https://doi.org/10.2352/ISSN.2470-1173.2020.14.COIMG-248

Crystallographic Symmetry for Data Augmentation in Detecting Dendrite Cores

Lan Fu[1], Hongkai Yu[2], Megna Shah[3], Jeff Simmons[3], Song Wang[1]
1 University of South Carolina, Columbia, SC 29208
2 University of Texas - Rio Grande Valley, Edinburg, TX 78539
3 Air Force Research Lab, Dayton, OH 45433
E-mail: lanf@email.sc.edu, hongkaiyu2012@gmail.com, megna.shah.1@us.af.mil, jeff.simmons.3@us.af.mil, and songwang@cec.sc.edu

Abstract

Accurately and rapidly detecting the locations of the cores of large-scale dendrites from 2D sectioned microscopic images helps quantify the microstructure of material components. This provides a critical link between the processing and properties of the material. Such a tool could be a critical part of a quality control procedure for manufacturing these components. In this paper, we propose to use Faster R-CNN, a convolutional neural network (CNN) model that considers both the detection accuracy and computational efficiency, to detect the dendrite cores with complex shapes. However, training CNN models usually requires a large number of images annotated with ground-truth locations of dendrite cores, which are usually obtained by highly labor-intensive manual annotations. In this paper, we leverage the crystallographic symmetry of dendrite cores for data augmentation – the cross sections of dendrite cores show, not perfect, but near four-fold rotation symmetry and we can rotate the image around the center of dendrite cores by specified angles to construct new training data without additional manual annotations. We conduct a series of experiments and the results show the effectiveness of the Faster R-CNN method with the proposed data augmentation strategy. Particularly, we find that we can reduce the number of the manually annotated training images by 75% while still maintaining the same detection accuracy of dendrite cores.

Introduction

Accurately detecting and quantifying microstructural features in manufactured components is a critical part of the digital manufacturing revolution. In materials science, aerospace turbine engines operate at extreme temperatures, so novel approaches have been taken to develop materials that can withstand the heat. The hottest components, the hot section turbine engine blades, are usually produced by directionally solidifying alloy materials with a cooling rate sufficiently slow as to produce single crystal components. This is similar to the process for developing semiconductor chips, except that turbine blades can be cooled faster than semiconductors, which results in a dendritic structure in the final product. Accurate and fast detection of the locations of the cores of large-scale dendrites from 2D sectioned microscopic images helps quantify the microstructure of components so produced, providing a critical link between the processing and properties of the material. Such a tool could be an indispensable part of a quality control procedure for manufacturing these components. In this paper, we investigate the problem of detecting dendrites formed in single crystal castings of Ni-based superalloys, as shown in Fig. 1(b). This problem is of great interest to material scientists since quantifying, and ultimately controlling the microstructure, will allow for both development of new processing routes based on quantitative data and quality control of manufactured parts such as turbine blades.

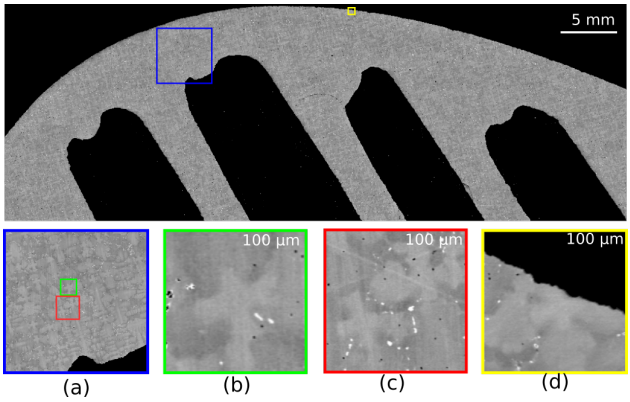

Figure 1. *An illustration of the microstructure of a cross-section crystal nickel-based superalloy turbine blade: top figure: cross section of a turbine blade, smaller figures are as follows. (a) layout of dendrite cores, (b) individual dendrite core, (c) overlap of neighboring dendrite cores, (d) a dendrite core near the edge of the material sample.*

Several complexities make the dendrite detection problem very challenging as is illustrated in Fig. 1. First, the dendrite cores to be detected on each image slice are usually of large scale and high spatial density, as shown in Fig. 1(a). Second, the dendritic microstructure and background are difficult to be separated because of their irregular appearance and diffuse boundaries, as shown in Fig. 1(b). Third, the cross-sections of dendrite microstructures are of complex, non-convex shape and many of them show much spatial overlap when represented by their bounding boxes, as seen in Fig. 1(c). Fourth, the dendrite cores near the edge of material sample pose further difficulty for detection due to different contextual image information, as shown in Fig. 1(d).

