

Received February 20, 2019, accepted March 7, 2019, date of publication March 11, 2019, date of current version April 2, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2904245

Discriminative Pattern Mining for Breast Cancer Histopathology Image Classification via Fully Convolutional Autoencoder

XINGYU LI¹, (Member, IEEE), MARKO RADULOVIC², KSENIJA KANJER², AND KONSTANTINOS N. PLATANIOTIS¹, (Fellow, IEEE)

¹The Edward S. Rogers Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada

²National Cancer Research Centre, Department of Experimental Oncology, Institute for Oncology and Radiology of Serbia, 11000 Belgrade, Serbia

Corresponding author: Xingyu Li (xingyu.li@mail.utoronto.ca)

This work was supported by the NSERC Discovery Grant.

ABSTRACT Accurate diagnosis of breast cancer in histopathology images is challenging due to the heterogeneity of cancer cell growth as well as a variety of benign breast tissue proliferative lesions. In this paper, we propose a practical and self-interpretable invasive cancer diagnosis solution. With minimum annotation information, the proposed method mines contrast patterns between normal and malignant images in a weak-supervised manner and generate a probability map of abnormalities to verify its reasoning. Particularly, a fully convolutional autoencoder is used to learn the dominant structural patterns among normal image patches. Patches that do not share the characteristics of this normal population are detected and analyzed by one-class support vector machine and one-layer neural network. We apply the proposed method to a public breast cancer image set. Our results, in consultation with a senior pathologist, demonstrate that the proposed method outperforms existing methods. The obtained probability map could benefit the pathology practice by providing visualized verification data and potentially leads to a better understanding of data-driven diagnosis solutions.

INDEX TERMS Breast cancer diagnosis, abnormality detection, convolutional autoencoder, discriminative pattern learning, histopathology image analysis.

I. INTRODUCTION

Breast cancer is the most common cancer in women. Invasive, malignant properties of breast cancer cell growth contribute to poor patient prognosis [1], and dictate precise early diagnosis and treatment, with an aim to reduce breast cancer morbidity rate. In this study, we particularly focus on the qualification of risky, aggressive characteristics of breast histomorphological patterns, as one of the basic features of invasiveness of breast carcinoma.

With the advance of imaging device and machine learning technology, digital histopathology image analysis becomes a promising approach to consistent and cost-efficient cancer diagnosis. Particularly for invasive breast cancer, based on the common knowledge that cancerous cells break through the basement membrane of ductulo-lobular structures and infiltrate into surrounding tissues - the feature of invasiveness [2] (as shown in Fig. 1), many algorithms were proposed

to classify breast histopathology images using nuclei's morphology and spatial-distribution features [3]. In literature, the most common solution to breast cancer image diagnosis is to train a classifier in a supervised learning manner. Then handcrafted features of a query image are passed to the trained algorithm for a yes/no label [4]–[10]. With the success of deep learning, data-driven methods, especially the end-to-end training of convolutional neural network, are adopted more often in recent breast cancer histopathology image classification studies [11]–[13]. Though breast cancer image diagnosis has achieved impressive progress, the issue of self-interpretability in existing diagnosis approaches is less addressed. Self-interpretability refers to the capability of an approach to explain and verify its reasoning and results. Without self-interpretability, attempts to improve histopathology image diagnosis is prone to be limited to trial-and-error.

To address the self-interpretability issue in breast cancer pathology image diagnosis, one solution is to generate labels for image pixels or small image patches in order

The associate editor coordinating the review of this manuscript and approving it for publication was Gustavo Callico.




FIGURE 1. Examples of hematoxylin and eosin stained images for (a) normal breast tissue and (b) invasive carcinoma with a magnification of 40 \times [14]. The left image corresponds to a normal tissue where normal epithelial cells lie on the membrane of ductulo-lobular structures; while in the right image malignant cells invade and spread into surrounding tissue.

to infer locations of suspected abnormalities in a query image. To this end, several studies made efforts to classify small image patches or pixels via supervised/semi-supervised learning [15]–[19]. It should be noted that these solutions require a large amount of images manually-annotated at the image pixel level. Due to the complexity and time-intensive properties of pathological annotations and privacy concerns in clinical practice, sufficient amount of well-labeled patches are difficult to collect.

This study attempts to tackle the self-interpretability issue in breast cancer diagnosis and presents a novel convolutional autoencoder-based contrast pattern mining approach to detect the invasive component of malignant breast epithelial growth in routine hematoxylin and eosin (H&E) stained histopathology images. As opposed to prior studies that require image sets with pixel annotation, our method requires only image labels as the minimal prior knowledge in training. By mining dominant patterns in images of normal breast tissues, the method generates a probability map to infer locations of abnormalities in an image. As a pathology image may contain both normal and cancerous tissues, the proposed method divides an image into small patches to facilitate local characteristics learning. It should be noted that due to the lack of pixel annotation indicating the locations of abnormal cell growth patterns in images, this problem is very challenging in two folds.

- 1) The algorithm is expected to learn contrast patterns between normal and malignant/invasive growth based on the knowledge of image labels. Effective differentiation between normal and abnormal histomorphology via weak-supervised learning is the key issue for the correct identification of cancerous growth.
- 2) As a histopathology image may contain both normal and cancerous tissue, labels of local patches may be inconsistent with the known image label. The method needs to learn a mapping function between local patches and image labels.

Note that though we do not know whether patches from images labeled as malignant really contain malignant cells/structures, patches from normal images do not contain cancerous cells certainly. So, we name a patch from a normal

image “true-normal” in this paper. Our original approach learns patterns in true-normal patches first and then assigns a normal/malignant label to a patch which resembles/deviates from those true normal ones. Intuition behind this originality is that in pathology, malignant cells and their growth patterns are diagnosed and graded by how different these cells are to normal cells. Specifically, to address the first challenge, we exploit the data-specific property of autoencoder (AE) networks [20], [21] and innovate to train an under-complete deep fully convolutional AE using small patches from pathology images annotated as normal. Since the network learns local patterns in true-normal patches only, its performance degrades when the input instance is different from training patches. Hence, autoencoder’s reconstruction residue suggests the similarity between the query instance and normal cases. It is noteworthy that different from standard autoencoders targeting to minimize mean square error (MSE) between input and output training instances, the proposed method trains the deep net by optimizing the structural similarity (SSIM) index [22], which enforces the network to learn the contrast and structural patterns in true-normal patches. In this study, the trained AE network is treated as a pattern mining and representation method and then combined with downstream classifiers to identify whether an image patch contains malignant cells.

To tackle the second challenge which is to infer whether a local patch contains morphological abnormalities derived by malignant cell growth, we cast the problem into the anomaly detection scenario, and introduce a novel malignant patch detector to distinguish patches containing cancerous cells from the normal ones. Particularly, the proposed detector makes the use of one class support vector machine (SVM) [23] to identify regions occupied by true-normal patches in the feature space and assigns abnormal labels to patches whose numerical features are located outside of the detected normal regions. Taking into account the obtained patch labels, the problem of breast cancer image classification with localization of abnormality areas is simplified to a patch-based supervised learning problem. Finally, a 1 layer neural network (NN) is trained to infer the existence of malignant tumor in a patch, followed by the generation of a probability map of abnormality in the query image.

In summary, this study proposes a practical, generalizable, and self-interpretable solution to pathology image based cancer diagnosis. With the minimal prior knowledge on whoselide-imaging (WSI) labels which can be easily acquired in clinic practice, the proposed method learns discriminative patterns in weak-supervised manner from histopathology images and explains its diagnosis results via inferring locations of abnormalities in an image. It is noteworthy that the proposed method is very user-friendly to pathologists, as the obtained abnormality map helps pathologists to understand and verify how machines make decisions. To the best of our knowledge, our work constitutes the first attempt in literature to tackle the self-interpretability issue in histopathology image classification.

TABLE 1. Table of notations.

