

Layout Hotspot Detection With Feature Tensor Generation and Deep Biased Learning

Haoyu Yang¹, Jing Su, Yi Zou, Yuzhe Ma¹, Bei Yu¹, *Member, IEEE*, and Evangeline F. Y. Young

Abstract—Detecting layout hotspots is a key step in the physical verification flow. Although machine learning solutions show benefits over lithography simulation and pattern matching-based methods, it is still hard to select a proper model for large scale problems and inevitably, performance degradation occurs. To overcome these issues, in this paper, we develop a deep learning framework for high performance and large scale hotspot detection. First, we use feature tensor generation to extract representative layout features that fit well with convolutional neural networks while keeping the spatial relationship of the original layout pattern with minimal information loss. Second, we propose a biased learning (BL) algorithm to train the convolutional neural network to further improve detection accuracy with small false alarm penalties. In addition, to simplify the training procedure and seek a better tradeoff between accuracy and false alarms, we extend the original BL to a batch BL algorithm. Experimental results show that our framework outperforms previous machine learning-based hotspot detectors in both ICCAD 2012 Contest benchmarks and large scale industrial benchmarks. Source code and trained models are available at <https://github.com/pdhyang007/dlhsd>.

Index Terms—Design for manufacturability, feature extraction, lithography hotspot detection, machine learning.

I. INTRODUCTION

AS TRANSISTOR feature size enters the nanometer era, manufacturing yield is drastically affected by lithographic process variations caused by the limitations of the conventional 193-nm wavelength lithography system. Even with various resolution enhancement techniques, manufacturing defects are still likely to happen for some sensitive layout patterns, referred to as hotspots, as shown in Fig. 1. Thus hotspot detection during physical verification stage is very important. The conventional hotspot detection flow includes optical proximity correction (OPC) and lithography simulation on multiple process windows, which



Fig. 1. (a) Hotspot and (b) nonhotspot examples.

has a high accuracy but suffers from runtime overhead issues.

To quickly and correctly recognize hotspots during physical verification, two major methodologies were heavily developed: 1) pattern matching [1]–[3] and 2) machine learning [4]–[9]. In the pattern matching solutions, [1] facilitated the verification flow by integrating the density-based layout encoding, principle components analysis and customized city-block distance. Yu *et al.* [2] developed a design rule-based pattern matching method to recognize hotspots. In [3], an improved tangent space-based distance metric was proposed to perform hotspot pattern analysis and classification. Although pattern matching is a direct and fast method to detect layout characteristics, it has a high error rate for unknown patterns. On the other hand, machine learning techniques are capable of learning hidden relations between layout patterns and their defect characteristics, and can greatly improve detection accuracy. Drmanac *et al.* [9] proposed a histogram-based layout representation and an unsupervised support vector machine (SVM) model to predict the variability of the lithography process. Ding *et al.* [7] incorporated hierarchical artificial neural networks and SVM models to reduce the false alarm rates. Ding *et al.* [8] constructed a meta classifier with basic pattern matching and machine learning classifiers to achieve better detection performance. Yu *et al.* [4] presented a hotspot classification flow with a multikernel SVM and critical feature extraction. In [5], AdaBoost and decision tree are adopted for fast hotspot detection. Very recently, Zhang *et al.* [6] achieved tremendous performance improvements on the ICCAD Contest 2012 benchmark suite [10] by applying an optimized concentric circle sampling (CCS) feature [11] and an online learning scheme. However, there are several aspects that previous works

Manuscript received August 3, 2017; revised March 8, 2018; accepted May 1, 2018. Date of publication May 16, 2018; date of current version May 20, 2019. This work was supported by the Research Grants Council of Hong Kong under Project CUHK24209017. The preliminary version has been presented at the Design Automation Conference (DAC) in 2017. This paper was recommended by Associate Editor L. Behjat. (*Corresponding author: Haoyu Yang.*)

H. Yang, Y. Ma, B. Yu, and E. F. Y. Young are with the Department of Computer Science and Engineering, Chinese University of Hong Kong, Hong Kong (e-mail: hyyang@cse.cuhk.edu.hk).

J. Su and Y. Zou are with the Advanced Technology Group, ASML Brion Inc., San Jose, CA 95131 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCAD.2018.2837078

do not take into account, especially when targeting a very large scale problem size.

- 1) *Scalability*: As integrated circuits develop to an ultralarge scale, very large scale integration layout becomes more and more complicated and traditional machine learning techniques do not satisfy the scalability requirements for printability estimation of a large scale layout. That is, it may be hard for machine learning techniques to correctly model the characteristics of a large amount of layout patterns.
- 2) *Feature Representation*: The state-of-the-art layout feature extraction approaches, including density [5] and CCS [11], inevitably suffer from spatial information loss, because extracted feature elements are flattened into 1-D vectors, and ignore potential spatial relations.

To overcome the conventional machine learning methodology limitations and issues, in this paper we develop a deep learning-based framework targeting high performance and large scale hotspot detection. Because of the automatic feature learning technique and highly nonlinear neural networks, deep learning is highly successful in image classification tasks [12], [13].

Several attempts were made to detect layout hotspots using deep neural networks. Matsunawa *et al.* [14], Shin and Lee [15], Yang *et al.* [16]–[18] demonstrated that ordinary convolutional neural networks can offer promising hotspot detection results. To make the deep neural networks more suitable and efficient for layout hotspot detection tasks, we consider the following two aspects. First, because layout hotspots are related to light diffraction, the input clips of a hotspot detector usually have a resolution over 1000×1000 , which is much larger than the image size (e.g., $\sim 200 \times 200$) in traditional object recognition tasks. As a result, the corresponding neural networks are not computational and storage efficient. Inspired by the feature map in deep neural networks, we utilize the feature tensor concept, i.e., a multidimensional representation of an original layout pattern, that performs compression on input layout images to facilitate learning while keeping the spatial relationship to avoid drastic information loss. Feature tensor compatibility with convolutional neural networks makes our framework more efficient for large scale layout patterns. Second, a good hotspot detector is expected to have a satisfactory tradeoff between detection accuracy and false alarms. Note that in the physical verification flow, a missing hotspot region means a possible failure while a false alarm simply induces additional lithographic simulation time. Therefore, hotspot detection accuracy plays a more important role when balancing the tradeoff. In our preliminary work [19], we developed a bias learning technique that is easily embedded into traditional neural network training procedures and achieve high detection accuracy with a low false alarm penalty. However, the biased learning (BL) requires multiple rounds of end-to-end training and manually determined bias terms, which result in a complex training procedure. To address these concerns, a batch-BL (BBL) algorithm is proposed so that we can train the network with simple back-propagation and single-round mini-batch gradient descent (MGD). Because BBL considers the batch loss of each iteration, we can obtain better

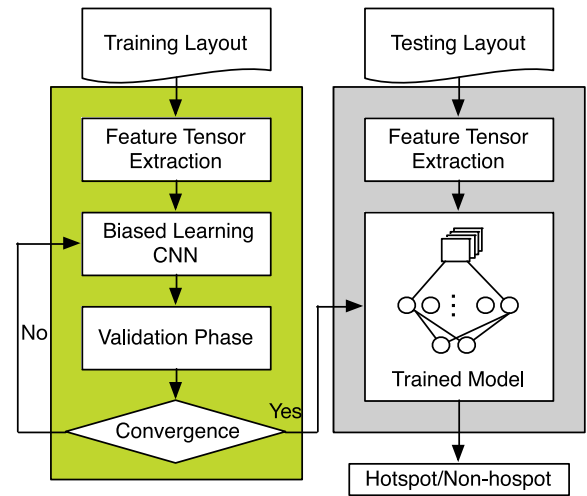


Fig. 2. Proposed hotspot detection flow.

tradeoffs between accuracy and false alarms. The proposed hotspot detection flow is illustrated in Fig. 2.

The main contributions of this paper are listed as follows.

- 1) A feature tensor extraction method is proposed, where the new feature is compatible with the emerging deep learning structure and can dramatically speed up feed-forward and back-propagation.
- 2) A BL technique is embedded into the deep learning framework, that offers significant improvements on hotspot detection accuracy with a minor cost of false alarms.
- 3) We develop a BBL algorithm to adaptively adjust the biased ground truth with respect to the batch training loss that facilitates the training procedure and provides better tradeoffs between accuracy and false alarms.
- 4) Experimental results show that our proposed methods have a great advantage over existing machine learning solutions and achieve a 6.3% accuracy improvements on average for very large scale industrial layouts.

The rest of this paper is organized as follows. Section II introduces basic concepts and problem formulation. Sections III and IV cover feature tensor generation and BL algorithm, respectively. Section V lists the experimental results, followed by the conclusion in Section VI.

II. PRELIMINARIES

In this section, we will introduce some terminologies used in layout hotspot detection. Designed layout patterns are transferred onto silicon wafers through a lithographic process, which involves a lot of variations. Some patterns are sensitive to lithographic process variations and may reduce the manufacturing yield due to potential open or short circuit failures. Layout patterns with a smaller process window and sensitive to process variations are defined as *hotspots*.