Recently, Convolutional Neural Networks (CNNs) have become the state-of-the-art methods for handling general-purpose object detection problems. Based on the CNNs, many object detectors, e.g., Faster R-CNN [9], SSD [6], and YOLO [8], have achieved promising detection performance in many real-world ap-

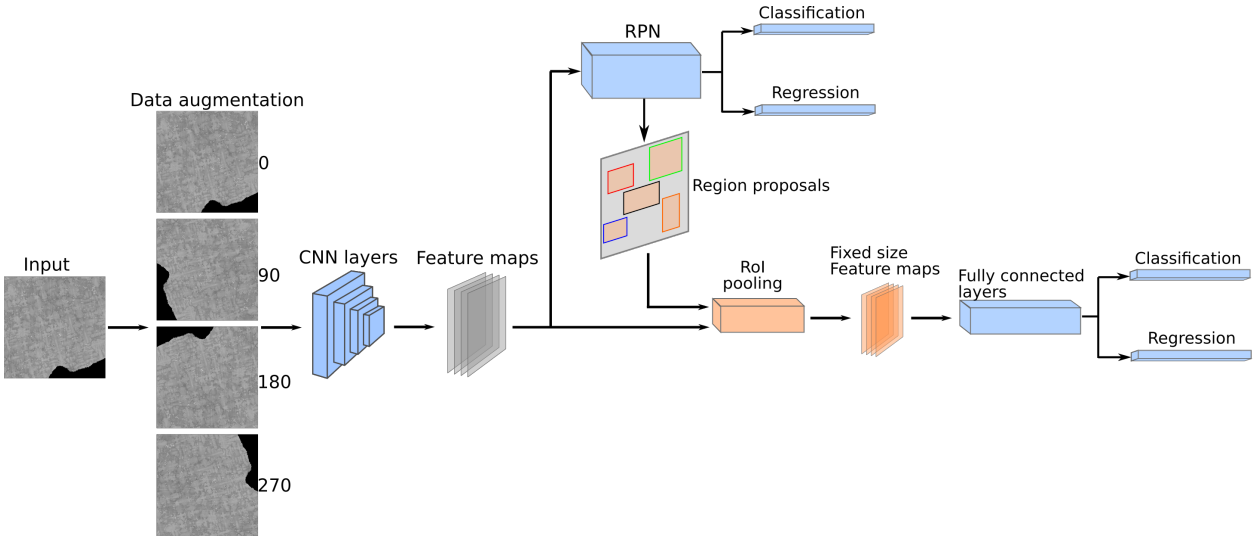

Figure 2. *Dendrite-core detection pipeline based on Faster R-CNN.*

plications. However, these CNN-based detectors usually require a large amount of training data with labeled ground truth. This is particularly undesirable in this work because manually annotating the ground truth is very labor intensive and acquisition of large-scale image data in materials science is usually very costly. In this paper, we investigate new data-augmentation strategies to generate more training samples of dendrite-structure images from a small number of labeled image data. The purpose of this work is to develop a learning-based algorithm that can automatically detect the locations of dendrite cores from 2D serial-sectioned microscopic image slices with high accuracy, eliminating the need for large-scale hand annotations by subject matter experts.

The main contributions of this work are:

1. We proposed to use Faster R-CNN to detect large-scale dendrite cores from 2D cross-sectioned microscopic images of Ni alloy samples.
2. We used the four-fold crystallographic symmetry to augment training data without introducing further burdens of manual annotation.
3. We conducted a series of experiments on real image data to verify the effectiveness of the proposed method.