Notations	Explanations
\mathcal{A}, \mathcal{B}	Trainable parameters of 1-layer NN
$\mathcal{D}()$	Decoder of AE
$\mathcal{E}()$	Encoder of AE
$\mathcal{F}()$	Patch labeling function
$\mathcal{G}()$	Decision function of one-class SVM
\mathcal{I}	Histopathology image
\mathcal{L}	Loss function of AE
M	Number of normal images
N	Number of malignant images
T	Number of training patches
c, ν	Hyper-parameter of one-class SVM
$k()$	Gaussian kernel
w_i, v_i	Trainable parameters of AE
x	Input of AE (i.e. gray-scale true-normal patches)
\tilde{x}	Output of AE
\hat{x}	residue of AE's reconstruction
y	Patch labels
z	Input of one-class SVM
\dagger	Histopathology image Label
α, β, γ	Hyper-parameters of SSIM
$\delta()$	Dirac delta function
λ	Lagrangian multiplier of one-class SVM

The rest of this paper is organized as follows. Section II provides brief introduction of machine learning techniques exploited in the proposed method and the public breast cancer biopsy image set used in this study. The problem's formal statement with notations and implementation details are presented in Section III and Section IV, respectively. Experimental results and discussions are presented in Section V, followed by conclusions in Section VI.

II. BACKGROUND

In this section, we will first introduce notations used in this study in Table 1. Then brief description of fully convolutional autoencoder and one-class support vector machine is presented, followed by information on the public image set used to evaluate the proposed method in this study.

A. FULLY CONVOLUTIONAL AE NETWORKS

Fully convolutional network is defined as the neural network composed of convolutional layers without any fully-connected layers [24]. It learns representations and makes decisions based on local spatial knowledge only. Because of its efficient learning, fully convolutional net has been popular in many image-to-image inference tasks, e.g. semantic segmentation.

Fully convolutional autoencoder is one instance of fully convolutional neural networks. The net takes input of arbitrary size and produces corresponding-sized output. Specifically, it encodes an image data x of arbitrary size into a low-dimensional representation \hat{x} such that the important properties of the original data can be reconstructed and maintained in the output \tilde{x} . Mathematically, a fully convolutional autoencoder is composed of an encoder $\mathcal{E}()$

and a decoder $\mathcal{D}()$, each of which is a composition of a sequence of C layers, i.e.

$$\begin{aligned}\tilde{x} &= \mathcal{D}(\hat{x}; v_1, \dots, v_C) \\ &= \mathcal{D}(\mathcal{E}(x; w_1, \dots, w_C); v_1, \dots, v_C),\end{aligned}\quad (1)$$

where $\tilde{x} = \mathcal{D}(\hat{x}; v_1, \dots, v_C) = \mathcal{D}_C(\cdot; v_C) \circ \dots \circ \mathcal{D}_1(\hat{x}; v_1)$ and $\hat{x} = \mathcal{E}(x; w_1, \dots, w_C) = \mathcal{E}_C(\cdot; w_C) \circ \dots \circ \mathcal{E}_1(x; w_1)$. $\mathcal{D}() \circ \mathcal{E}(x) = \mathcal{D}(\mathcal{E}(x))$ and w_i and v_i are the weights and bias for the i^{th} encoder layer $\mathcal{E}_i()$ and decoder layer $\mathcal{D}_i()$, respectively. Conventionally, $\mathcal{E}_i()$ performs one of the following operations: a) convolution with a bank of filters, b) downsample by spatial pooling, and c) non-linear activation; and $\mathcal{D}_i()$ takes actions including: d) convolution with a bank of deconvolution filters, e) upsample by interpolations, and f) non-linear activation. Given a set of T training sample $\{x_1, \dots, x_T\}$, the parameter set of autoencoder $\{w_k, v_k, 0 < k \leq C\}$ is optimized such that reconstruction \tilde{x} resembles input x :

$$\arg \min_{w_k, v_k, 0 < k \leq C} \frac{1}{T} \sum_{i=1}^T \mathcal{L}(x_i, \tilde{x}_i), \quad (2)$$

where \mathcal{L} is a loss function measuring the similarity between x_i and \tilde{x}_i , e.g. MSE.

B. ONE-CLASS SUPPORT VECTOR MACHINE

One-class SVM is an approach for semi-supervised anomaly detection. It models the normal data as a single class that occupies a dense subset of the feature space corresponding to the kernel and aims to find the "normal" regions. A test instance that resides in such a region is accepted by the model whereas anomalies are not [23]. That is, it returns a function for input z that takes the value +1 in the small region capturing most of normal points, and -1 elsewhere. With the training set $\{z_1, \dots, z_T\}$, the dual problem of the one-class SVM solution can be formulated by

$$\min_{\alpha_i, 0 < i \leq T} \frac{1}{2} \sum_{i,j=0}^T \lambda_i \lambda_j k(z_i, z_j) \quad (3)$$

$$\text{s.t. } 0 < \lambda_i \leq \frac{1}{vT}, \quad \sum_{i=0}^T \lambda_i = 1, \quad (4)$$

where λ_i is a Lagrangian multiplier for sample z_i , $k(z_i, z_j) = e^{-\|z_i - z_j\|^2/c}$ is the Gaussian kernel with parameter c , and $v \in (0, 1)$ is a hyper-parameter that controls training errors in this optimization problem [23]. The samples $\{z_i : \lambda_i > 0\}$ are support vectors which lay on the "normal" region boundary. For a new point z , SVM computes the corresponding decision function $\mathcal{G}(z) = \sum_{i=0}^T \lambda_i k(z_i, z) - \rho$ and the label of the new point, y , is evaluated by the function's sign, i.e.

$$y = \text{sgn}(\mathcal{G}(z)) = \text{sgn}\left(\sum_{j=0}^T \lambda_j k(z_j, z) - \rho\right), \quad (5)$$

$$\text{where } \rho = \sum_{j=0}^T \lambda_j k(z_j, z_i), \forall z_i : \lambda_i > 0. \quad (6)$$

It should be noted that performance of one-class SVM strongly depends on the settings of their hyper-parameters ν and c [25]. However, these two parameters are application-dependent, and their settings in an efficient and weak-supervised manner is still an open research problem [26].

C. BREAST CANCER BIOPSY IMAGE SET

The breast cancer benchmark biopsy dataset collected from clinical samples was published by the Israel Institute of Technology [14]. The image set consists of 361 samples, of which 119 were classified by a pathologist as normal tissue, 102 as carcinoma in situ, and 140 as invasive carcinoma. The samples were generated from patients' breast tissue biopsy slides, stained with H&E. They were photographed using a Nikon Coolpix 995 attached to a Nikon Eclipse E600 at magnification of $40\times$ to produce images with resolution of about $5\mu\text{m}$ per pixel. No calibration was made, and the camera was set to automatic exposure. The images were cropped to a region of interest of 760×570 pixels and compressed using the lossy JPEG compression. The resulting images were again inspected by a pathologist in the Institute to ensure that their quality was sufficient for diagnosis.

III. METHODOLOGY

A. SYSTEM OVERVIEW

Given a dataset $\{(\mathcal{I}_1, \dagger_1), (\mathcal{I}_2, \dagger_2), \dots, (\mathcal{I}_K, \dagger_K)\}$ with K samples, where \mathcal{I}_i is an image and $\dagger_i \in \{-1, 1\}$ is a class label indicating whether the corresponding image contains malignant tumor, the goal is to predict the label \dagger for an query image \mathcal{I} and at the same time to generate a probability map indicating suspected abnormal regions in image \mathcal{I} . For simplification, \mathcal{I}_i^- and \mathcal{I}_i^+ denote that \mathcal{I}_i is a normal or malignant image in following sections, respectively. Without loss of generality, assume that there are N normal images and M invasive breast cancer images in the dataset, $0 < N, M < K$ and $N + M = K$, and the normal images are ordered before the malignant ones. That is, the training set is organized as $\{\mathcal{I}_1^-, \mathcal{I}_2^-, \dots, \mathcal{I}_N^-, \mathcal{I}_{N+1}^+, \mathcal{I}_{N+M-1}^+, \mathcal{I}_K^+\}$.

To generate the probability map of cancerous cells, we propose a patch-based learning solution, whose schematic diagram is depicted in Fig. 2. Specifically, for each training image \mathcal{I}_i , we extract l_i overlapping image patches, denoted by $\{x_{i1}, \dots, x_{il_i}\}$. Patches from normal image \mathcal{I}_i^- are assigned label $y_{i,j} = -1$. However, since patches from malignant image \mathcal{I}_i^+ may contain normal tissues only, patch labels $y_{i,j}$ are unknown with a positive constraint that at least one patch contains cancerous cells, i.e. $\max y_{i,j} = 1$ for $0 < j \leq l_i$. If we collect all image patches into a patch set $\{(x_{11}, y_{11}), \dots, (x_{i,j}, y_{i,j}), \dots, (x_{K,l_K}, y_{K,l_K})\}$, then it is evident that in the total $T = T_N + T_M$ patches, the first $T_N = \sum_{i=1}^N l_i$ patches are true-normal ones from the normal histopathology images while the remaining $T_M = \sum_{i=N+1}^K l_i$ patches are from malignant images.