The main objectives of the hotspot detection procedure are identifying as many real hotspots as possible, avoiding incorrect predictions on nonhotspot clips, and reducing run-time. In this paper, we use the following metrics to evaluate performance of a hotspot detector.




Fig. 3. Feature tensor generation example ($n = 12$). The original clip ($1200 \times 1200 \text{ nm}^2$) is divided into 12×12 blocks and each block is converted to a 100×100 image representing a $100 \times 100 \text{ nm}^2$ subregion of the original clip. Feature tensor is then obtained by encoding on first k DCT coefficients of each block.

Definition 1 (Accuracy [10]): The ratio between the number of correctly predicted hotspot clips and the number of all real hotspot clips.

Definition 2 (False Alarm [10]): The number of nonhotspot clips that are predicted as hotspots by the classifier.

In the actual design flow, detected hotspots (including false positive patterns) are required to perform lithographic simulation, it is reasonable to account for false alarms the overall estimation flow runtime. Therefore, an overall detection and simulation time (ODST) is defined as follows.

Definition 3 (ODST [6]): The sum of the lithography simulation time for layout patterns detected as hotspots (including real hotspots and false alarms) and the learning model evaluation time.

With the above definitions, we can formulate the hotspot detection problem as follows.

Problem 1 (Hotspot Detection): Given a set of clips consisting of hotspot and nonhotspot patterns, the objective of hotspot detection is training a classifier that can maximize accuracy and minimize false alarms.

III. FEATURE TENSOR EXTRACTION

Finding a good feature representation is a key procedure in image classification tasks, and so is layout pattern classification. Local density extraction and CCS were widely explored in previous hotspot detection and OPC research [11], and were proved to be efficient on hotspot detection tasks because of the embedded lithographic prior knowledge. It is notable that layout hotspots are associated with light diffraction, therefore whether a layout pattern contains hotspots is not only determined by the pattern itself, but is also affected by the surrounding patterns. Therefore, to analyze clip characteristics, we must be aware of the spatial relations of its local regions. However, all of these existing features finally are flattened into 1-D vectors that limit hotspot detection accuracy due to a large amount of spatial information loss.

To address the above issue, we propose a feature tensor extraction method that provides a lower scale representation of the original clips while keeping the spatial information of the clips. After feature tensor extraction, each layout image \mathbf{I} is converted into a hyper-image (image with a customized number of channels) \mathbf{F} with the following properties: 1) size of each channel is much smaller than \mathbf{I} and 2) an approximation of \mathbf{I} can be recovered from \mathbf{F} .

Spectral analysis of mask patterns for wafer clustering was recently explored in [20] and [21] and achieved good clustering performance. Inspired by that work, we express the subregion as a finite combination of different frequency components. High sparsity of the discrete cosine transform (DCT) makes it preferable over other frequency representations in terms of spectral feature extraction, and it is consistent with the expected properties of the feature tensor.

To sum up, the process of feature tensor generation contains the following steps.

Step 1: Divide each layout clip into $n \times n$ subregions, then obtain feature representations of all subregions for multilevel perceptions of layout clips.

Step 2: Convert each subregion of the layout clip $\mathbf{I}_{i,j}$ ($i, j = 0, 1, \dots, n-1$) into a frequency domain

$$\mathbf{D}_{i,j}(m, n) = \sum_{x=0}^{B-1} \sum_{y=0}^{B-1} \mathbf{I}_{i,j}(x, y) \cos \left[\frac{\pi}{B} \left(x + \frac{1}{2} \right) m \right] \times \cos \left[\frac{\pi}{B} \left(y + \frac{1}{2} \right) n \right]$$

where $B = (N/n)$ is subregion size, (x, y) and (m, n) are original layout image and frequency domain indexes, respectively. Particularly, the left-upper side of DCT coefficients in each block correspond to low frequency components, that contain high density information, as depicted in Fig. 3.

Step 3: Flatten $\mathbf{D}_{i,j}$ s into vectors in Zig-Zag form [22] with the larger index being higher frequency coefficients as follows:

$$\mathbf{C}_{i,j}^* = [\mathbf{D}_{i,j}(0, 0), \mathbf{D}_{i,j}(0, 1), \mathbf{D}_{i,j}(1, 0), \dots, \mathbf{D}_{i,j}(B, B)]^T. \quad (1)$$

Step 4: Pick the first $k \ll B \times B$ elements of each $\mathbf{C}_{i,j}^*$

$$\mathbf{C}_{i,j} = \mathbf{C}_{i,j}^*[:k] \quad (2)$$

and combine $\mathbf{C}_{i,j}$, $i, j \in \{0, 1, \dots, n-1\}$ with their spatial relationships unchanged. Finally, the feature tensor is given as follows:

$$\mathbf{F} = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} & \mathbf{C}_{13} & \dots & \mathbf{C}_{1n} \\ \mathbf{C}_{21} & \mathbf{C}_{22} & \mathbf{C}_{23} & \dots & \mathbf{C}_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_{n1} & \mathbf{C}_{n2} & \mathbf{C}_{n3} & \dots & \mathbf{C}_{nn} \end{bmatrix} \quad (3)$$

where $\mathbf{F} \in \mathbb{R}^{n \times n \times k}$. By reversing above procedure, an original clip can be recovered from an extracted feature tensor.

The nature of DCT ensures that high frequency coefficients are near zero. As shown in Fig. 3, large responses only present at the entries with smaller indexes, i.e., low frequency regions. Therefore, most information is kept even when a large amount of elements in $\mathbf{C}_{i,j}^*$ are dropped.

The feature tensor also has the following advantages when applied in neural networks: 1) highly compatible with the data packet transference in convolutional neural networks and 2) forward propagation time is significantly reduced when compared with using an original layout image as input, because the scale of the neural network is reduced with the smaller input size.

IV. TRAINING THE NEURAL NETWORKS

This section discusses the convolutional neural network details. First, we introduce the basis and the architecture of convolutional neural networks. Then, we present a customized training procedure that looks for better tradeoffs between accuracy and false alarms. Finally, we list additional training and testing strategies.

A. Convolutional Neural Network Architecture

To address the weak scalability of traditional machine learning techniques, we introduce the convolutional neural network (CNN) as preferred classifier. CNN is built with several convolution stages and fully connected (FC) layers, where convolution stages perform feature abstraction and FC layers generate the probability of testing instances drawn from each category (Fig. 4).

In this paper, our convolutional neural network has two convolution stages followed by two FC layers, and each convolution stage consists of two convolution layers: 1) an ReLU layer and 2) a max-pooling layer. In each convolution, a set of kernels perform convolution on a tensor \mathbf{F} as follows:

$$\mathbf{F} \otimes \mathbf{K}(j, k) = \sum_{i=1}^c \sum_{m_0=1}^m \sum_{n_0=1}^m \mathbf{F}(i, j - m_0, k - n_0) \mathbf{K}(i, m_0, n_0) \quad (4)$$

where $\mathbf{F} \in \mathbb{R}^{c \times n \times n}$, and kernel $\mathbf{K} \in \mathbb{R}^{c \times m \times m}$. In this paper, the convolution kernel size is set to 3×3 and the numbers of output feature maps in two convolution stages are 16 and 32, respectively. ReLU is an element-wise operation that follows each convolution layer as a replacement of the traditional sigmoid activation function. As shown in (5), ReLU ensures that the network is nonlinear and sparse

$$\text{ReLU}(x) = \begin{cases} x, & \text{if } x > 0 \\ 0, & \text{if } x \leq 0. \end{cases} \quad (5)$$

The max-pooling layer performs 2×2 down-sampling on the output of the previous layer and is applied as the output layer of each convolution stage. Following the two convolution stages are two FC layers with output node numbers of 250 and 2, respectively. A 50% dropout is applied on the first FC layer during training to alleviate overfitting. The second FC layer is the output layer of the entire neural network, where two output nodes generate the predicted probabilities of an input

TABLE I
NEURAL NETWORK CONFIGURATION

Layer	Kernel Size	Stride	Output Node #
conv1-1	3	1	$12 \times 12 \times 16$
conv1-2	3	1	$12 \times 12 \times 16$
maxpooling1	2	2	$6 \times 6 \times 16$
conv2-1	3	1	$6 \times 6 \times 32$
conv2-2	3	1	$6 \times 6 \times 32$
maxpooling2	2	2	$3 \times 3 \times 32$
fc1	-	-	250
fc2	-	-	2

instance being hotspot and nonhotspot. Detailed configurations are shown in Table I.

B. Mini-Batch Gradient Descent

Determining the gradient of each neuron and the parameter updating strategy in the neural network are two key mechanisms in the training procedure. Back-propagation [23] is widely applied to calculate gradients when training large neural networks. Each training instance \mathbf{F} has a corresponding gradient set $\mathcal{G} = \{\mathbf{G}_1, \mathbf{G}_2, \dots, \mathbf{G}_v\}$, where each element is a gradient matrix associated with a specific layer and v is the total layer number. All the neural network parameters are then updated with the obtained \mathcal{G} .