Related work

CNN-based object detectors. CNN-based object detectors include one-stage detector algorithms, e.g., SSD [6] and YOLO [8], and two-stage detectors, e.g., Faster R-CNN [9] and Mask R-CNN [3]. Although both achieved promising detection performance, two-stage detectors are the state of the art due to their proposal-driven mechanism. However, one of the main problems of these CNN-based methods lies in that they require a large amount of manually annotated data for CNN training, which are generally labor-intensive and hard to collect. For example, there are about 250 dendrite cores in each image slice shown in the top part of Fig. 1. It is time-consuming to annotate these many dendrite cores, given their irregular appearance and complicated background. Insufficient training data, however, generally hurts the generalization ability of CNN models. Data augmentation is an alternative approach to enlarge training dataset for training better CNN models.

Data augmentation. Lack of training data is a widely-existing problem in real-world applications, which renders trained models to suffer from the over-fitting problem. Data augmentation, increasing the amount of training data only utilizing the existing information in training dataset, is an effective way to improve the generalization ability of CNN models in many computer vision tasks [17, 5, 4, 10, 16]. Typical data augmentation technique includes various kinds of image transformations, such as random cropping [5], flipping [10], translation, rotation, adding noises, etc. Random flipping and cropping are the two most effective strategies used as data augmentation for training CNN models. The former intends to randomly flip the original image in horizontal direction, while the latter aims at extracting the subpatch of the input image to construct more training data without any further manual annotations. In this specific dendrite core detection task, we propose a novel data augmentation strategy based on physical and geometry properties of the dendrite structures to train a more robust CNN model.

Dendrite core detection. Accurate detection and characterization of dendrite cores automatically from 2D sectioned microscopic image slices is important for quality control and design of novel materials. A traditional segmentation-based technique [12, 13] that incorporates multiple-scale information and the four-fold symmetry of the dendrites has been developed for dendrite-core detection. Given each pixel in the input image, multi-scale regions centered at this pixel are extracted and then rotated by 90, 180 and 270 degrees, respectively, resulting in four images (the original one together with three rotated ones). These four images are then averaged and the four-fold symmetry filter is obtained by comparing the averaged image with the original image to infer weighting factors, which acts as the dendrite core detector. However, in this method, a threshold parameter needs to be manually set to segment the four-fold symmetry filtered serial images, which is time-consuming and not adaptive. Furthermore, the larger size the selected region around each pixel, the more the required computational time. Different from this tradi-

tional method, our proposed CNN-based method could be adaptive and efficient by utilizing limited ground-truth labeled data for data augmentation.

Method

Faster R-CNN

In this work, we utilize Faster R-CNN [9] as the dendrite core detection pipeline, as shown in Fig. 2. Faster R-CNN, a typical and widely-used two-stage detector based on CNN, has achieved great performance in many detection related tasks [4, 3, 1, 11, 15, 14]. The Faster R-CNN network mainly includes two parts: Region Proposal Network (RPN) and R-CNN (Region-based CNN). RPN learns to tell the CNN model where to look at and outputs a list of region proposals, and the goal of R-CNN is to assign each region proposal a classification score and refine the localization of each region proposal.

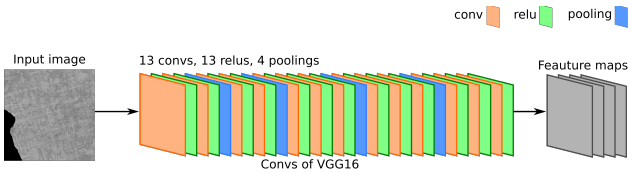

Figure 3. *An illustration of the convolutional layers of VGG-16 [10].*

Specifically, when an image is fed into Faster R-CNN, the backbone VGG-16 [10] is used for feature extraction, as shown in Fig. 3. It involves 13 convolutional layers, 13 Rectified Linear Unit (ReLU) layers and 4 pooling layers in the CNN architecture. Here, convolutional layers for feature extraction are shared by both RPN and R-CNN to improve the computation efficiency. The output feature maps are used for later classification and regression tasks in both RPN and R-CNN. VGG-16 also includes 3 fully connected layers in R-CNN.