In the training phase, the target is to learn a mapping function $\mathcal{F} : x \rightarrow \{-1, +1\}$ from the training set $\{(x_{i,j}, y_{i,j}), 0 < j \leq l_i, 0 < i \leq K\}$. Since labels of the last T_M patches from malignant images \mathcal{I}_i^+ are unavailable, we make the use of weak-supervised learning methods for discriminative patterns mining to classify image patches. As shown in Fig. 2, an under-complete deep convolutional autoencoder and a one-class SVM, both trained with true-normal patches, are used to implicitly mine dominant patterns in true-normal patches. As a representation method, autoencoder learns the common information in training patches and delivers contrast patterns in its reconstruction residues. Briefly, a normal patch has a low construction error while a malignant patch is with a high residue. Then the trained one-class SVM assigns a label $\{+1, -1\}$ to patch $x_{i,j}$ from malignant image \mathcal{I}_i^+ based on autoencoder's residues. Since the one-class SVM is incapable of generating a probability value, with the obtained decision function and labels generated by the SVM, a 1-layer NN is trained to obtain Platt's score [27] as patch-based posterior probabilities.

In the testing phase, l overlapping patches x_j for $0 < j \leq l$ are extracted from the query image \mathcal{I} . The learnt mapping function \mathcal{F} , achieved by the trained autoencoder and the one-class SVM with the 1-layer NN, is operated on each patch, generating a patch label and a value between $[0, 1]$ indicating the probability that the patch contains malignant tumor. Finally, image classification and a probability map are inferred from obtained patch labels.

B. CONTRAST PATTERN MINING VIA CONVOLUTIONAL AUTOENCODER

Though cell's spatial distribution is the one of the key features for invasive breast cancer diagnosis, this feature is not trivial to quantify. This is because specific structures and patterns of malignant cell clusters differ very much among different tumors and also locally within the same tumor. The incompleteness of local patch labels in this study makes the problem more challenging. Hence in this study, a data-driven solution, specifically, deep convolutional autoencoder, is used to learn the contrast patterns in the training data.

It is noteworthy that an autoencoder is used as a generator of normal patches in this study. In pathology, normal breast tissues usually share certain common patterns, whereas abnormal patterns are highly heterogeneous and features learned from limited quantity of malignant samples may not be descriptive for unseen samples. To overcome this challenge, our method proposes detection of histological abnormalities implicitly by identifying the common patterns in normal breast images. To this end, we make use of the data-specific property of autoencoder and train an autoencoder to learn histological knowledge in true-normal patches.

1) ARCHITECTURE

Since histopathology images are H&E stained and image patches from normal and malignant biopsy images share certain common features, efforts are made to enforce the

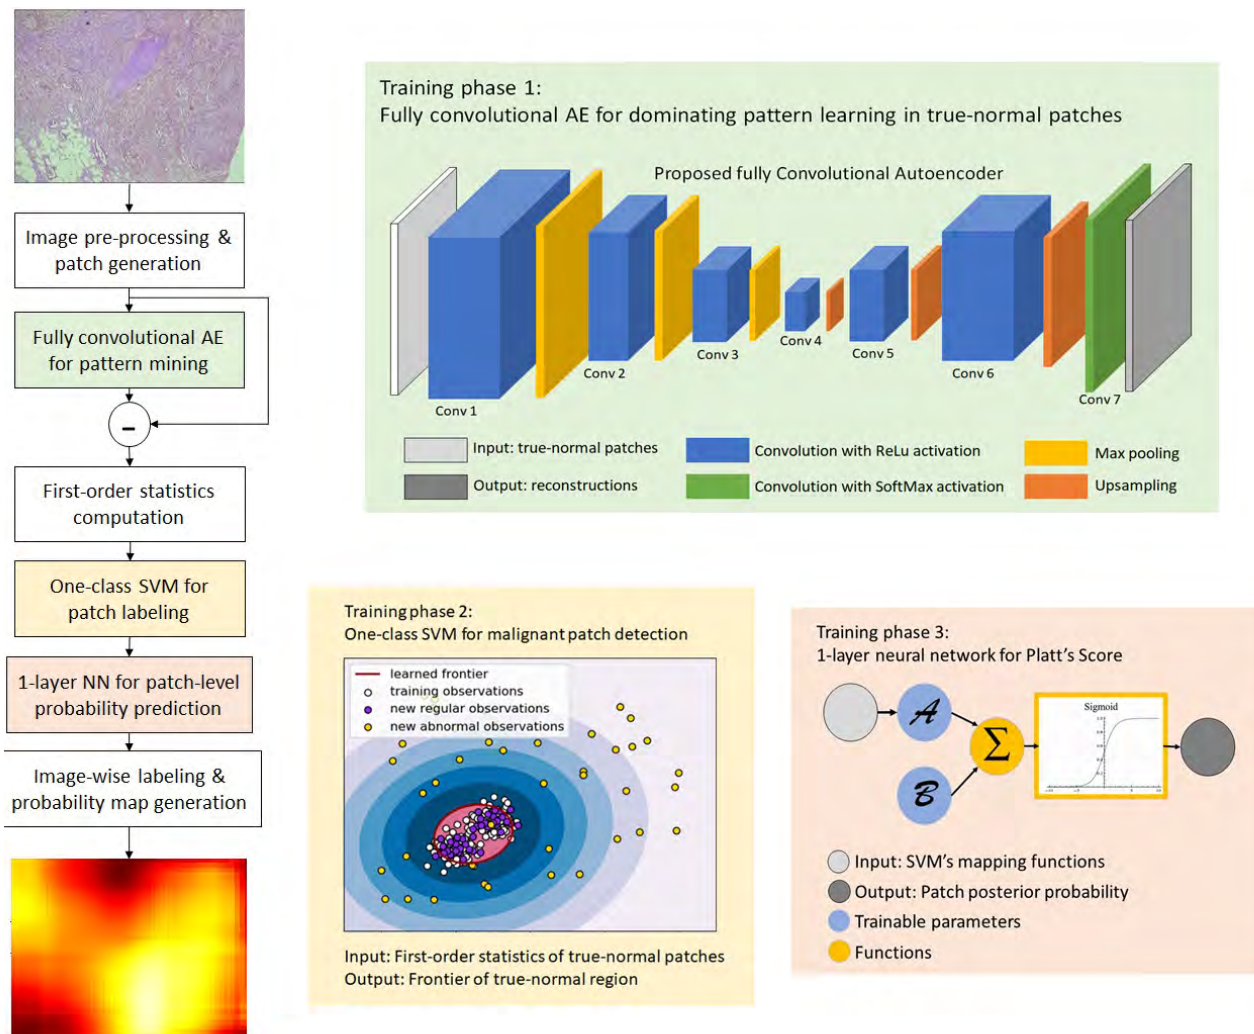


FIGURE 2. Overview of the proposed contrast pattern mining method for invasive breast cancer diagnosis in histopathology images. The output is a probability map of malignant cell clusters in the query image, bright pixel values representing high probability. The whole system is mainly composed of three learning phases. First, a fully convolutional AE is applied to true-normal patches, so that the common patterns shared by true normal patches are learned. Based on AE’s reconstruction residues, we propose to use one-class SVM to learn the regions taken by true-normal patches in the feature space. Finally, the distance to the normal region boundary in the feature space is feed to a 1-layer NN for posterior probability prediction.

autoencoder to learn discriminative structural patterns via designing autoencoder’s architecture. Particularly in this study, the experimental images from the Israel Institute of Technology image set [14] have a magnification of 40x where pixel size is approximately $5\mu m$. Since the diameters of breast epithelial cells’ nuclei stained by H&E are approximately $6\mu m$ [12], nuclei radii are between 1 and 3 pixels. Thus, we design the proposed convolutional autoencoder whose encoder $\mathcal{E}()$ and decoder $\mathcal{D}()$ both are with $C = 6$ and have 3 convolutional layers, such that the nuclei-scale features, nuclei organization features, and the tissue structural-scale features are explored. Table 2 provides detailed architecture of the proposed autoencoder and associates histological features with network layers. Note that the first 6 convolutional layers are composed of the Rectified Linear (Relu) activation unit $Relu(x) = \max(0, x)$ [28]. We select the sigmoid

function $sigmoid(x) = \frac{1}{1+e^{-x}}$ as the activation function in the last convolutional layer to generate a grayscale image in the range of [0, 1].