Stochastic gradient descent (SGD), where each training data instance is randomly presented to the machine learning model, has proved more efficient to train large data sets [24] than conventional batch learning, where a complete training set is presented to the model for each iteration. However, as a data set scales to the ultralarge level, e.g., millions of instances, SGD has difficulty to efficiently utilize computation resources. Therefore, it takes a long time for the model to cover every instance in the training set. A compromise approach called MGD [25] can be applied where a group of instances are randomly picked to perform gradient descent. Additionally, MGD is naturally compatible with the online method allowing it to facilitate convergence and avoid large storage requirements for training ultra large instances.

However, for large nonlinear neural networks, back-propagation and MGD do not have rigorous convergence criteria. A fraction, empirically 25%, of training instances (validation set) is separated out and is never shown to the network for weight updating. We then test the trained model on the validation set every few iterations. When the test performance on the validation set does not show much variation or starts getting worse, the training procedure is considered to be converged. To make sure MGD reaches a more granular solution, we reduce the learning rate along with the training process.

The details of MGD with learning rate decay are shown in Algorithm 1, where \mathbf{W} is the neuron weights, λ is the learning rate, $\alpha \in (0, 1)$ is the decay factor, k is the decay step, \mathbf{y}_h^* is the hotspot ground truth, and \mathbf{y}_n^* is the nonhotspot ground truth. MGD can be regarded as a function that returns the model with the best performance on the validation set. Indicator j will count up through iterations (line 4), and in each iteration, m training instances $\{\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_m\}$ are randomly




Fig. 4. Proposed convolutional neural network structure.

Algorithm 1 MGD

```

1: function MGD( $\mathbf{W}$ ,  $\lambda$ ,  $\alpha$ ,  $k$ ,  $\mathbf{y}_h^*$ ,  $\mathbf{y}_n^*$ )
2:   Initialize parameters  $j \leftarrow 0$ ,  $\mathbf{W} > \mathbf{0}$ ;
3:   while not stop condition do
4:      $j \leftarrow j + 1$ ;
5:     Sample  $m$  training instances  $\{\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_m\}$ ;
6:     for  $i \leftarrow 1, 2, \dots, m$  do
7:        $\mathcal{G}_i \leftarrow \text{backprop}(\mathbf{F}_i)$ ;
8:     end for
9:     Calculate gradient  $\bar{\mathcal{G}} \leftarrow \frac{1}{m} \sum_{i=1}^m \mathcal{G}_i$ ;
10:    Update weight  $\mathbf{W} \leftarrow \mathbf{W} - \lambda \bar{\mathcal{G}}$ ;
11:    if  $j \bmod k = 0$  then
12:       $\lambda \leftarrow \alpha \lambda$ ,  $j \leftarrow 0$ ;
13:    end if
14:  end while
15:  return Trained model  $f$ ;
16: end function
    
```

sampled from the training set (line 5). Gradients of these training instances (\mathcal{G}_i) are calculated using back-propagation (lines 6–8). Then neuron weights \mathbf{W} are updated by subtracting the average gradient of sampled instances $\lambda \bar{\mathcal{G}}$ scaled by learning rate γ (line 14). When j is an integer multiple of k , λ is reduced to $\alpha \lambda$, i.e., the learning rate decays every k iterations (lines 11–13). At the end of MGD, a trained model that has satisfactory performance on validation set will be returned (line 15).

Although SGD has shown an advantage in emerging machine learning techniques, it cannot fully utilize GPU resources. MGD, on the other hand, is more compatible with parallel computing and can speed up training procedures. To evaluate the efficiency of MGD, we train the neural network on the ICCAD benchmark using MGD and SGD separately with the configuration in Table III. The training procedure is shown in Fig. 5, where the x -axis is the elapsed time (s) in the training procedure and the y -axis is the cross-entropy loss on the validation set. The curve shows that MGD behaves much more stably than SGD, which indicates the neural network with the MGD learning strategy is more efficient and effective than conventional SGD.

C. Learning Toward Biased Target

Softmax cross entropy can provide speedup for back-propagation while attaining comparable network performance with the mean square error [25]. In a n -category classification task, the instance that belongs to class c has a ground

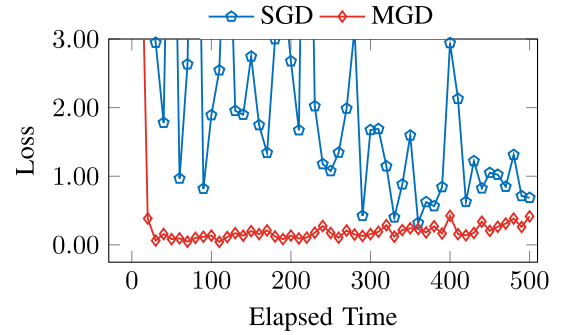


Fig. 5. SGD versus MGD.

truth $\mathbf{y}^* \in \mathbb{R}^n$ where \mathbf{y}^* has the property that $\mathbf{y}^*(c) = 1$ and $\sum_{i=1}^n \mathbf{y}^*(i) = 1$. Each entry of \mathbf{y}^* is regarded as the probability of the instance drawn from each category. The predicted label vector \mathbf{y} by classifier is defined similarly.

In the task of hotspot detection, $\mathbf{y}^* = \mathbf{y}_n^* = [1, 0]$ and $\mathbf{y}^* = \mathbf{y}_h^* = [0, 1]$ are assigned as the ground truths for nonhotspot and hotspot. To generate loss with respect to ground truth, score $\mathbf{x} = [x_h, x_n]$ predicted by the neural network is scaled to a $(0, 1)$ interval by the softmax function shown in

$$\mathbf{y}(0) = \frac{\exp x_h}{\exp x_h + \exp x_n}, \quad \mathbf{y}(1) = \frac{\exp x_n}{\exp x_h + \exp x_n} \quad (6)$$

and then, cross-entropy loss is calculated as follows:

$$l(\mathbf{y}, \mathbf{y}^*) = -(\mathbf{y}^*(0) \log \mathbf{y}(0) + \mathbf{y}^*(1) \log \mathbf{y}(1)). \quad (7)$$

In case of the situation we need to calculate $\log 0$, we define

$$\lim_{x \rightarrow 0} x \log x = 0. \quad (8)$$

Because each entry of softmax label \mathbf{y}_i is the probability of given instance \mathbf{F}_i being nonhotspot \mathcal{N} and hotspot \mathcal{H} , we have

$$\mathbf{F} \in \begin{cases} \mathcal{N}, & \text{if } \mathbf{y}(0) > 0.5 \\ \mathcal{H}, & \text{if } \mathbf{y}(1) > 0.5. \end{cases} \quad (9)$$

$$\mathbf{y}(0) + \mathbf{y}(1) = 1. \quad (10)$$

To improve the hotspot detection accuracy, a straightforward approach is shifting the decision boundary, as shown

$$\mathbf{F} \in \begin{cases} \mathcal{N}, & \text{if } \mathbf{y}(0) > 0.5 + \lambda \\ \mathcal{H}, & \text{if } \mathbf{y}(1) > 0.5 - \lambda \end{cases} \quad (11)$$

where $\lambda > 0$ is the shifting level. However, this method can cause a large increase in false alarms.




Fig. 6. Training loss versus accuracy. The loss in (a) and (b) are calculated with respect to the hotspot ground truth. The loss in (c) and (d) are calculated with respect to the nonhotspot ground truth. An optimal loss does not guarantee a better discriminant performance: (a) illustrates a smaller loss than (b) but suffers a lower hotspot prediction accuracy (75% versus 100%); (c) illustrates a smaller loss than (b) but suffers a higher false alarm (1 versus 0).

The conventional training procedure applies ground truth label $\mathbf{y}_n^* = [1, 0]$ for nonhotspot instances and $\mathbf{y}_h^* = [0, 1]$ for hotspot instances. For non hotspot instances, the classifier is trained toward \mathbf{y}_n^* . If the training procedure meets the stop criteria, then for most nonhotspot clips, f will predict them to have a high probability, close to 1, to be nonhotspots. However, as can be seen in (9), the instance would be predicted as a nonhotspot as long as the predicted probability is greater than 0.5. Thus, to some extent, the classifier is too confident as expected.

Intuitively, a too confident classifier is not necessary to give a good prediction performance and on the contrary, may induce more training pressure or even overfitting problem. We exemplify the case using a linear classifier. As illustrated in Fig. 6, a more confident classifier results in an optimized loss, but cannot guarantee higher classification accuracy. Therefore, an assumption can be made that the hotspot detection accuracy can be further improved by sacrificing the training loss of non-hotspot samples. Meanwhile, the induced false alarm penalties are expected to be lower than directly shifting the decision boundary.