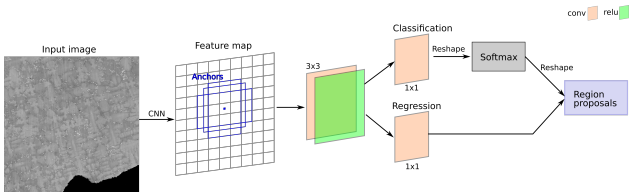

Figure 4. *An illustration of RPN stage.*

In RPN, for each pixel in the input feature map, 9 anchors with different scales (4, 8, 16 pixels) and various aspect ratios (1:1, 1:2, 2:1) are generated for training RPN to obtain region proposals, as shown in Fig. 4. Each anchor over the backbone feature map passes through a 3×3 convolution layer and 1 ReLU layer, and followed by two sub-networks: one 1×1 convolutional layer with 2×9 units for recognizing whether the anchor is the foreground (dendrite core region) or background (non-dendrite core region); another 1×1 convolutional layer with 4×9 units for the bounding box regression. Before obtaining the classification score, the features need to pass through a softmax layer to get the classification probability. Anchors with Intersection-over-Union (IoU) overlaps with corresponding ground-truth bounding box above 0.7 are viewed as positive samples, and those with IoU below 0.3 are viewed as negative samples. In training RPN, we sample 256 anchors (128 as positive and 128 as negative) for each image.

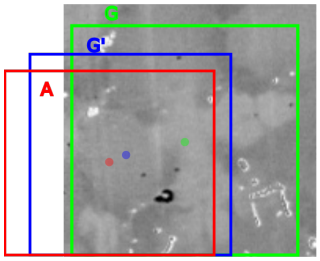

Figure 5. *An illustration of regression from a region proposal to a ground-truth bounding box. A: anchor, G' : region proposal, G: ground-truth bounding box.*

In R-CNN, region proposals pass through Non-Maximum Suppression (NMS) to filter out duplicate proposals regressed to the same object. The first 2,000 proposals per image with high classification scores are kept for later classification and regression tasks. Since fully connected layers can only accept feature maps with fixed size, Region of Interest (RoI) pooling is adopted here to transform different-size feature maps of region proposals to fixed-size feature maps. Then they pass through three fully connected layers for classification and regression, respectively. Regression task helps to refine the localization of region proposals. Since regression either from an anchor to a region proposal or from a region proposal to a ground-truth bounding box is viewed as a kind of linear transformation, as shown in Fig. 5, we introduce general regression offsets from a region proposal to its ground-truth bound box as

$$t_x^* = \frac{1}{w^r}(x^* - x^r) \qquad t_y^* = \frac{1}{h^r}(y^* - y^r) \tag{1}$$

$$t_w^* = \log\frac{w^*}{w^r} \qquad t_h^* = \log\frac{h^*}{h^r} \tag{2}$$

where $[x^r, y^r, w^r, h^r]$ and $[x^*, y^*, w^*, h^*]$ are 2D center location, width, height of the region proposal and the ground-truth bounding box over dendrite core, respectively.

Faster R-CNN [9] mainly includes two kinds of loss functions for training RPN and R-CNN by taking the predictions and the manually labeled ground truth as input. The first one is the classification loss, denoted as L_c, which is used to evaluate the misalignment of classification. It is defined as:

$$L_c = \frac{1}{N_c}\sum_i -(y_i \times \log P_i + (1 - y_i) \times \log(1 - P_i)) \tag{3}$$

where N_c is the number of anchors (256), y_i is its manually labeled ground truth (1 for dendrite core and 0 for non-dendrite-core) and P_i is the probability of the i-th proposal to be a dendrite core. Another one is the regression loss, denoted as L_r, which is used to evaluate the region proposal's localization misalignment. It is defined as:

$$L_r = \frac{1}{N_r}\sum_i y_i \times Smooth_{L_1}(B_i^* - B_i^r) \tag{4}$$

$$Smooth_{L_1}(x) = \begin{cases} 0.5x^2 & if\,|x| < 1 \\ |x| - 0.5 & otherwise, \end{cases} \tag{5}$$

where N_r is the number of proposals (about 2,000), and $Smooth_{L_1}$ [2] is a type of the loss function, B_i^r is the predicted bounding box location (4 parameterized coordinates of the bounding box) of the i-th proposal, B_i^* is the location of corresponding ground truth bounding box .

The total loss function, denoted as L_t, of Faster R-CNN integrates the above two loss functions, which is defined as:

$$L_t = L_c + w \times L_r \tag{6}$$

where w is a balance weight to combine the classification loss L_c and the regression loss L_r. Throughout our experiments, w is set to 1. The whole Faster R-CNN can be trained in an end-to-end way by gradient descent in backpropagation.