2) SSIM-BASED LOSS FUNCTION

The loss function \mathcal{L} is the effect driver of the neural network’s learning, and the loss function in an autoencoder network generally defaults to MSE. However, MSE is prone to lead to a smooth/blur reconstruction which may lose some structural information in the original signal [29]. Structural information refers to the knowledge about the structure of objects, e.g. spatially proximate, in the visual scene [22]. Particularly in this study, structural information mainly refers to the spatial organization of cells in H&E stained breast cancer histopathology images. It should be noted that since the multicellular structural information is a key for invasive

TABLE 2. Architecture of the proposed convolutional autoencoder.

	Layer	Layer type	Filter Size	Activation	Input/Output Dimension	Histological association with breast cancer images in magnification of 40x
Encoder $\mathcal{E}()$	0	input			$256 \times 256 \times 1$	
	1	Convolutional	$3 \times 3 \times 16$	Relu	$256 \times 256 \times 16$	nuclei & edge
	2	Max-polling	2×2		$128 \times 128 \times 16$	
	3	Convolutional	$3 \times 3 \times 8$	Relu	$128 \times 128 \times 8$	nuclei organization
	4	Max-polling	2×2		$64 \times 64 \times 8$	
	5	Convolutional	$3 \times 3 \times 8$	Relu	$64 \times 64 \times 8$	Structure & tissue organization
6	Max-polling	2×2		$32 \times 32 \times 8$		
7	Convolutional	$3 \times 3 \times 8$	Relu	$32 \times 32 \times 8$		
8	Upsampling	2×2		$64 \times 64 \times 8$		
Decoder $\mathcal{D}()$	9	Convolutional	$3 \times 3 \times 8$	Relu	$64 \times 64 \times 8$	nuclei organization
	10	Upsampling	2×2		$128 \times 128 \times 8$	
	11	Convolutional	$3 \times 3 \times 16$	Relu	$128 \times 128 \times 16$	nuclei & edge
	12	Upsampling	2×2		$256 \times 256 \times 16$	
	13	Convolutional/Output	$3 \times 3 \times 1$	Sigmoid	$256 \times 256 \times 1$	

breast cancer diagnosis [12], it should be learned and maintained in autoencoder's output. To this end, we make use of the SSIM index to compose the loss function for AE's training, i.e.

$$\mathcal{L}(x_{i,j}, \tilde{x}_{i,j}) = 1 - SSIM(x_{i,j}, \tilde{x}_{i,j}), \quad (7)$$

which facilitate the autoencoder to maintain structural information in image patches. SSIM index is defined as

$$SSIM(x, \tilde{x}) = [l(x, \tilde{x})]^\alpha \times [c(x, \tilde{x})]^\beta \times [s(x, \tilde{x})]^\gamma, \quad (8)$$

where $l()$, $c()$, and $s()$ are the luminance comparison function, contrast comparison function, and structural comparison functions, respectively. α , β , and γ are used to adjust the relative importance of the three components.

3) PATTERN LEARNING WITH TRUE-NORMAL PATCHES

It is noteworthy that in this study, instead of training the network with all training patch $\{x_{11}, \dots, x_{K,L_K}\}$, the autoencoder is trained with only true-normal patches $\{x_{11}, \dots, x_{N,L_N}\}$. The motivation behind this innovation is the data-specific property of autoencoder. That is, an autoencoder has low reconstruction errors for samples following training data's generating distribution, while having a large reconstruction error otherwise. Specifically, in our study, the autoencoder learns the common properties and dominant patterns among true-normal patches. Thus, the trained autoencoder is capable of recovering the common content shared by query and true-normal patches. The smaller the residue is, the similar the query patch is to true-normal patches. However, since patches containing invasive tumor have some distinct patterns so that the autoencoder cannot represent well, large construction residue is generated. In other words, the discriminative and contrast patterns in this problem are embedded in autoencoder's residues Δx , which is quantified by the absolute value of the difference between AE's input and output.

$$\Delta x = |x - \tilde{x}| = |x - \mathcal{D}(\mathcal{E}(x))|. \quad (9)$$

To facilitate downstream patch labeling, discriminative patterns embedded in Δx are summarized by several compact

numerical descriptors. Motivated by the radiomics analysis [30], we compute 16 patch-wise first-order statistics to describe the distribution of intensities within AE's residue Δx , which are energy, minimum, maximum, 10^{th} percentile, 90^{th} percentile, mean, median, interquartile range, full range, mean absolute deviation, robust mean absolute deviation, variation, skewness, kurtosis, entropy, and histogram uniformity. We denote the obtained numerical feature set by $\{z_{11}, \dots, z_{K,L_K}\}$, where $z_{i,j}$ is the description vector of the residue corresponding to patch $x_{i,j}$ and the first T_N elements in the feature set come from the true-normal patches.

C. PATCH LABELING BY ONE-CLASS SVM

With the discriminative representation $\{z_{i,j}\}$ generated by a deep convolutional autoencoder, we precede to investigate the mapping function $\mathcal{F} : z_{i,j} \rightarrow \{-1, +1\}$ from numerical features of true-normal patches. Note that for malignant images with $\dagger_i = 1$, patch labels are not necessarily consistent with image labels; in addition, the number of patches having malignant tumor may be much smaller compared to the quantity of true-normal patches in the training set. As only true-normal patches with their labels are reliable, we cast the problem into the problem of semi-supervised anomaly detection. Intuitively, if one can find regions in the feature space where true-normal patches cluster, patches falling out of the "normal" regions are highly likely to be abnormal.

Due to the good performance of one-class SVM in medical anomaly detection [31], [32], we select one-class SVM to approximate the distribution of true-normal patches for anomaly patch detection. It should be noted that in one-class SVM, normal patches are labeled as +1, which is opposite to the patch labels defined in this study. Hence, based on the features of true-normal patches $\{z_{i,j} : 0 < i < T_n\}$, a patch label can be obtained using a mapping function

$$\mathcal{F}(z) = -\text{sgn}(\mathcal{G}(z)) = -\text{sgn}\left(\sum_{i=0}^{T_n} \lambda_i k(z_i, z) - \rho\right), \quad (10)$$

where ρ is defined in (6).

D. MALIGNANT TUMOR PROBABILITY MAP GENERATION

Now we obtain a labeled training set $\{(z_{i,j}, y_{i,j}) : 0 < j \leq l_i, 0 < i \leq T\}$, where the last T_M samples have labels generated by the one-class SVM. It should be noted that in precision medicine, patches having cancerous tissues definitely (locating far from SVM's hyperplane) and those suspected containing abnormality (residing near SVM's hyperplane) should be distinguished. However, SVM does not provide a posterior probability $p(y_{i,j}|z_{i,j})$ and the resulting labels cannot differentiate these cases. Hence, for any data sample $z_{i,j}$, we make use of SVM's decision function $\mathcal{G}(z_{i,j})$ to compute Platt's score for posterior probability approximation [27]. Platt's score is defined as a sigmoid function on SVM's decision function. That is,

$$p(y_{i,j}|z_{i,j}) \approx p(y_{i,j}|\mathcal{G}(z_{i,j})) = \frac{1}{1 + e^{(A\mathcal{G}(z_{i,j})+B)}}, \quad (11)$$

where A and B are parameters trained using sample labels.

Examining the Platt's score in (11), we notice that Platt's score can be implemented by a 1-layer neural network with a sigmoid activation function for two reasons. First, from a theoretical point of view, sigmoid function is a good candidate to generate a probability from a real value (i.e. SVM's decision function $\mathcal{G}(z_{i,j})$ in this study) because its output can be interpreted as the posterior probability for the most general categorical distribution: Bernoulli distribution. Second, from the application point of view, though there is a "−" sign difference between Platt's score in (11) and the standard sigmoid function in machine learning, the sign difference can be easily compensated by the trainable parameters in deep learning. Consequently, the recursive optimization of Platt's score is realized by training a 1-layer NN with sigmoid activation function. For the 1-layer NN, input is one-class SVM's decision function. Parameters of the network, a 1×1 transformation matrix A and a bias B , can be optimized using training set $\{(\mathcal{G}(z_i), y_i) : 0 < i \leq T\}$.