Assumption 1: Given a pretrained convolutional neural network model with ground truth $\mathbf{y}_n^* = [1, 0]$ and $\mathbf{y}_h^* = [0, 1]$ and hotspot detection accuracy a on a given test set. Fine-tune the network with $\mathbf{y}_n^\epsilon = [1 - \epsilon, \epsilon]$, $\epsilon \in [0, 0.5)$, we can obtain the hotspot detection accuracy a' and false alarm fs' of the new model. Shifting the decision boundary of the original model to reach the detection accuracy a' , the corresponding false alarm is fs'' . We have $a' \geq a$ and $fs'' \geq fs'$.

Because it is hard to have a solid proof of the above assumption due to the uncertainty of the deep neural networks, we conduct a sketch explanation of $a' \geq a$ by analyzing the training actions of the FC layer.

Proof: Consider a trained classifier f with $O_{l-1}(\mathbf{F}_i)$ as the output of the second last layer, and the neurons have weight \mathbf{W}_l . Then the output can be expressed as follows:

$$\mathbf{x}_i = \mathbf{W}_l^T O_{l-1}(\mathbf{F}_i) \quad (12)$$

where \mathbf{W}_l is learned toward the target $\mathbf{y}_n^* = [1, 0]$. We will show in (29)–(32) that training gradients of nonhotspot instances with biased labels are smaller than those without bias. Considering the fine-tune process with $\mathbf{y}_n^\epsilon = [1 - \epsilon, \epsilon]$, training instances with predicted probability less than $1 - \epsilon$

are only supposed to generate minor gradient to update the network, since they can not even make prominent difference with the target \mathbf{y}_n when the stopping criteria is met. For those nonhotspot instances that have predicted probability in $(1 - \epsilon, 1)$ (confident instances), neuron weights are updated along the gradient generated by confident instances.

Also, gradient vanishing theory [26] indicates a later layer in the neural network learns faster, therefore within a limited number of iterations, updates of the front layer can be ignored. We assume the layers before O_{l-1} are fixed for some iterations. Let the updated weight for the output layer be \mathbf{W}'_l , and the current output for confident instance \mathbf{F}_c is

$$\mathbf{x}'_c = \mathbf{W}'_l{}^T O_{l-1}(\mathbf{F}_c). \quad (13)$$

Before adjusting the ground truth, we have

$$\mathbf{x}_c = \mathbf{W}_l^T O_{l-1}(\mathbf{F}_c). \quad (14)$$

Note that $\mathbf{x} \in \mathbb{R}^2$, therefore \mathbf{W}_l has two columns $\mathbf{W}_{l,1}$ and $\mathbf{W}_{l,2}$. Similarly, $\mathbf{W}'_l = [\mathbf{W}'_{l,1}, \mathbf{W}'_{l,2}]$. We define \mathbf{w} and \mathbf{w}' as follows:

$$\mathbf{w} = \mathbf{W}_{l,1} - \mathbf{W}_{l,2}, \quad \mathbf{w}' = \mathbf{W}'_{l,1} - \mathbf{W}'_{l,2}. \quad (15)$$

Here \mathbf{w}' is updated from \mathbf{w} through gradient descent

$$\begin{aligned} \mathbf{w}' &= \mathbf{w} - \alpha \nabla_{\mathbf{w}} (\mathbf{x}_c(0) - \mathbf{x}_c(1)) (\nabla_{\mathbf{x}_c(0)} L_c - \nabla_{\mathbf{x}_c(1)} L_c) \\ &= \mathbf{w} - \alpha O_{l-1}(\mathbf{F}_c) (\nabla_{\mathbf{x}_c(0)} L_c - \nabla_{\mathbf{x}_c(1)} L_c) \end{aligned} \quad (16)$$

where L_c is the cross-entropy loss and $\alpha > 0$ is the learning rate. Besides,

$$L_c = -\mathbf{y}_n^\epsilon(0) \log \mathbf{y}_c(0) - \mathbf{y}_n^\epsilon(1) \log \mathbf{y}_c(1) \quad (17)$$

$$\mathbf{y}_c(0) = \frac{\exp \mathbf{x}_c(0)}{\exp \mathbf{x}_c(0) + \exp \mathbf{x}_c(1)} \quad (18)$$

$$\mathbf{y}_c(1) = \frac{\exp \mathbf{x}_c(1)}{\exp \mathbf{x}_c(0) + \exp \mathbf{x}_c(1)}. \quad (19)$$

Substitute (18) and (19) into (17)

$$L_c = -\mathbf{y}_n^\epsilon(0) \mathbf{x}_c(0) - \mathbf{y}_n^\epsilon(1) \mathbf{x}_c(1) + \log(\exp \mathbf{x}_c(0) + \exp \mathbf{x}_c(1)), \quad (20)$$

$$\nabla_{\mathbf{x}_c(0)} L_c = -\mathbf{y}_n^\epsilon(0) + \mathbf{y}_c(0) \quad (21)$$

$$\nabla_{\mathbf{x}_c(1)} L_c = -\mathbf{y}_n^\epsilon(1) + \mathbf{y}_c(1). \quad (22)$$

$$\nabla_{\mathbf{x}_c(0)} L_c - \nabla_{\mathbf{x}_c(1)} L_c > 0. \quad (23)$$

For hotspot instances

$$\begin{aligned} \mathbf{w}'^\top O_{l-1}(\mathbf{F}_h) &= \mathbf{w}^\top O_{l-1}(\mathbf{F}_h) - \alpha O_{l-1}^\top(\mathbf{F}_c) \\ &\quad \times (\nabla_{\mathbf{x}_c(0)} L_c - \nabla_{\mathbf{x}_c(1)} L_c) O_{l-1}(\mathbf{F}_h). \end{aligned} \quad (24)$$

Because O_{l-1} is ReLU output, we have $O_{l-1}(\mathbf{F}_c) > \mathbf{0}$ and $O_{l-1}(\mathbf{F}_h) > \mathbf{0}$. Therefore,

$$\mathbf{w}'^\top O_{l-1}(\mathbf{F}_h) > \mathbf{w}^\top O_{l-1}(\mathbf{F}_h) \quad (25)$$

which indicates that

$$\mathbf{x}_h(0) - \mathbf{x}_h(1) > \mathbf{x}'_h(0) - \mathbf{x}'_h(1) \quad (26)$$

$$\Rightarrow \frac{\exp \mathbf{x}_h(1)}{\exp \mathbf{x}_h(0) + \exp \mathbf{x}_h(1)} < \frac{\exp \mathbf{x}'_h(1)}{\exp \mathbf{x}'_h(0) + \exp \mathbf{x}'_h(1)}. \quad (27)$$

Therefore, the predicted probability of hotspot instances being real hotspots is expected to be greater, and the classifier is more confident about those wrongly detected patterns that have predict probability around 0.5. In other words, $a' \geq a$. ■

From a physical point of view, the biased term ϵ can be regarded as a force “dragging” the decision boundary closer to nonhotspot instances. However, in the space defined by the pretrained model, all the instances are located in different locations and have different distances to the decision boundary. For any nonhotspot instance, the loss with and without the biased term ϵ is given by (28) and (29), respectively

$$l = -\log \mathbf{y}(0) \quad (28)$$

$$\begin{aligned} l_b &= -(1 - \epsilon) \log \mathbf{y}(0) - \epsilon \log \mathbf{y}(1) \\ &= -(1 - \epsilon) \log \mathbf{y}(0) - \epsilon \log(1 - \mathbf{y}(0)). \end{aligned} \quad (29)$$

The training speed of the neural network is determined by the gradient of the loss with respect to neuron weights. Because the relationships between the prediction score \mathbf{y} and neuron weights are the same regardless of the loss function, the training speeds of the two cases are determined by the following equations:

$$\frac{\partial l}{\partial \mathbf{y}(0)} = -\frac{1}{\mathbf{y}(0)} \quad (30)$$

$$\frac{\partial l_b}{\partial \mathbf{y}(0)} = \frac{\epsilon + \mathbf{y}(0) - 1}{\mathbf{y}(0)(1 - \mathbf{y}(0))}. \quad (31)$$

Observe that when $\mathbf{y}(0) \leq 0.5$, i.e., the network makes an incorrect prediction

$$\left| \frac{\partial l}{\partial \mathbf{y}(0)} \right| > \left| \frac{\partial l_b}{\partial \mathbf{y}(0)} \right| \quad (32)$$

which indicates if the bias is directly applied at the random initialized networks, the network with bias updates slower than the network without bias. Therefore, we apply multiple rounds fine-tuning on the pretrained model instead of directly training the network with bias.