Four-fold symmetry-based data augmentation

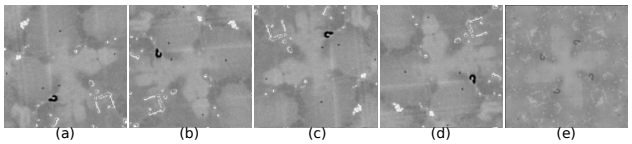

Figure 6. *An illustration of data augmentation by averaging the original and the rotated images of the dendrite structures. (a) an original dendrite core, (b-d) The rotated images of (a) by 90, 180, and 270 degrees, respectively. (e) Average image.*

To address the issue of limited training data (with ground-truth label of the dendrite cores), we augment the training data by considering the physical and geometry properties of the dendrite structures. First, all the dendrite structures for the Ni-based superalloys show near four-fold rotation symmetry [12, 13], as shown in Fig. 6(a). We can rotate each image in training dataset by 90, 180, and 270 degrees, and assume that the ground-truth locations and bounding-boxes of all the present dendrite structures are also rotated accordingly. This way, we can increase the number of training data to a factor of 4 without any further manual annotations. Different from commonly used data augmentation strategy-random rotation, we utilize the four-fold symmetry property as the rotation criterion for data augmentation without additional manual annotations. The proposed four-fold symmetry-based data augmentation method is denoted as DA_{4F} in this work. The statistical distribution of the dendrite structures of Ni-based superalloys shows perfect four-fold rotation symmetry shown in Fig. 6(e), as a consequence of the Curie principle [7]. Based on this assumption, we can take a tight square bounding box of a dendrite structure, rotate it by 90, 180 and 270 degrees respectively and then average all four versions of this dendrite structure to get an average image, as shown in Fig. 6(e). This average image will also exhibit a reasonable dendrite structure if the center of the bounding box is aligned to the center location of dendrite core.

Dataset and experimental setting

Dataset. Given the high resolution of the original image slice of the material sample, $16,257 \times 32,771$ in our case, we first crop it into smaller-size images ($1,000 \times 1,000$) to construct training and testing data sets. In the experiments, we partition a $16,257 \times 32,771$ dendrite-structure image with partial overlap to construct 366 images of size $1,000 \times 1,000$. Among them, we randomly select 166 images, rotate each of them by 90, 180, and 270 degrees, and combine them as the testing image set. We conduct three experiments to show the effectiveness of the data augmentation for dendrite core detection. In the first experiment, we use the 200 remaining images (excluding the 166 test images) for training, without using any further data augmentation. For comparison, we also randomly select 50 images from these 200 images and then rotate them by 90, 180, and 270 degrees as data augmentation for constructing 200 training images. In the second experiment, we further reduce the original training images to 80 and compare it to the case of using 20 original training images augmented with three kinds of rotations. In the third experiment, we take the 200 remaining images used in the first experiment as training dataset and train them with and without DA_{4F} strategy for comparison.

Experimental setting. We use VGG16 [10] as our backbone for feature extraction and the proposed method is implemented based on the Faster R-CNN network in PyTorch. We utilize Stochastic Gradient Descent (SGD) with a weight decay of 0.1 and momentum of 0.9 to optimize all models and set the initial learning rate to 0.001, reduce it by the weight decay after 20 epochs.

The evaluation method is Precision-Recall curve and the Average Precision (AP) metric. Precision measures the proportion of the predictions that are true positives, while recall quantifies the percentage of all the ground-truth dendrites cores that are correctly detected according to predictions. Precision-recall curve presents the trade off between precision and recall for various threshold. Average precision, $\in$ [0,1], measures the area under the precision-recall curve. Higher average precision indicates better detection performance with both higher precision and recall.