To infer an image label from obtained patch labels, patch majority voting, where the image label is selected as the most common patch label, is the most common method in literature of breast cancer histopathology image diagnosis [11], [12]. However, in clinical practice, any abnormalities, suspect lesions in particular, should trigger an alarm. Based on this belief, instead of using majority voting, we propose a much stricter rule to combine patch diagnosis results, that is, an image is labeled as benign when all patches are classified as normal. Finally, we proceed to generate a probability map of abnormality for an image. With the obtained probability $p(y_i|z_i)$, if an image pixel is only contained in one patch, the probability is assigned to the pixel directly. Otherwise (a pixel is in multiple overlapping patches), the probability at the pixel is obtained by averaging probability values of the overlapping patches.

IV. SYSTEM IMPLEMENTATION AND TRAINING DETAILS

A. SYSTEM IMPLEMENTATION AND HYPER-PARAMETER SETTING

The proposed method is implemented using python 3.6.6. Each histopathology image is normalized using the illuminant normalization method [33]. Then the normalized image is converted to the grayscale version and rescaled to $[0, 1]$. The two-step pre-processing mitigates the effect of color variations usually observed in histopathology images on downstream discriminative pattern mining. The autoencoder and the computation of Platt's score are realized using the Keras library which uses TensorFlow as its backend. The One-class SVM is called from the scikit-learn library.

To compute the SSIM index when training autoencoder, we use the Gaussian filter with size 11×11 to smooth image patches and fuse the luminance comparison function, contrast comparison function, and structural comparison functions with $\alpha = \beta = \gamma = 1$ following SSIM's original paper [22].

One-class SVM has application-dependent hyper-parameters, ν and c . In this study, we propose the use of the whole training set to select the optimal hyper-parameters. Given a pair of ν and c , after training over true-normal patches, one-class SVM generates a label for each training patch. Though we cannot directly assess patch classification accuracy due to the lack of patch annotation, we can correspond image labels to evaluate the SVM model indirectly. Specifically, all patches from normal images are normal and an image is labeled as malignant when at least one of its patches contains cell's malignant growth pattern, i.e. $\max y_{i,j} = \dagger_i, \forall j \in [1, l_i]$. Hence, the one-class SVM's classification accuracy in image level, denoted by ACC_{img} , can be measured by

$$ACC_{img} = \frac{1}{T} \sum_{i=1}^T \delta(\dagger_i - \max_j y_{i,j}), \quad (12)$$

where $\delta(\cdot)$ is the Dirac delta function, i.e. $\delta(x) = 1$ for $x = 0$ and $\delta(x) = 0$ otherwise. By comparing all obtained ACC_{img} , the one-class SVM with highest image classification accuracy is selected, i.e.

$$\nu_{opt}, c_{opt} = \arg \max_{\nu, c} ACC_{img}. \quad (13)$$

B. TRAINING DATA AUGMENTATION

For histopathology images in the training set, each pre-processed image is divided into 35 patches, each having 256×256 pixels with 30% overlap at most. Different from conventional data augmentation methods that generates a fixed augmented training set, data augmentation in this study is performed in an online manner with the support of Keras. Specifically, to learn a rotation-invariant AE network, data augmentation operations in this study include patch rotation with an angle randomly drawn from $[0, 180)$ degrees, vertical reflection, and horizontal flip. At each learning epoch, transformations with randomly selected parameters among the augmentation operations are generated and applied to original

training patches. Then the augmented patches are feed to the network. When the next learning epoch is started, the original training data is once again augmented by applying specified transformations. That is, the number of times each training data is augmented is equal to the number of learning epochs. In this way, the AE network almost never sees two exactly identical training patches, because at each epoch training patches are randomly transformed. For example, with the breast cancer image data set used in this study, we got a basic set of 3750 true-normal patches for AE's training. After 100 epoch learning, the network had seen 375,000 augmented patches in total. We believe the online augmented method helps to fight against network's over-fitting in this study.

C. NETWORK INITIALIZATION AND TRAINING

For both the deep convolutional autoencoder and 1-layer neural network for Platt's score, we initialized all weights with zero-mean Xavier uniform random numbers [34]. All biases were set to zero. The networks were trained using Adam stochastic optimization with learning rate 0.001, and the exponential decay rates for the first and second moment estimations are set to 0.9 and 0.999. To enforce the autoencoder to learn the dominant patterns in true-normal patches, the training ran 100 epochs. The 1-layer neural network for Platt's score was trained with 25 epochs. We used 10% of the training data for validation. The optimal networks for autoencoder and Platt's score were selected based on the proposed SSIM-based loss function and the binary-classification cross-entropy on the validation sets, respectively.

V. EXPERIMENTATION

In this study, the proposed method is evaluated using the 119 images of the morphologically normal breast tissue and the 140 images of invasive breast cancer in the breast cancer benchmark data set published by the Israel Institute of Technology¹. We will first compare image patches reconstructed by autoencoder with loss functions of SSIM and MSE. Then we qualitatively assess the effectiveness of contrast pattern mining by visualizing obtained features and their distribution in a manifold space. Finally image classification and the obtained abnormality maps are examined and compared to prior arts.

A. PATCH RECONSTRUCTION USING SSIM AND MSE

The proposed method exploits an autoencoder as a generator of normal patches. A loss function should be selected such that the autoencoder can reconstruct a normal patch as much as possible. In this experiment, we compare reconstructed patches generated by different loss functions, SSIM and MSE, and quantify their effects on normal patch generation. After 100-epoch training over 4165 patches generated from the 119 normal breast tissue images, energies of reconstruction residues over all true-normal patches are calculated and averaged. Specifically for SSIM and MSE based loss functions, the average energies per patch are 194.756 and 219.785, respectively. That is, the SSIM-based function drives the




FIGURE 3. Comparison of patch reconstruction using different loss functions, SSIM and MSE. The SSIM-driven reconstructions are sharper than the MSE-driven images.

autoencoder to learn more from its inputs. Fig. 3 presents examples of patch reconstructions associated with loss functions of SSIM and MSE, where reconstructed patches associated with MSE is more blurred.

B. VISUALIZATION OF CONTRAST PATTERN MINING

A fully convolutional autoencoder net is used to mine the common patterns in normal histopathology image patches. Fig. 4 presents several examples of autoencoder's input (in the left column) and their corresponding reconstruction residues (in the right column), where AE's residues are represented as heatmaps for visualization. As shown in the figure, residue images (e)-(f) that correspond to the malignant patches (a)-(b) have brighter values; on contrast, the true-normal patches (c)-(d) have relatively small reconstruction errors (g)-(h). Thanks to the data-specific property of the autoencoder, the dominant patterns in normal image patches are well summarized, while the abnormal patterns among malignant breast cancer images are maintained in AE's residues.

To visualize the distribution of learnt contrast patterns, we project the obtained high dimensional feature set into 2-D domain via T-SNE [35] and illustrated in Fig. 5(a). A green sample is associated with a true-normal patch and a red sample represents a patch from a malignant image. From the figure, more than half red samples share different characteristics from green samples. Fig 5(b) visualizes




FIGURE 4. Examples of the deep autoencoder’s inputs (256×256 grayscale patches from images with a magnification of $40\times$) (a)-(d) and their reconstruction residues (e)-(h), where brighter values in the residue heatmaps represent large reconstruction errors. The upper two patches from malignant images (a)-(b) contain abnormal cell growth patterns, and the lower two (c)-(d) are extracted from normal breast tissue histopathology images.

the performance of one-class SVM, where green sample corresponds to a true-normal patch, while yellow and red samples are associated with patches from malignant images in the 2-D T-SNE domain. The difference is that yellow samples are classified as normal by the one-class SVM, while red data represent patches that are labeled as containing malignant cell clusters.