The bias term ϵ cannot be increased without limitations, because at some point, most of the nonhotspot patterns will cross the middle line, where the probability is 0.5, causing a significant increase of false alarms. Because the approach improves the performance of hotspot detection at the cost of confidence on nonhotspots, we call it BL. As the uncertainty exists for large CNN, a validation procedure is applied

Algorithm 2 BL

Require: $\epsilon, \delta\epsilon, t, \mathbf{W}, \lambda, \alpha, k, \mathbf{y}_h^*, \mathbf{y}_n^*$,

1: $i \leftarrow 0, \epsilon \leftarrow 0, \mathbf{y}_h^* \leftarrow [0, 1]$;

2: **while** $i < t$ **do**

3: $\mathbf{y}_n^* \leftarrow [1 - \epsilon, \epsilon]$;

4: $f_\epsilon \leftarrow \text{MGD}(\mathbf{W}, \lambda, \alpha, k, \mathbf{y}_n^*, \mathbf{y}_h^*)$;

5: $i \leftarrow i + 1, \epsilon \leftarrow \epsilon + \delta\epsilon$;

6: **end while**

▷ Algorithm 1

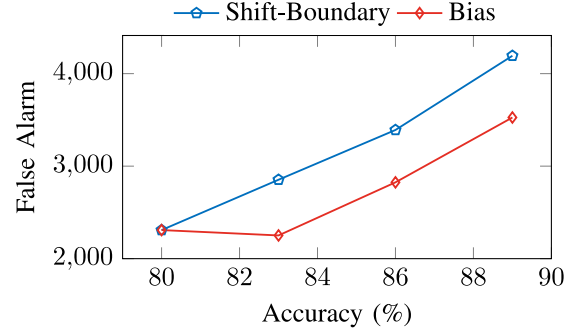


Fig. 7. Bias learning shows a smaller false alarm penalty to obtain the same hotspot detection accuracy.

to decide when to stop BL. To sum up, BL is iteratively carrying out normal MGD with changed nonhotspot ground truths, as shown in Algorithm 2.

Here ϵ is the bias, $\delta\epsilon$ represents the bias step, and t is the maximum iteration of BL. In BL, the hotspot ground truth is fixed at $[0, 1]$ while the nonhotspot truth is $[1 - \epsilon, \epsilon]$. Initially, the normal MGD is applied with $\epsilon = 0$ (line 1). After getting the converged model, fine-tune it with ϵ updated by $\epsilon = \epsilon + \delta\epsilon$. Repeat the procedure until the framework reaches the maximum bias adjusting time t (lines 2–6).

Assumption 1 shows that the BL algorithm can improve hotspot detection accuracy by taking advantage of the ReLU property. Because BL is applied through training, the false alarm penalty on the improvement of hotspot accuracy is expected to be limited. Here we evaluate the BL algorithm by training the neural network with $\epsilon = 0$ to obtain an initial model and fine-tuning with $\epsilon = 0.1, 0.2, 0.3$. Then we perform boundary shifting on the initial model to achieve the same test accuracy with three fine-tuned models. As shown in Fig. 7, BL has 600 less false alarm penalties for the same improvement of hotspot detection accuracy equivalent to saving 6000s of ODST consumption, which demonstrates the validity of Assumption 1.

D. Batch Biased Learning

In the BL algorithm, we perform fine-tuning on the pretrained models with a fixed and biased ground truth until meeting the stop condition. From the deduction above, we also notice that it might not be suitable to apply the same biased ground truth on all the nonhotspot instances. Therefore, to dynamically adjust the bias for different instances, we define a bias function as follows:

$$\epsilon(l) = \begin{cases} \frac{1}{1 + \exp(\beta l)}, & \text{if } l \leq 0.3 \\ 0, & \text{if } l > 0.3 \end{cases} \quad (33)$$

Algorithm 3 BBL

Require: \mathbf{W} , λ , α , k , \mathbf{y}_h^* , \mathbf{y}_n^* , β ;

- 1: Initialize parameters, $\mathbf{y}_h^* \leftarrow [0, 1]$;
- 2: **while not** stop condition **do**
- 3: Sample m non-hotspot instances $\{\mathbf{N}_1, \mathbf{N}_2, \dots, \mathbf{N}_m\}$;
- 4: Sample m hotspot instances $\{\mathbf{H}_1, \mathbf{H}_2, \dots, \mathbf{H}_m\}$;
- 5: Calculate average loss of non-hotspot samples l_n with ground truth [1, 0];
- 6: $\mathbf{y}_n^* \leftarrow [1 - \epsilon(l_n), \epsilon(l_n)]$;
- 7: **for** $i \leftarrow 1, 2, \dots, m$ **do**
- 8: $\mathcal{G}_{h,i} \leftarrow \text{backprop}(\mathbf{H}_i)$;
- 9: $\mathcal{G}_{n,i} \leftarrow \text{backprop}(\mathbf{N}_i)$;
- 10: **end for**
- 11: Calculate gradient $\bar{\mathcal{G}} \leftarrow \frac{1}{2m} \sum_{i=1}^m (\mathcal{G}_{h,i} + \mathcal{G}_{n,i})$;
- 12: Update weight $\mathbf{W} \leftarrow \mathbf{W} - \lambda \bar{\mathcal{G}}$;
- 13: **if** $j \bmod k = 0$ **then**
- 14: $\lambda \leftarrow \alpha \lambda$, $j \leftarrow 0$;
- 15: **end if**
- 16: **end while**

where l is the training loss of the current instance or batch in terms of the unbiased ground truth and β is a manually determined hyper-parameter that controls how much the bias is affected by the loss. Because the training loss of the instance at the decision boundary is $-\log 0.5 \approx 0.3$, we set the bias function to take effect when $l \leq 0.3$. With the bias function, we can train the neural network in a single-round MGD and obtain a better model performance. Because $\epsilon(l)$ is fixed within each training step, no additional computing effort is required for back-propagation, as indicated by (31).

The training procedure is summarized as Algorithm 3, where β is a hyper-parameter defined in (33). Similar to MGD, we initialize the neural network (line 1) and update the weight until meeting the convergence condition (lines 2–16). Within each iteration, we first sample the same amount of hotspot and nonhotspot instances to make sure the training procedure is balanced (lines 3 and 4); we then calculate the average loss of nonhotspot instances to obtain the bias level and the biased ground truth (lines 5 and 6); the gradients of the hotspot and nonhotspot instances are calculated separately (lines 8 and 9); the rest of the steps are the normal weight update through back-propagation and learning rate decay (lines 11–15).

E. Data Augmentation and Ensemble Testing

Many studies have shown that data preprocessing provides a greater generalization ability of the deep learning model [12], [16], [27]. Observe that: 1) the orientation of a clip does not affect its property under most illumination settings and 2) usually there are more nonhotspot clips than hotspot clips, we include several data augmentation techniques in the training stage accordingly.

- 1) We randomly perform flipping (top-bottom transformation) and mirroring (left-right transformation) on each feature tensor along its last axis. It is easy to derive that the convolution operation is not flipping invariant, therefore random flipping and mirroring will introduce

TABLE II
BENCHMARK STATISTICS

Benchmarks	Training Set		Testing Set		Size/Clip (μm^2)
	HS#	NHS#	HS#	NHS#	
ICCAD	1204	17096	2524	13503	3.6×3.6
Industry0	3629	80299	942	20412	1.2×1.2
Industry1	34281	15635	17157	7801	1.2×1.2
Industry2	15197	48758	7520	24457	1.2×1.2
Industry3	24776	49315	12228	24817	1.2×1.2

diversity to the dataset and increase the trained model’s generalization ability (how well the model fits the testing data [28], [29]).

- 2) We force the number of hotspot and nonhotspot instances to be equal in each mini-batch. As shown in the experiment of [16], a highly imbalanced dataset causes performance degradation. Sampling equal amount of instances from different categories is expected to benefit both the training progress and the model performance.

To make better use of the flipping variance property, an ensemble method is applied in the testing phase.

- 1) We do prediction on four directions of each clip and take the average of the prediction scores as the final prediction result.

Although ensemble testing will induce additional testing runtime, it offers better model performance.

V. EXPERIMENTAL RESULTS

A. Experimental Setup

We implement our deep BL framework in Python with the TensorFlow library [30], and test it on a platform with a Xeon E5 processor and Nvidia Graphic card. To fully evaluate the proposed framework, we employ four test cases. Because the individual test cases in the ICCAD 2012 contest [10] are not large to verify the scalability of our framework, we merge all the 28-nm patterns into a unified test case ICCAD.

Additionally, we adopt four more complicated industry test cases: Industry0 – Industry3. The details for all test cases are listed in the Table II.

Columns “Train HS#” and “Train NHS#” list the total number of hotspots and the total number of nonhotspots in the training set. Columns Test NHS# and Test HS# list the total number of hotspots and total number of nonhotspots in the testing set. Images in the testing set of ICCAD have a resolution of 3600×3600 which is larger than the images in industry benchmarks (1200×1200). Therefore, during the testing phase, each clip in ICCAD Testing Set is divided into nine 1200×1200 blocks before feeding into the testing flow. Mask images in the four benchmarks are from different OPC stages, and have different complexities, as shown in Fig. 8. Note that classic hotspot detection problems aim to find and revise problematic designs at an early stage of the whole layout verification flow which corresponds to cases ICCAD benchmarks and Industry0 that contains before-OPC patterns and are labeled based on the results of the entire layout verification flow, including OPC and lithography simulation. Industry1–Industry3 are benchmark sets




Fig. 8. Benchmark examples. (a) and (b) are before-OPC patterns from ICCAD and Industry0, respectively. (b)–(d) correspond to Industry1–Industry3, which are from intermediate OPC results. The pattern becomes more complicated after OPC and is more challenging for machine learning based hotspot detectors.