Results

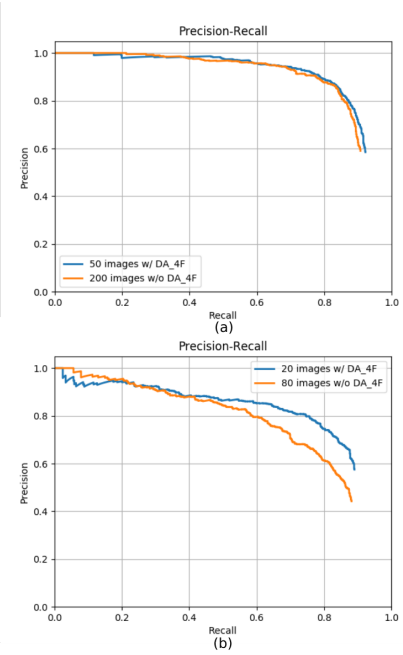

Figure 7. *Precision-Recall curves of the testing results when using (a) 20 and (b) 80 images for network training.*

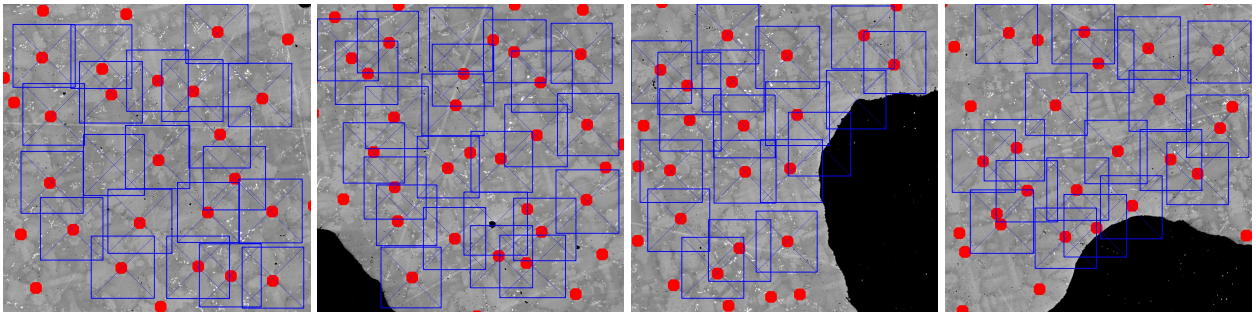

Figure 8. *Sample dendrite-core detection results of training 50 images with DA_{4F}. Blue bounding boxes indicate the detection results while red dots indicate the ground-truth locations of dendrite cores.*

Table 1. Comparison results in terms of average precision.

Experiment 1				*AP*
200	*images*	*w/o*	DA_{4F}	86.24%
50	*images*	*w/*	DA_{4F}	**87.72%**
Experiment 2				***AP***
80	*images*	*w/o*	DA_{4F}	73.43%
20	*images*	*w/*	DA_{4F}	**77.50%**
Experiment 3				***AP***
200	*images*	*w/o*	DA_{4F}	86.24%
200	*images*	*w/*	DA_{4F}	**88.76%**

The comparison results for the first experiment are shown in Fig. 7(a) and Table 1. We can see that only using 50 original images with manual annotations, together with 150 augmented image data which do not require further ground-truth annotations, can train a network with comparable or even better detection performance, with average precision 87.72%, than using all the 200 original training images.

Figure 7(b) and Table 1 show the comparison results for the second experiment. We can see that the proposed data augmentation produces much better performance while reducing the required human annotations by 75%. Specifically, training 20 images with data augmentation achieves performance with average precision 77.50% and converges in a faster speed, while training 80 original images without data augmentation obtains 4% lower average precision 73.43%. Several qualitative detection results are shown in Fig. 8.

We also compare the results of training 200 images with four-fold symmetry-based data augmentation strategy and without any data augmentation, which are shown in Fig. 9 and Table 1. It shows that training with four-fold symmetry-based data augmentation strategy not only achieves better performance with average precision 88.76% but also converges much faster, which verifies the contribution of the four-fold symmetry for CNN models to learn the feature representation of the input images. We also use 360 images for training a CNN model with four-fold symmetry-based data augmentation, and the average precision reaches 94.26%.

Conclusion and future work

In this paper, we proposed to utilize the four-fold crystallographic symmetry to augment training data for CNN-based dendrite core detection from 2D cross-sectioned microscopic images of Ni alloy samples by avoiding further labor-intensive manual annotations. The Faster R-CNN was utilized as the detection pipeline and experiments showed the effectiveness of the proposed data augmentation strategy not only in improving the detection performance but also in improving the detection efficiency. In the future, we intend to utilize the statistical distribution of dendrite cores and the average dendrite image, for further data augmentation.

