheluvraju, and T. Rai Dastidar, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018) pp. 2230–2238.
- M. Ge and H. L. Xin, *Microscopy and Microanalysis* **24**, 504 (2018).
- M. C. Daniel, L. Atzrodt, F. Bucher, K. Wacker, S. Böhringer, T. Reinhard, and D. Böhringer, *Scientific reports* **9**, 1 (2019).
- P. R. Gudla, G. Zaki, S. Shachar, T. Misteli, and G. Pegoraro, *Microscopy and Microanalysis* **25**, 1376 (2019).
- A. Khadangi, T. Boudier, and V. Rajagopal, *bioRxiv* (2020).



- E. Gómez-de Mariscal, M. Maška, A. Kotrbová, V. Pospíchalová, P. Matula, and A. Muñoz-Barrutia, *Scientific reports* **9**, 1 (2019).
- N. S. Punj and S. Agarwal, *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **16**, 1 (2020).
- R. Bermúdez-Chacón, P. Márquez-Neila, M. Salzmann, and P. Fua, in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)* (IEEE, 2018) pp. 400–404.
- M. Hÿtch, *Scanning Microsc* **11**, 53 (1997).
- P. L. Galindo, S. Kret, A. M. Sanchez, J.-Y. Laval, A. Yanez, J. Pizarro, E. Guerrero, T. Ben, and S. I. Molina, *Ultramicroscopy* **107**, 1186 (2007).
- X. Li, W. Cai, J. An, S. Kim, J. Nah, D. Yang, R. Piner, A. Velamakanni, I. Jung, E. Tutuc, *et al.*, *science* **324**, 1312 (2009).
- H. I. Rasool, E. B. Song, M. J. Allen, J. K. Wassei, R. B. Kaner, K. L. Wang, B. H. Weiller, and J. K. Gimzewski, *Nano letters* **11**, 251 (2011).
- H. I. Rasool, C. Ophus, Z. Zhang, M. F. Crommie, B. I. Yakobson, and A. Zettl, *Nano letters* **14**, 7057 (2014).
- C. Ophus, H. I. Rasool, M. Linck, A. Zettl, and J. Ciston, *Advanced structural and chemical imaging* **2**, 15 (2017).
- M. A. O’Keefe, *Ultramicroscopy* **108**, 196 (2008).
- B. Ni, T. Zhang, J. Li, X. Li, and H. Gao, *Handbook of Graphene, Volume 2: Physics, Chemistry, and Biology*, 1 (2019).
- G. Farin, *Curves and Surfaces for CAGD: A Practical Guide*, 5th ed. (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2001).
- X. Pan, W. D. Kaplan, M. Rühle, and R. E. Newnham, *Journal of the American Ceramic Society* **81**, 597 (1998).
- D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” (2014), [arXiv:1412.6980 \[cs.LG\]](https://arxiv.org/abs/1412.6980).
- A. Taha and A. Hanbury, *BMC Med Imaging* (2015), <https://doi.org/10.1186/s12880-015-0068-x>.
- J. K. Udupa, V. R. LaBlanc, H. Schmidt, C. Imielinska, P. K. Saha, G. J. Grevera, Y. Zhuge, L. M. Currie, P. Molholt, and Y. Jin, in *Medical Imaging 2002: Image Processing*, Vol. 4684, edited by M. Sonka and J. M. Fitzpatrick, International Society for Optics and Photonics (SPIE, 2002) pp. 266 – 277.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Journal of Machine Learning Research* **12**, 2825 (2011).