TABLE III
TRAINING CONFIGURATIONS

Configurations	Value
Optimizer	Adam [31]
Initial Learning Rate (λ)	0.001
Learning Rate Decay (α)	0.75
LR Decay Step (k)	10,000
Bias Function Coefficient (β)	6 and 43
Batch Size	32
Feature Tensor Channel	32

that contain layouts from intermediate OPC results which are labeled hotspot or nonhotspot based on the lithography simulation results of current OPC step. The motivation of introducing intermediate OPCed layouts is to show some potential of embedding efficient hotspot detectors into OPC engines and facilitate the procedure.

B. Model Training

We train five individual models for each benchmark set following Algorithm 3. Table III lists the details of the training configurations. “Adam” is an improved optimizer to conduct a gradient descent proposed in [31] that is proved to converge faster. Parameters α and k denote that the learning rate (λ) drops to $\alpha\lambda$ every k steps. β is the coefficient that appears in (33) and controls how much the bias is affected by the loss of a nonhotspot batch, therefore, it also controls the tradeoffs between accuracy and false alarms. Here we use $\beta = 6$ for the datasets (ICCAD, Industry0, and Industry1) with more regular patterns and $\beta = 43$ for the datasets (Industry2 and Industry3) with complicated patterns. Because it takes more effort to fit the complicated patterns with the neural networks, we pick a larger β to avoid additional perturbation on the neural networks at early training stages (Fig. 9). Feature tensor channel represents the number of remaining elements after dropping high frequency components in the feature tensor extraction procedure.

We visualize the training loss of BBL and BL of five benchmarks in Fig. 9, where Fig. 9(a)–(e) correspond to ICCAD, Industry0, Industry1, Industry2, and Industry3, respectively. The “Loss” is the average cross-entropy loss given by (7) with respect to the unbiased ground truth. For the ICCAD and Industry0 dataset, the BBL converges quickly within 20 000 steps, while the BL requires manual adjustment of the bias and finally converges at a higher loss level [see the blue curve in Fig. 9(a) and (b)]. Also, for the benchmarks Industry1–Industry3, the BBL (red curve) has better convergent result than the BL (blue).

C. Model Testing

To evaluate the effectiveness of the BBL, we first compare the test results on four datasets (statistics are listed in Table II) with our preliminary results in [19], as shown in Tables IV and V. “Accu (%)” and “FA #” denote the hotspot detection accuracy (Definition 1) and the false alarm number (Definition 2). “FA (%)” is the percentage representation of false alarms that are coherent with FA #. “SPIE’17 [16]+BL” corresponds to the results obtained by inserting BL (Algorithm 2) in SPIE’17 [16] neural networks model; “SPIE’17 [16]+BBL” lists the results obtained by embedding BBL (Algorithm 3) in SPIE’17 [16] neural networks model; “BL [19]” corresponds to the original BL algorithm (Algorithm 2) that is applied in our preliminary work; column “BL [19] + AUG” contains the results obtained from the original BL and the data augmentation mentioned in Section IV-E; “BBL” lists the testing results of the BBL as in Algorithm 3; and “BBL + AUG” also includes the data augmentation in the BBL procedure. Note that original SPIE’17 [16] also employs data augmentation on raw layout images.

In this paper, we propose a BBL algorithm that can train the neural network in one round MGD and seek a better tradeoff than the BL algorithm. Table IV lists the testing results of target layouts ICCAD and Industry0, which show that BBL surpasses BL on both average detection accuracy and false alarm. With the aid of data augmentation, detection accuracy is further improved from 97.8% to 98.8% with ignorable false alarm penalty. For intermediate OPCed layouts Industry1–3 in Table V, BBL also significantly reduces average false alarm from 2347 to 1981 when achieving almost the same detection accuracy. We can also notice that for OPCed layouts with data augmentation, BBL does not exhibit as efficiently as on target layouts. One reason falls on the complexity of OPC patterns which induce more challenge on neural network training, when unbiased labels dominate the training procedure compared to BL which forces nonhotspot labels to be biased.

We also embed the proposed BL and BBL algorithms into large neural networks designed in SPIE’17 [16] which takes raw layout images as input. Compared with original SPIE’17 [16] (as shown in Tables VI and VII), BL and BBL can significantly improve the hotspot detection performance. For target layouts, average detection accuracy increased from 97.63% to 98.17% and 98.76% with BL and BBL,

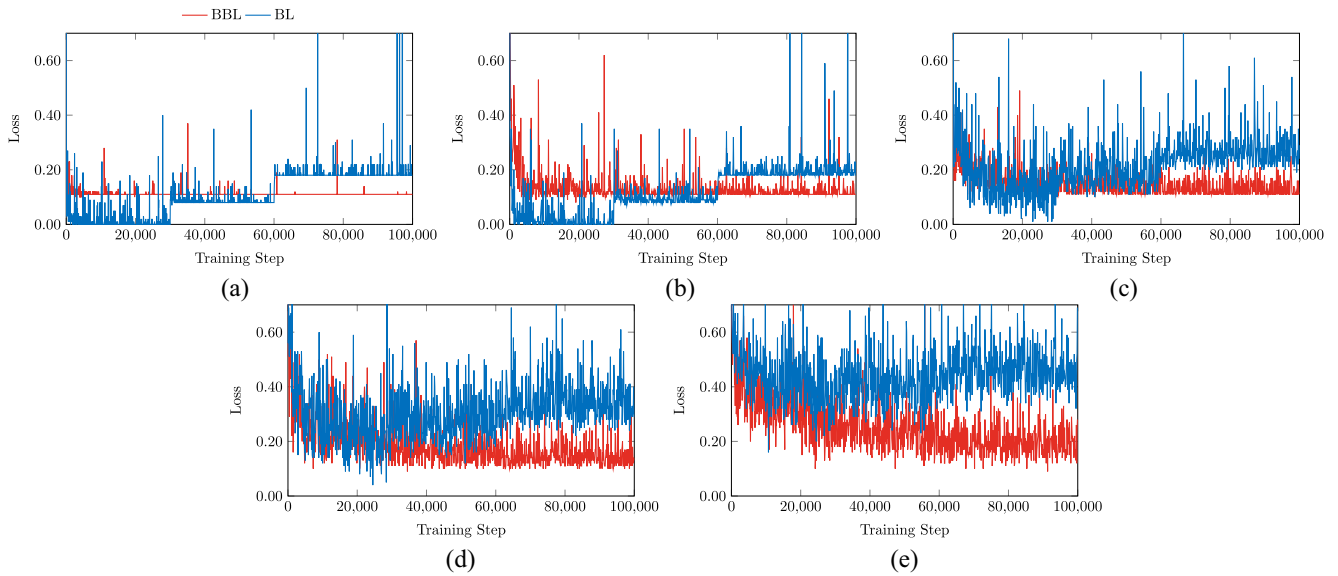


Fig. 9. Visualize training of the BBL (red) and the BL (blue). On each benchmarks, BBL exhibits lower loss than the BL at convergence. (a) ICCAD. (b) Industry0. (c) Industry1. (d) Industry2. (e) Industry3.

TABLE IV
PERFORMANCE COMPARISON BETWEEN THE BL AND THE BBL ON TARGET LAYOUTS

Benchmarks	SPIE'17 [16]+BL		SPIE'17 [16]+BBL		BL [19]		BL [19]+AUG		BBL		BBL+AUG	
	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#
ICCAD	97.60	2054	98.04	3237	98.20	3413	98.40	3878	98.00	2745	98.40	3535
Industry0	98.73	316	99.47	469	97.35	142	99.26	246	97.66	136	99.36	387
Average	98.17	1185	98.76	1853	97.78	1778	98.83	2062	97.83	1441	98.88	1961
Ratio	0.993	0.604	0.999	0.945	0.989	0.906	0.999	1.052	0.989	0.735	1.000	1.000

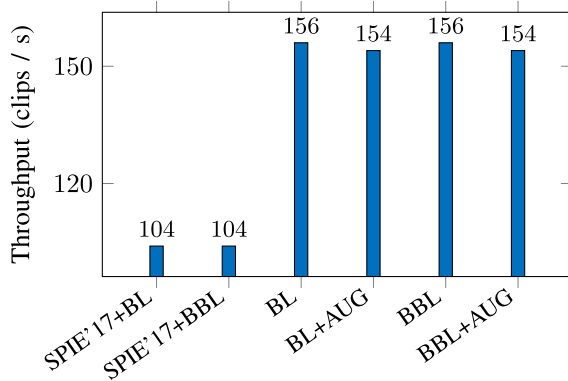


Fig. 10. Throughput comparison of different neural network models.

respectively. BL also reduces the false alarm count from 1394 to 1185. For OPCed layouts, BL and BBL dramatically increase the detection accuracy from 88.17% to 94.39% and 95.84%, respectively. In particular, BBL can further improve the hotspot detection accuracy when inserted in large SPIE'17 nets at relatively large false alarm penalty, because biased labels dominate during the training procedure of BBL for neural networks with larger capacity. Fig. 10 illustrates the testing speed of different neural network models that correspond to Tables IV and V. From the experimental results, we can clearly see that the proposed methods can achieve similar hotspot detection accuracy compared to larger neural network models but with much less computing costs (104 clips/s of

deep networks in SPIE'17 versus 156 clips/s in our proposed network architecture). Although data augmentation and ensemble testing slightly induces computing costs, we can still observe great advantage of the proposed models over SPIE'17 nets.