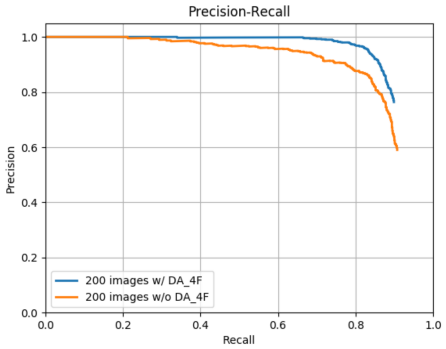

Figure 9. *Precision-Recall curves of the testing results when using 200 images for network training.*

References

[1] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and LucVan Gool. Domain adaptive faster r-cnn for object detection in the wild. InProceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3339–3348, 2018.

[2] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 580–587, 2014.

[3] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In Proceedings of the IEEE International Conference on Computer Vision, pages 2961–2969, 2017.

[4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 770–778, 2016.

[5] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances

in Neural Information Processing Systems, pages 1097–1105, 2012.

[6] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In Proceedings of the European Conference on Computer Vision, pages 21–37. Springer, 2016.

[7] Norihiro Nakamura and Hiroyuki Nagahama. Curie symmetry principle: does it constrain the analysis of structural geology? FORMA-TOKYO-, pages 87-94, 2000.

[8] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In Proceedings of the IEEE International Conference on Computer Vision, pages 779–788, 2016.

[9] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In Advances in Neural Information Processing Systems, pages 91–99, 2015.

[10] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.

[11] Xudong Sun, Pengcheng Wu, and Steven CH Hoi. Face detection using deep learning: An improved faster rcnn approach. Neurocomputing, 299:42–50, 2018.

[12] MA Tschopp, MA Groeber, R Fahringer, JP Simmons, AH Rosenberger, and C Woodward. Symmetry-based automated extraction of microstructural features: Application to dendritic cores in single crystal ni-based superalloys. Scripta Materialia, 62(6):357–360, 2010.

[13] MA Tschopp, MA Groeber, JP Simmons, AH Rosenberger, and C Woodward. Automated extraction of symmetric microstructure features in serial sectioning images. Materials Characterization, 61(12):1406–1417, 2010.

[14] Qing Wang, Xiaodong Zhang, Guanzhou Chen, Fan Dai, Yuanfu Gong, and Kun Zhu. Change detection based on faster r-cnn for high-resolution remote sensing images. Remote sensing letters, 9(10):923–932, 2018.

[15] Tao Wang, Xiaopeng Zhang, Li Yuan, and Jiashi Feng. Few-shot adaptive faster r-cnn. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 7173–7182, 2019.

[16] Sebastien C Wong, Adam Gatt, Victor Stamatescu, and Mark D McDonnell. Understanding data augmentation for classification: when to warp? In 2016 International Conference on Digital Image Computing: Techniques and Applications, pages 1–6. IEEE, 2016.

[17] Yan Xu, Ran Jia, Lili Mou, Ge Li, Yunchuan Chen, Yangyang Lu, and Zhi Jin. Improved relation classification by deep recurrent neural networks with data augmentation. arXiv preprint arXiv:1601.03651, 2016.

Author Biography

Lan Fu *received her B.S. and M.S. from the Department of Biomedical Engineering, Yanshan University, Qinhuangdao, China, in 2011 and Tianjin University, Tianjin, China, in 2014, respectively. She is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, University of South Carolina, SC, USA. Her current research interests include computer vision and machine learning.*

Hongkai Yu *received the B.S. and M.S. degrees from the Department of Automation, Chang'an University, Xi'an, China, in 2009 and 2012, respectively. He is currently an Assistant Professor in the Department of Computer Science at University of Texas - Rio Grande Valley, Edinburg, TX, USA. His current research interests include computer vision, machine learning, and intelligent transportation system.*

Megna Shah *received her B.S. from Ohio State University and her Ph.D. from Northwestern University, both in Materials Science and Engineering. After receiving her Ph.D., she worked as a post-doc at the Air Force Research Lab (AFRL), working to automate 3D serial sectioning. Since then, she spent some time working in industry, including GE Avaition and the RJ Lee Group. She returned to AFRL in November of 2019, where she is interested in working data analysis and collection problems related to digitizing microstructure characterization.*