C. BREAST CANCER IMAGE CLASSIFICATION

1) EVALUATION PROTOCOL

To evaluate the proposed method, stratified 10-fold cross-validation is performed. Specifically, the image set is

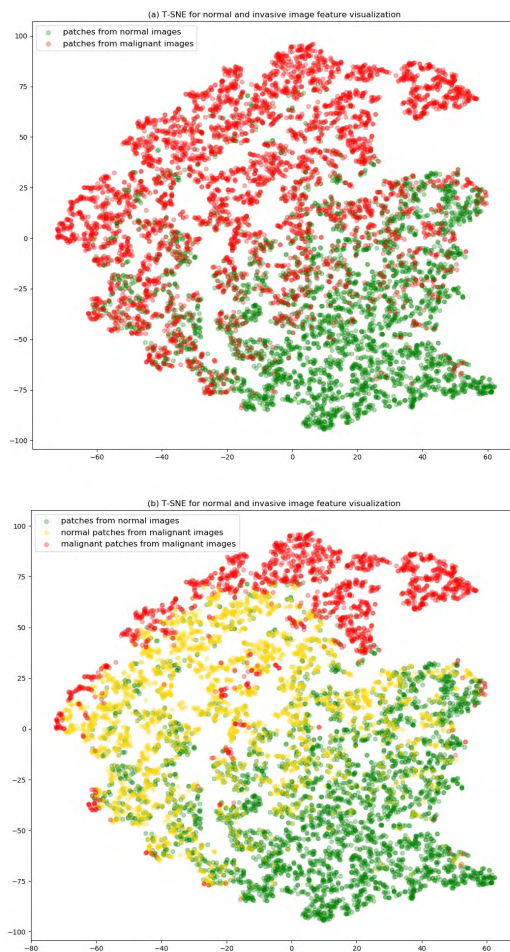


FIGURE 5. Feature set visualization via T-SNE. (a) Some patches from malignant images overlap with true-normal patches in the low-dimensional T-SNE domain. (b) The yellow samples represent patches that are extracted from malignant images and classified as normal by the one-class SVM.

randomly partitioned into 10 equal-size folds, where each fold contains roughly the same proportions of normal and malignant labels. The cross-validation is repeated 10 times where each fold is used as the test set once and images in remaining 9 folds are used as training data. Then the obtained 10 diagnosis results are averaged to estimate classification performance. In each round of cross-validation, images in the training set are processed as described in the section of data augmentation. In the testing phase, image patches are extracted every 16 pixels, i.e. the centers of two patches may be only 16 pixels apart in an image. The distance of 16 pixels is a trade-off between generating a fine probability map and maintaining computation efficiency.

In this experiment, we first perform a quantitative evaluation on image classification. Particularly, to measure image classification performance, we use the most common medical diagnosis assessments, which include classification accuracy $ACC \in [0, 1]$, F-measure score $F_1 \in [0, 1]$, positive/negative likelihood ratios $LR+ \in [1, \infty)$ and $LR- \in [0, 1]$, and diagnostic odds ratio $DOR \in [1, \infty)$. ACC is one of the most common classification performance measurements.

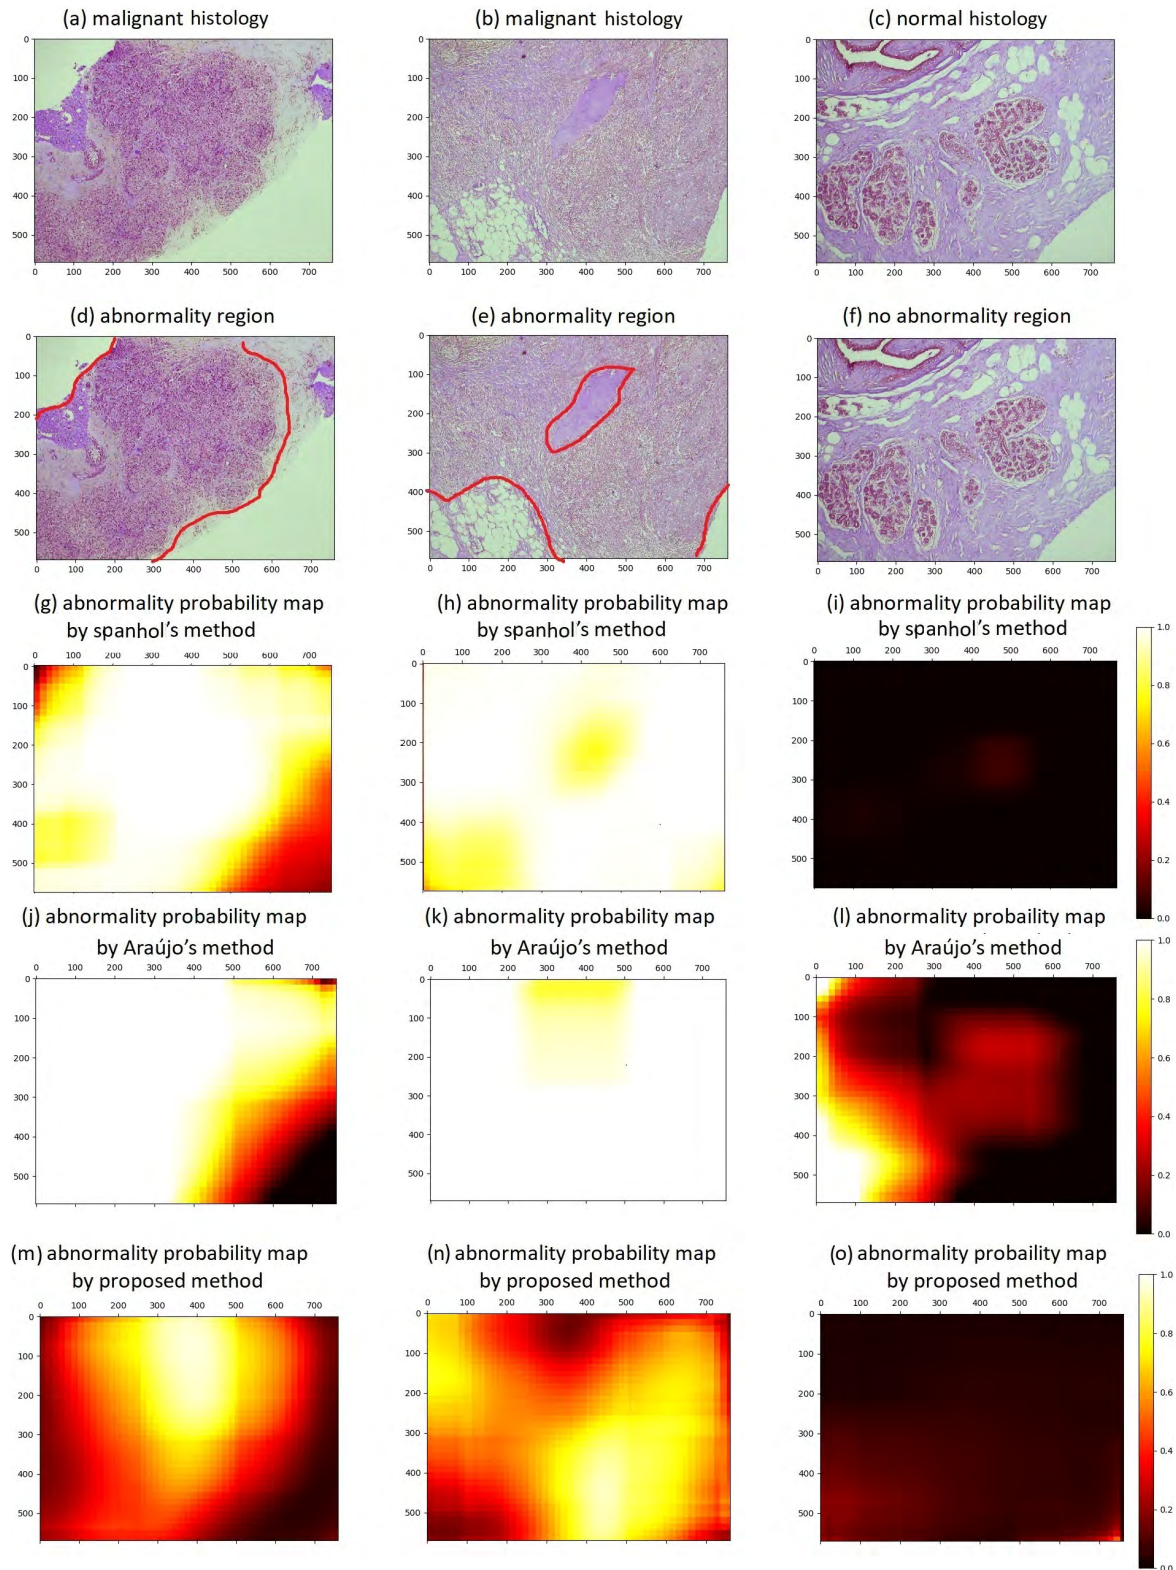


FIGURE 6. Examples of breast tissue images and their corresponding abnormality probability maps, where probability value greater than 0.5 indicates abnormalities in this study. In images in the second row, abnormality regions derived by malignant cell growth were delineated red. (a) Malignant histology. (b) Malignant histology. (c) Normal histology. (d) Abnormality region. (e) Abnormality region. (f) No abnormality region. (g) Abnormality probability map by Spanhol's method. (h) Abnormality probability map by Spanhol's method. (i) Abnormality probability map by Spanhol's method. (j) Abnormality probability map by Araújo's method. (k) Abnormality probability map by Araújo's method. (l) Abnormality probability map by Araújo's method. (m) Abnormality probability map by proposed method. (n) Abnormality probability map by proposed method. (o) Abnormality probability map by proposed method.