We then compare the hotspot detection results with four state-of-the-art hotspot detectors in Tables VI and VII. "SPIE'15 [5]" is a traditional machine learning-based hotspot detector that applies the density-based layout features and the AdaBoost [32]—DecisionTree model. "ICCAD'16 [6]" takes the lithographic properties into account during feature extraction and adopts the more robust Smooth Boosting [33] algorithm. "SPIE'17 [16]" is another deep learning solution for hotspot detection that takes the original layout image as input and contains more than 20 layers. "SOCC'17 [18]" employs a deep neural networks that replace all pooling layers with strided convolution layers and contains the same number of layers as SPIE'17.

Overall, our framework performs better than traditional machine learning techniques (SPIE'15 [5] and ICCAD'16 [6]) with at least a 2% advantage for the detection accuracy (98.88% of BBL versus 96.89% of ICCAD'16) on target layouts. Traditional machine learning models are effective for the benchmarks with regular patterns (ICCAD and Industry0) with the highest detection accuracy of 97.7% on the ICCAD and 96.07% on the Industry0 achieved by [6]. However, manually designed features, including the density-based features and the CCAS, have difficulties to grasp the attributes

TABLE V
PERFORMANCE COMPARISON BETWEEN THE BL AND THE BBL ON OPCED LAYOUTS

Benchmarks	SPIE'17 [16]+BL		SPIE'17 [16]+BBL		BL [19]		BL [19]+AUG		BBL		BBL+AUG	
	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#
Industry1	98.61	311	98.30	362	98.90	680	99.00	558	98.20	402	98.20	409
Industry2	94.62	1156	95.49	1550	93.60	2165	94.80	1483	94.70	1702	95.50	1691
Industry3	89.95	2541	93.72	4361	91.30	4196	91.30	3917	90.00	3840	91.40	3888
Average	94.39	1336	95.84	2091	94.60	2347	95.03	1986	94.30	1981	95.03	1996
Ratio	0.993	0.669	1.008	1.048	0.995	1.176	1.000	0.995	0.992	0.993	1.000	1.000

TABLE VI
PERFORMANCE COMPARISON WITH STATE-OF-THE-ART HOTSPOT DETECTORS ON TARGET LAYOUTS

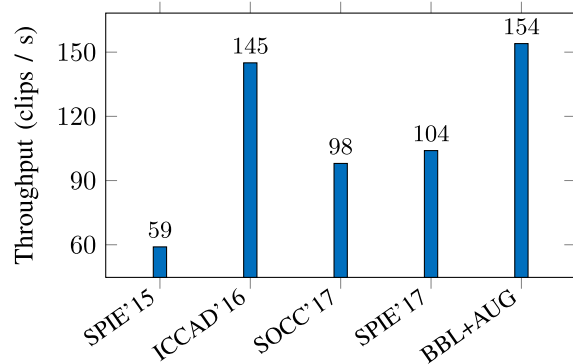
Benchmarks	SPIE'15 [5]		ICCAD'16 [6]		SOCC'17 [18]		SPIE'17 [16]		BBL+AUG	
	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#
ICCAD	84.20	2919	97.70	4497	96.90	1960	97.70	2703	98.40	3535
Industry0	93.63	30	96.07	1148	97.77	100	97.55	85	99.36	387
Average	88.92	1475	96.89	2823	97.34	1030	97.63	1394	98.88	1961
Ratio	0.899	0.752	0.980	1.439	0.984	0.525	0.987	0.711	1.000	1.000

TABLE VII
PERFORMANCE COMPARISON WITH STATE-OF-THE-ART HOTSPOT DETECTORS ON OPCED LAYOUTS

Benchmarks	SPIE'15 [5]		ICCAD'16 [6]		SOCC'17 [18]		SPIE'17 [16]		BBL+AUG	
	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#	Accu (%)	FA#
Industry1	93.20	2204	89.90	1136	98.50	318	97.74	519	98.20	409
Industry2	44.80	1320	88.40	7402	91.1	710	89.62	760	95.50	1691
Industry3	44.00	3144	82.30	8609	84.01	1628	77.14	1966	91.40	3888
Average	60.67	2223	86.87	5716	91.20	885	88.17	1082	95.03	1996
Ratio	0.638	1.114	0.914	2.864	0.960	0.444	0.928	0.542	1.000	1.000

of the post-OPC mask layouts. For OPCed layouts, [5] suffers a large performance degradation with only approximately 45% detection accuracy on the most complicated case Industry3. Although prior knowledge and a more robust Smooth Boosting [33] algorithm are applied in [6], the hotspot detection accuracy inevitably drops around 8%. Deep learning solutions [16], [18] and BBL+AUG exhibit better performances with hotspot detection accuracy of 91.20%, 88.17%, and 95.03%, respectively. It is notable that the architecture of our framework contains significantly fewer layers than the framework in SPIE'17 and SOCC'17. Although the shallow architecture results in acceptable more false alarms than deep models, our framework can still offer significant higher accuracy on ICCAD, Industry0, Industry2, and Industry3.

Although neural networks are not as computational efficient as traditional machine learning methods, with the aid of parallel computing units (e.g., GPU), we are able to achieve comparable and acceptable processing speed. Fig. 11 presents the detecting speed of different hotspot detectors adopted in this paper, where SPIE'15 and ICCAD'16 are tested on CPU only and neural network implementations are tested on graphic cards. Runtime reports show that although neural network models are not computational friendly due to the complicated convolutional operations, we are still able to complete the task with comparable and acceptable time, which also show the potential of enhancing the layout verification flow with dedicated computing units instead of CPU only.



- SPIE'15 & ICCAD'16: tested on 3.3GHz Quad-Core Intel processor;
- SOCC'17, SPIE'17 & BBL+AUG: tested on NVIDIA GPU.

Fig. 11. Throughput comparison with state-of-the-art hotspot detectors.

VI. CONCLUSION

To address the existing problems of machine learning-based printability estimation techniques, we first propose a high-dimensional feature (feature tensor) extraction method that can reduce the size of training instances while keeping the spatial information. The feature tensor is also compatible with powerful convolutional neural networks. Additionally, to improve hotspot detection accuracy, we first develop a BL algorithm, which takes advantage of the ReLU function in CNN to prominently increase accuracy while reducing false alarm penalties. We further propose a BBL algorithm to

automatically adjust the training ground truth according to the current batch loss, that can offer better tradeoffs between accuracy and false alarms. The experimental results show that the BBL algorithm is more efficient during training and our framework outperforms the other hotspot solutions on complicated benchmarks. Source code and trained models are available at <https://github.com/phdyang007/dllhsd>. As the technology node keeps shrinking down, we hope this paper can be a demonstration that deep learning has the potential to provide satisfactory solutions for advanced design for manufacturability research.