Jeff Simmons *is a materials and imaging scientist in the Metals Branch, Materials and Manufacturing Directorate of the Air Force Research Laboratory (AFRL). He received the B.S. degree in metallurgical engineering from the New Mexico Institute of Mining and Technology, Socorro, NM, USA, and M.E. and Ph.D. degrees in Metallurgical Engineering and Materials Science and Materials Science and Engineering, respectively, from Carnegie Mellon University, Pittsburgh, PA, USA. After receiving the Ph.D. degree, he began work at AFRL as a post-doctoral research contractor. In 1998, he joined AFRL as a Research Scientist. His research interests are in computational imaging for microscopy and has developed advanced algorithms for analysis of large image datasets. Other research interests have included phase field (physics-based) modeling of microstructure formation, atomistic modeling of defect properties, and computational thermodynamics. He has lead teams developing tools for digital data analysis and computer resource integration and security. He has overseen execution of research contracts on computational materials science, particularly in prediction of machining distortion, materials behavior, and thermodynamic modeling. He has published in the Materials Science, Computer Vision, Signal Processing, and Imaging Science fields. He is a member of ACM and a senior member of IEEE.*

Song Wang *received the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana–Champaign (UIUC), Champaign, IL, USA, in 2002. He was a Research Assistant with the Image Formation and Processing Group, Beckman Institute, UIUC, from 1998 to 2002. In 2002, he joined the Department of Computer Science and Engineering, University of South Carolina, Columbia, SC, USA, where he is currently a Professor. His current research interests include computer vision, image processing, and machine learning. Dr. Wang is a senior member of IEEE and a member of the IEEE Computer Society. He is currently serving as the Publicity/Web Portal Chair for the Technical Committee of Pattern Analysis and Machine Intelligence of the IEEE Computer Society and as an Associate Editor for the IEEE Transactions on Pattern Analysis and Machine Intelligence, Pattern Recognition Letters,and Electronics Letters.*

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

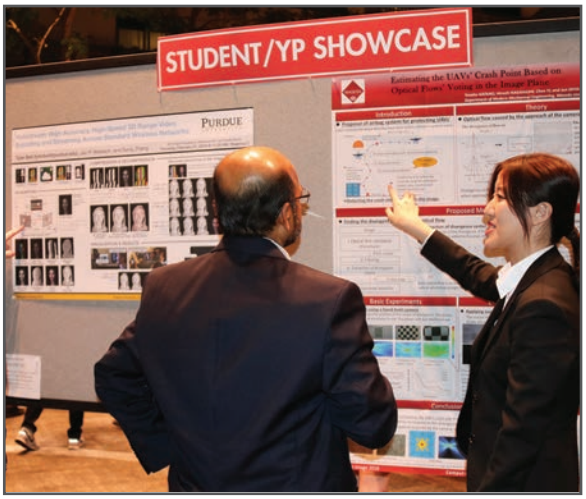

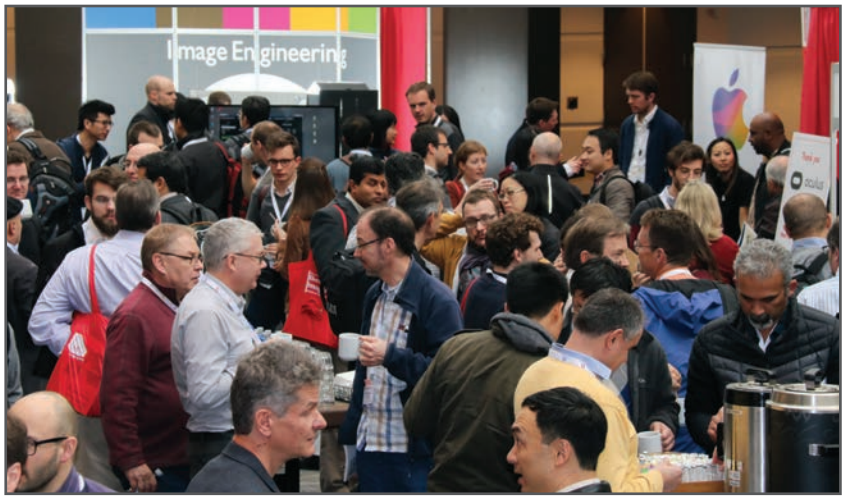

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org