TABLE 3. Image-level classification evaluation. The sign * indicates that the performance difference between the prior method and the proposed method is of statistical significance at the 5% significance level.

	Spanhol's [11]	Araujo's [12]	proposed
ACC	0.700*	0.710	0.760
F_1	0.766*	0.763*	0.777
$LR+$	1.540*	1.554*	2.645
$LR-$	0.229*	0.281*	0.304
DOR	8.563	8.132	12.876

It represents the proportion of accurate diagnoses, but it is impacted by disease prevalence. F_1 is the harmonic average of the precision and sensitivity and $F_1 = 1$ corresponds to a perfect binary classification. Likelihood ratios use the sensitivity and specificity of the test to determine its diagnostic performance. They are believed as good institutions of AUC when ROC analysis is infeasible. DOR combines sensitivity and specificity and equals to the ratio of positive and negative likelihood ratio. Among the five measurements, F_1 score, likelihood ratio, and DOR are independent of test prevalence, with higher values indicating a better discriminative performance. Since the breast cancer image set does not delineate the specific locations of malignant cell clusters, we perform a qualitative assessment of the obtained probability maps by comparing it to abnormality regions derived by malignant cell growth.

2) OTHER APPROACHES

To the best of our understanding, the proposed method constitutes the first attempt in literature of breast cancer diagnosis to infer locations of abnormalities from image labels. Since there is no such breast cancer diagnosis study in literature, we compare the performance of the proposed method to the latest patch-based deep-learning breast cancer histopathology image classification methods proposed by Spanhol *et al.* [11] and Araujo *et al.* [12]. Spanhol's method divides an image into 64×64 image patches and uses an 8-layer convolutional neural network to classify image patches. Then three fusion rules, majority voting, malignant patch detection (i.e. maximum probability), and sum of probability, are used to obtain the final image classification. Araujo's method divides images into 512×512 patches and enforces training-patch labels consistent with image labels. Based on training patches and their newly-assigned labels, a 13-layer convolutional neural net is trained and used to classify unseen image patches. The final image classification is also achieved by combining all inferred patch labels using one of the three rules used in Spanhol's method. Noted that in both studies, malignant patch detection is reported to achieve worse performance than the other two fusing rules for image diagnosis. However, based on the belief that in a medical alert system, any suspected alterations should trigger an alarm, we select the fusing rule of malignant patch detection for both image classification methods in this comparison experiment.

3) RESULTS AND DISCUSSIONS

Table 3 lists the image classification performance. The sign * indicates that the performance difference between the prior method and the proposed method is of statistical significance at the 5% significance level. The better performance of the proposed method is mainly contributed by its practicality and generalizability in contrast pattern mining. First, the minimum prior knowledge to train the proposed method is WSI label which can be easily acquired in clinic practice. But this information is not enough to train Spanhol's method [11] and Araujo's method [12]; instead, patch-label or even pixel-wise annotation is required. It should be noted that these prior deep models have more than ten thousands of trainable parameters. Acquisition of sufficient amount of well-labeled data for these models is fairly prohibitive, if not impossible, in practice due to the expensive and time-consuming properties of pathological annotations. To address the shortage of well-labeled training patches for supervised learning, Araujo's method makes an assumption that patch-labels are identical to their image labels [12]. However, this assumption hardly holds in practice because a tumor usually takes 0.01%-70% (median 2%) areas of a WSI image [36]. The less practical requirement/assumption on training data in prior arts limits their diagnosis performance in practice. Second, in pathology, normal breast tissues usually share common histological patterns, whereas structures and patterns of malignant cell clusters are heterogeneous. Consequently, quantification of the normal patterns is relatively feasible, but representation learning among histological irregularities is more challenging. In prior studies, cell's abnormal patterns are usually learnt directly from training samples. However, due to the limited amount of training data, variations in histological abnormalities may not be fully represented. As a result, the generalizability of these deep diagnosis models to unseen malignant cancer images is still in question. To overcome the challenge of abnormality representation in digital pathology, the proposed method learns the common patterns in normal breast images first and diagnoses malignant cells by similarity of these cells and their growth patterns to normal ones. In this way, detection of histological abnormalities is simplified to identification of common patterns in normal breast images. Consequently, the proposed method is less dependent on specific malignant image samples and can generalize well.

Examples of the obtained probability maps with their corresponding H&E images are demonstrated in Fig. 6. Abnormality regions derived by malignant cell growth in the query images were delineated by our senior pathologist and highlighted in images at the second row. A probability map is presented in the form of a heat-map where bright pixels represent high probabilities of abnormalities. It provides an insight and verification of the image diagnosis result by inferring locations of abnormalities in an image. In this sense, it even conveys more information compared to the classification result itself. Since Spanhol's's approach and Araujo's method are

TABLE 4. Comparison summary of examined methods.

	Spanhol's method [11]	Araujo's method [12]	proposed method
Training procedure	1 end-to-end training	1 end-to-end training	3 training phases
Practicality	no	no	yes
Generalizability	less generalizable	less generalizable	generalizable
classification accuracy	fair	fair	better
self-interpretability	no	no	yes

also based on patch processing, as a comparison, the obtained patch-level probabilities are used to form the corresponding probability maps following the method proposed in this study. As shown in the figure, the two prior methods are prone to yield large probabilities in background areas of invasive breast cancer histopathology images.

In summary, the advantages of the proposed method are contributed by its practicality, generalizability, and self-interpretability. First, the minimal prior knowledge on whoselide-imaging (WSI) labels for system training is easily acquired in clinic practice. Second, the proposed method detects discriminative patterns in images in weak-supervised manner. Because the method is less dependent on specific malignant image samples, it generalizes well on unseen images. Third, the obtained probability map infers locations of abnormalities in an image. The insightful information explains the final diagnosis result and helps pathologists to verify the diagnosis reasoning. Table 4 summarizes a comprehensive comparison of examined methods. The major limitation of this study is the size of the experimental image set and the absence of the external validation group. However, the carefully designed experimentation and the involvement of pathologist's expertise in this study support the reliability of the obtained results. In addition, the public-accessibility of the experimental image set facilitates other scholars to reproduce our study.

VI. CONCLUSION

In this study, we presented a discriminative pattern mining approach for invasive carcinoma diagnosis in routine H&E stained breast tissue histopathology images. By learning contrast patterns between normal and malignant breast cancer images, the proposed method was capable to identify suspected regions of malignant cell clusters in an image. The evaluation was conducted on a public histopathology image set and experimentation demonstrated that the proposed method outperformed prior arts. Particularly, the superiority of the proposed method was its practicality, generalizability, and self-interpretability. The obtained probability map would facilitate a better understanding of the proposed pattern mining and diagnosis solution.


In this study, heterogeneity of histological abnormalities posed a big challenge in pattern mining. we noted that there was still a large room to improve the diagnosis performance by investigating more efficient pattern mining methods. On the other hand, though we tried to tackle the problem of self-interpretability in machine learning and proposed a

diagnosis system which was capable to generate visualized information to support and verify its decision, the black-box property of deep learning in terms of data representation was still less-touched. Following the work of CAM [37] and Grad-CAM [38], we would investigate how to interpret the internal reasoning of a deep diagnosis system in future.