REFERENCES

- [1] W.-Y. Wen, J.-C. Li, S.-Y. Lin, J.-Y. Chen, and S.-C. Chang, "A fuzzy-matching model with grid reduction for lithography hotspot detection," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 33, no. 11, pp. 1671–1680, Nov. 2014.
- [2] Y.-T. Yu, Y.-C. Chan, S. Sinha, I. H.-R. Jiang, and C. Chiang, "Accurate process-hotspot detection using critical design rule extraction," in *Proc. ACM/IEEE Design Autom. Conf. (DAC)*, San Francisco, CA, USA, 2012, pp. 1167–1172.
- [3] F. Yang, S. Sinha, C. C. Chiang, X. Zeng, and D. Zhou, "Improved tangent space-based distance metric for lithographic hotspot classification," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 36, no. 9, pp. 1545–1556, Sep. 2017.
- [4] Y.-T. Yu, G.-H. Lin, I. H.-R. Jiang, and C. Chiang, "Machine-learning-based hotspot detection using topological classification and critical feature extraction," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 34, no. 3, pp. 460–470, Mar. 2015.
- [5] T. Matsunawa, J.-R. Gao, B. Yu, and D. Z. Pan, "A new lithography hotspot detection framework based on AdaBoost classifier and simplified feature extraction," in *Proc. SPIE*, vol. 9427. San Jose, CA, USA, 2015, Art. no. 94270S.
- [6] H. Zhang, B. Yu, and E. F. Y. Young, "Enabling online learning in lithography hotspot detection with information-theoretic feature optimization," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Austin, TX, USA, 2016, pp. 1–8.
- [7] D. Ding, A. J. Torres, F. G. Pikus, and D. Z. Pan, "High performance lithographic hotspot detection using hierarchically refined machine learning," in *Proc. IEEE/ACM Asia South Pac. Design Autom. Conf. (ASPAC)*, Yokohama, Japan, 2011, pp. 775–780.
- [8] D. Ding, B. Yu, J. Ghosh, and D. Z. Pan, "EPIC: Efficient prediction of IC manufacturing hotspots with a unified meta-classification formulation," in *Proc. IEEE/ACM Asia South Pac. Design Autom. Conf. (ASPAC)*, Sydney, NSW, Australia, 2012, pp. 263–270.
- [9] D. G. Drmanac, F. Liu, and L.-C. Wang, "Predicting variability in nanoscale lithography processes," in *Proc. ACM/IEEE Design Autom. Conf. (DAC)*, San Francisco, CA, USA, 2009, pp. 545–550.
- [10] A. J. Torres, "ICCAD-2012 CAD contest in fuzzy pattern matching for physical verification and benchmark suite," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, San Jose, CA, USA, 2012, pp. 349–350.
- [11] T. Matsunawa, B. Yu, and D. Z. Pan, "Optical proximity correction with hierarchical Bayes model," in *Proc. SPIE*, vol. 9426. San Jose, CA, USA, 2015, Art. no. 94260X.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Conf. Neural Inf. Process. Syst. (NIPS)*, 2012, pp. 1097–1105.
- [13] T. Xiao, H. Li, W. Ouyang, and X. Wang, "Learning deep feature representations with domain guided dropout for person re-identification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, 2016, pp. 1249–1258.
- [14] T. Matsunawa, S. Nojima, and T. Kotani, "Automatic layout feature extraction for lithography hotspot detection based on deep neural network," in *Proc. SPIE Adv. Lithography*, vol. 9781. San Jose, CA, USA, 2016, Art. no. 97810H.
- [15] M. Shin and J.-H. Lee, "Accurate lithography hotspot detection using deep convolutional neural networks," *J. Micro/Nanolithography MEMS MOEMS*, vol. 15, no. 4, 2016, Art. no. 043507.
- [16] H. Yang, L. Luo, J. Su, C. Lin, and B. Yu, "Imbalance aware lithography hotspot detection: A deep learning approach," in *Proc. SPIE Adv. Lithography*, vol. 10148, 2017, Art. no. 1014807.
- [17] H. Yang, L. Luo, J. Su, C. Lin, and B. Yu, "Imbalance aware lithography hotspot detection: A deep learning approach," *J. Micro/Nanolithography MEMS MOEMS*, vol. 16, no. 3, 2017, Art. no. 033504.
- [18] H. Yang, Y. Lin, B. Yu, and E. F. Y. Young, "Lithography hotspot detection: From shallow to deep learning," in *Proc. IEEE Int. Syst. Chip Conf. (SOCC)*, Munich, Germany, 2017, pp. 233–238.
- [19] H. Yang, J. Su, Y. Zou, B. Yu, and E. F. Y. Young, "Layout hotspot detection with feature tensor generation and deep biased learning," in *Proc. ACM/IEEE Design Autom. Conf. (DAC)*, Austin, TX, USA, 2017, pp. 1–6.
- [20] W. Zhang, X. Li, S. Saxena, A. Strojwas, and R. Rutenbar, "Automatic clustering of wafer spatial signatures," in *Proc. ACM/IEEE Design Autom. Conf. (DAC)*, Austin, TX, USA, 2013, pp. 1–6.
- [21] S. Shim, W. Chung, and Y. Shin, "Synthesis of lithography test patterns through topology-oriented pattern extraction and classification," in *Proc. SPIE*, vol. 9053. San Jose, CA, USA, 2014, Art. no. 905305.
- [22] G. K. Wallace, "The JPEG still picture compression standard," *IEEE Trans. Consum. Electron.*, vol. 38, no. 1, pp. 18–34, Feb. 1992.
- [23] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [24] W. A. Gardner, "Learning characteristics of stochastic-gradient-descent algorithms: A general study, analysis, and critique," *Signal Process.*, vol. 6, no. 2, pp. 113–133, 1984.
- [25] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>
- [26] S. Hochreiter, "The vanishing gradient problem during learning recurrent neural nets and problem solutions," *Int. J. Uncertainty Fuzziness Knowl. Based Syst.*, vol. 6, no. 2, pp. 107–116, 1998.
- [27] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015, pp. 1–14.
- [28] C. Zhang *et al.*, "Theory of deep learning III: Generalization properties of SGD," Center Brains Minds Mach., Massachusetts Inst. Technol., Rep., 2017.
- [29] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein, "On the expressive power of deep neural networks," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Sydney, NSW, Australia, 2017, pp. 2847–2854.
- [30] M. Abadi *et al.*, "TensorFlow: A system for large-scale machine learning," in *Proc. USENIX Symp. Oper. Syst. Design Implement. (OSDI)*, Savannah, GA, USA, 2016, pp. 265–283.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015.
- [32] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *Proc. Eur. Conf. Comput. Learn. Theory*, Barcelona, Spain: Springer, 1995, pp. 23–37.
- [33] R. A. Servedio, "Smooth boosting and learning with malicious noise," *J. Mach. Learn. Res.*, vol. 4, pp. 633–648, Sep. 2003.



Haoyu Yang received the B.E. degree from Qiushi Honors College, Tianjin University, Tianjin, China, in 2015. He is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, Chinese University of Hong Kong, Hong Kong.

His current research interests include very large scale integration design for manufacturability and machine learning.



Jing Su received the B.S. degree in physics from the University of Science and Technology of China, Hefei, China, in 2008 and the Ph.D. degree in theoretical and computational AMO physics from the University of Colorado at Boulder, Boulder, CO, USA, in 2014.

He is currently a Senior Design Engineer with the Advanced Technology Development Group, ASML-Brion, San Jose, CA, USA. His current research interests include wide range of topics in lithography, including advanced mask optimization, data mining, and machine learning.



Yi Zou received the B.S. degree in control science and engineering from Zhejiang University, Hangzhou, China, in 1998, the M.Phil. degree in industrial engineering from the Hong Kong University of Science and Technology, Hong Kong, in 2000, and the Ph.D. degree in microchip simulation from Boston University, Boston, MA, USA, in 2004.

He is currently the Director of Engineering at ASML, San Jose, CA, USA. His current research interests include optical proximity correction and lithography scanner modeling and control via advanced data analytics.



Bei Yu (S'11–M'14) received the Ph.D. degree from the University of Texas at Austin, Austin, TX, USA, in 2014.

He is currently an Assistant Professor with the Department of Computer Science and Engineering, Chinese University of Hong Kong, Hong Kong.

Dr. Yu was a recipient of Five Best Paper Awards from *Integration, the VLSI Journal* in 2018, the International Symposium on Physical Design 2017, SPIE Advanced Lithography Conference 2016, International Conference on Computer Aided Design 2013, and Asia and South Pacific Design Automation Conference 2012, and Four International Conference on Computer Aided Design/International Symposium on Physical Design contest awards. He has served in the editorial boards of *Integration, the VLSI Journal* and *IET Cyber-Physical Systems: Theory & Applications*.



Evangeline F. Y. Young received the B.Sc. degree in computer science from the Chinese University of Hong Kong (CUHK), Hong Kong and the Ph.D. degree from the University of Texas at Austin, Austin, TX, USA, in 1999.

She is currently a Professor with the Department of Computer Science and Engineering, CUHK. Her current research interests include optimization algorithms and very large scale integration CAD. She researches actively on floorplanning, placement, routing, DFM, and EDA on physical design in

general.

Dr. Young's research group has won several championships and prizes in renowned EDA contests, including the 2016, 2015, 2013, and 2012 CAD Contests at International Conference on Computer Aided Design (ICCAD), Design Automation Conference (DAC) 2012, and International Symposium on Physical Design (ISPD) 2011 Routability-Driven Placement Contests and ISPD 2010 High Performance Clock Network Synthesis Contest. She has served on the organization committees of ISPD, Applied Reconfigurable Computing Symposium, and International Conference on Field-Programmable Technology and on the program committees of conferences, including DAC, ICCAD, ISPD, Asia and South Pacific Design Automation Conference, System Level Interconnect Prediction Workshop, Design, Automation and Test in Europe Conference, and Great Lakes Symposium on VLSI. She also served on the editorial boards of the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, ACM TODAES, and *Integration, the VLSI Journal*.



Yuzhe Ma received the B.E. degree from the Department of Microelectronics, Sun Yat-sen University, Guangzhou, China, in 2016. He is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, Chinese University of Hong Kong, Hong Kong.

His current research interests include very large scale integration design for manufacturing, physical design, and machine learning on chips.