REFERENCES

- [1] D. Dillon, A. Guidi, and S. Schnitt, "Pathology of invasive breast cancer," in *Diseases Breast*, J. Harris, M. Lippman, M. Morrow, and C. Osborne, Eds. Philadelphia, PA, USA: Lippincott Williams Wilkins, 2010, ch. 8, pp. 374–407.
- [2] S. Robertson, H. Azizpour, K. Smith, and J. Hartman, "Digital image analysis in breast pathology—from image processing techniques to artificial intelligence," *Transl. Res.*, vol. 194, pp. 19–35, Apr. 2018.
- [3] M. Veta, J. Pluim, P. van Diest, and M. Viergever, "Breast cancer histopathology image analysis: A review," *IEEE Trans. Biomed. Eng.*, vol. 61, no. 5, pp. 1400–1411, May 2014.
- [4] P. Filipczuk, T. Fevens, A. Krzyzak, and R. Monczak, "Computer-aided breast cancer diagnosis based on the analysis of cytological images of fine needle biopsies," *IEEE Trans. Med. Imag.*, vol. 32, no. 12, pp. 2169–2178, Dec. 2013.
- [5] Y. M. George, H. H. Zayed, M. I. Roushdy, and B. M. Elbagoury, "Remote computer-aided breast cancer detection and diagnosis system based on cytological images," *IEEE Syst. J.*, vol. 8, no. 3, pp. 949–964, Sep. 2014.
- [6] M. Kandemir, C. Zhang, and F. Hamprecht, "Empowering multiple instance histopathology cancer diagnosis by cell graphs," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent*, 2014, pp. 228–235.
- [7] S. H. Bhandari, "A bag-of-features approach for malignancy detection in breast histopathology images," in *Proc. IEEE Int. Conf. Image Process.*, Sep. 2015, pp. 4932–4936.
- [8] W. Li, J. Zhang, and S. J. McKenna, "Multiple instance cancer detection by boosting regularised trees," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent*, 2015, pp. 645–652.
- [9] X. Li and K. Plataniotis, "Toward breast cancer histopathology image diagnosis using local color binary pattern," in *Proc. 14th Annu. Imaging Netw. Ontario Symp.*, Sep. 2016, pp. 25–36.
- [10] F. Spanhol, L. S. Oliveira, C. Petitjean, and L. Heutte, "A dataset for breast cancer histopathological image classification," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 7, pp. 1455–1463, Jul. 2016.
- [11] F. A. Spanhol, L. S. Oliveira, C. Petitjean, and L. Heutte, "Breast cancer histopathological image classification using convolutional neural networks," in *Proc. Int. Joint Conf. Neural Netw.*, Jul. 2016, pp. 2560–2567.
- [12] T. Araújo et al., "Classification of breast cancer histology images using convolutional neural networks," *PLoS ONE*, vol. 12, no. 6, Jun. 2017, Art. no. e0177544.
- [13] F. A. Spanhol, P. R. Cavalin, L. S. Oliveira, C. Petitjean, and L. Heutte, "Deep features for breast cancer histopathological image classification," in *Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC)*, Oct. 2017, pp. 1868–1873.
- [14] *Breast Cancer Imageset*. [Online]. Available: ftp://ftp.cs.technion.ac.il/pub/projects/medicimage/breast_cancer_data/
- [15] A. Cruz-Roa et al., "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks," *Proc. SPIE*, vol. 9041, 2014, Art. no. 904103.
- [16] K. Sirinukunwattana, S. E. A. Raza, Y.-W. Tsang, D. R. J. Snead, I. A. Cree, and N. M. Rajpoot, "Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1196–1206, May 2016.

- [17] A. Cruz-Roa et al., "Accurate and reproducible invasive breast cancer detection in whole-slide images: A deep learning approach for quantifying tumor extent," *Sci. Rep.*, vol. 7, Apr. 2017, Art. no. 46450.
- [18] R. Bidart et al., "Localization and classification of cell nuclei in post-neoadjuvant breast cancer surgical specimen using fully convolutional networks," *Proc. SPIE*, vol. 10581, 2018, Art. no. 105810O.
- [19] M. Peikari, S. Salama, S. Nofech-Mozes, and A. Martel, "A cluster-then-label semi-supervised learning approach for pathology image classification," *Sci. Rep.*, vol. 8, no. 1, p. 7193, May 2018.
- [20] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
- [21] M. A. Ponti, L. S. F. Ribeiro, T. S. Nazare, T. Bui, and J. Collomosse, "Everything you wanted to know about deep learning for computer vision but were afraid to ask," in *Proc. 30th SIBGRAPI Conf. Graph., Patterns Images Tuts.*, Oct. 2017, pp. 17–41.
- [22] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [23] B. Schölkopf, J. C. Platt, J. C. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Comput.*, vol. 13, no. 7, pp. 1443–1471, 2001.
- [24] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Sep. 2015, pp. 1–8.
- [25] Z. Ghafoori, S. M. Erfani, S. Rajasegarar, J. C. Bezdek, S. Karunasekera, and C. Leckie, "Efficient unsupervised parameter estimation for one-class support vector machines," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 10, pp. 5057–5070, Oct. 2018.
- [26] W. Liu, G. Hua, and J. R. Smith, "Unsupervised one-class learning for automatic outlier removal," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 3826–3833.
- [27] J. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," *Adv. Large Margin Classifiers*, vol. 10, no. 3 pp. 61–74, Mar. 1999.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NIPS*, vol. 25, 2012, pp. 1106–1114.
- [29] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Trans. Comput. Imaging*, vol. 3, no. 1, pp. 47–57, Mar. 2017.
- [30] J. Griethuysen et al., "Computational radiomics system to decode the radiographic phenotype," *Cancer Res.*, vol. 77, no. 21, pp. e104–e107, Nov. 2017.
- [31] S. Dreiseitl, M. Osl, C. Scheibböck, and M. Binder, "Outlier detection with one-class SVMs: An application to melanoma prognosis," in *Proc. Amer. Med. Informat. Assoc. Annu. Symp.*, 2010, p. 172.
- [32] P. Cichosz et al., "Novelty detection for breast cancer image classification," *Proc. SPIE*, vol. 10031, Sep. 2016, Art. no. 1003135.
- [33] X. Li and K. N. Plataniotis, "A complete color normalization approach to histopathology images using color cues computed from saturation-weighted statistics," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 7, pp. 1862–1873, Jul. 2015.
- [34] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. 13th Int. Conf. Artif. Intell. Statist.*, Mar. 2010, pp. 249–256.
- [35] L. van der Maaten and G. Hinton, "Visualizing high-dimensional data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.
- [36] Y. Liu et al. (2017). "Detecting cancer metastases on gigapixel pathology images." [Online]. Available: <https://arxiv.org/abs/1703.02442>
- [37] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Sep. 2016, pp. 2921–2929.
- [38] R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis.*, Sep. 2017, pp. 618–626.



XINGYU LI (M'18) received the B.Sc. degree in electronics engineering from Peking University, China, in 2007, the M.Sc. degree in electrical and computer engineering from the University of Alberta, Canada, in 2009, and the Ph.D. degree in electrical and computer engineering from the University of Toronto, Canada, in 2017. In 2012, she joined the Multimedia Lab, University of Toronto, Canada, where she is currently a Postdoctoral Research Fellow. Her research interests include color image processing, signal processing for big data, machine learning, and statistical pattern recognition, with applications to computational health informatics.



MARKO RADULOVIC received the Ph.D. degree in immunology from the Institute for Oncology and Radiology of Serbia, Belgrade, Serbia, where he is currently a Research Professor. He was with the Max Planck Institute for Experimental Medicine, Germany, University College London, and Imperial College. His research interests include computational analysis of tumour histomorphology and investigation of molecular cell dynamics by use of high-throughput proteomics.



KSENIJA KANJER received the Ph.D. degree in oncology from the Institute for Oncology and Radiology of Serbia, Belgrade, Serbia, where she is currently a Research Associate. After completed residency in pathology and master's exam for pathologist, she is employed as the Consultant Pathologist at the Laboratory for receptors and growth of malignant tumors—Department of Experimental Oncology. Her research interest includes the regulation of malignant cell growth and death, along with the identification of biomarkers for disease progression and therapy resistance, most notably in breast cancer.



KONSTANTINOS N. PLATANIOTIS (S'93–M'95–SM'03–F'12) is currently a Professor and the Bell Canada Chair in multimedia with the ECE Department, University of Toronto.

Dr. Plataniotis is a registered Professional Engineer in Ontario and a Fellow of the Engineering Institute of Canada. He was the Technical Co-Chair for ICASSP 2013 and the General Co-Chair for the 2017 IEEE GlobalSIP. He co-chairs the 2018 IEEE International Conference on Image Processing and the 2021 IEEE International Conference on Acoustics, Speech, and Signal Processing. He has served as a Signal Processing Society Vice President for Membership, from 2014 to 2016, and as the Editor-in-Chief for the IEEE SIGNAL PROCESSING LETTERS, from 2009 to 2011.

...