

Case-Based Histopathological Malignancy Diagnosis using Convolutional Neural Networks

Qicheng Lao
qi_lao@encs.concordia.ca
Thomas Fevens
fevens@encs.concordia.ca

Department of Computer Science and Software Engineering,
Concordia University,
Montréal, Québec, Canada

Abstract

In practice, histopathological diagnosis of tumor malignancy often requires a human expert to scan through histopathological images at multiple magnification levels, after which a final diagnosis can be accurately determined. However, previous research on such classification tasks using convolutional neural networks primarily determine a diagnosis for a single magnification level. In this paper, we propose a case-based approach using deep residual neural networks for histopathological malignancy diagnosis, where a case is defined as a sequence of images from the patient at all available levels of magnification. Effectively, through mimicking what a human expert would actually do, our approach makes a diagnosis decision based on features learned in combination at multiple magnification levels. Our results show that the case-based approach achieves better performance than the state-of-the-art methods when evaluated on BreakHis, a histopathological image dataset for breast tumors.

1 Introduction

Histopathology is regarded as the gold standard method for cancer diagnosis, including almost all types of cancers, such as breast, lung, colon and prostate cancer [6, 9, 14]. Suspicious tissues are biopsied and the biopsy undergoes fixation, sectioning, and finally mounting on a slide. The biopsy section then is subjected to haematoxylin and eosin (H&E) staining which is a routinely used staining procedure that enhances tissue structure and cell morphology. A pathologist would then thoroughly examine the H&E stained slides under a microscope at multiple magnification levels, searching for morphological signatures indicating the onset or progression of cancerous tissues whose presence determines whether the tumor should be diagnosed as benign or malignant. The whole process, however, can be very time-consuming, since it is often required that the pathologist switch between magnification levels and jump among different image locations [14]. In addition, the diagnosis from a pathologist can sometimes be subjective and heavily dependent on the experience of the pathologist [9].

In order to address the above problems, computer aided diagnosis (CAD) systems have been proposed to facilitate cancer diagnosis, not only to reduce labor work for the pathologist, but also to improve objectivity and consistency. Despite the work that has been done

in the last few decades [6, 6, 9, 20], tumor malignancy classification remains still a challenge for most automatic cancer diagnosis applications due to the tremendous complexity of histopathological images, which can be due to various reasons including the staining variations in specimen treatment process [24] and the diversity of tissue characteristics in different cancers. Therefore, a robust and reliable CAD system for cancer diagnosis has to be designed to capture all discriminative features in histopathological images effectively. However, as has been pointed out by many researchers [6, 6, 9, 20], when using traditional classification approaches, the feature engineering step can be very difficult that requires a fair amount of expert domain knowledge.

Recently, a wide variety of new deep learning technologies [10, 22], such as the convolutional neural network (CNN), first developed by LeCun *et al.* [11], have achieved great success on various computer vision and pattern recognition tasks. Indeed, CNN has become the state-of-the-art method for image based classification problems, consistently outperforming traditional machine learning methods. More importantly, CNN can automatically extract discriminative features from images by itself. As a result, no hand-crafted feature engineering step is required anymore, which saves considerable efforts in most applications including histopathological image classification.

2 Previous work

Due to its superior performance compared to traditional machine learning methods, CNN has been widely applied to histopathological cancer diagnosis problems. Cireřan *et al.* [9] use CNN to detect mitosis in breast cancer histological images and won the ICPR 2012 mitosis detection competition. Sirinukunwattana *et al.* [18] propose a spatially constrained CNN for nucleus detection and then a Neighboring Ensemble Predictor (NEP) coupled with CNN for nucleus classification in colon cancer histological images, and achieve the highest average F1 score for this problem compared to other methods. Although both of the above two papers are not directly working on tumor malignancy classification, their results could undoubtedly benefit cancer diagnosis, since both mitosis and nuclear characteristics are important indicators for cancerous tissue detection. Direct work on malignancy classification have also been published. For example, Cruz-Roa *et al.* [9] show that a CNN classifier achieves a balanced accuracy of 84.23% for the detection of invasive ductal carcinoma, where the best performance of methods using handcrafted features and classifiers is 78.74%. Similarly, Litjens *et al.* [13] also demonstrate that CNN improves the efficacy of prostate cancer diagnosis.

We note that the previous work mentioned above on histopathological image classification using convolutional neural networks are done on whole slide images (WSI), and the patches used for training are extracted from the original images at a certain fixed magnification level. However, an experienced pathologist would not choose to determine a diagnosis decision based on a single magnification level. In practice, it is often required that the pathologists evaluate the histopathological slides at multiple magnification levels [15, 16], as different magnifications give different features. For instance, lower magnification gives global texture information and tissue structure while higher magnification resolve more on cellular morphology and sub-cellular details [9]. Sometimes it is difficult to determine a diagnosis merely based on a single magnification level. Only by integrating all the features at multiple magnification levels, a confident diagnosis can be determined.

Recently, an image dataset BreakHis is released [19], which provides histopathological images of breast tumor at multiple magnification levels (40 \times , 100 \times , 200 \times and 400 \times). Both

traditional methods using handcrafted features [19] and CNN method [20] have been applied on this dataset for malignancy classification, and it has been shown that by combining different CNNs using fusion rules, the CNN performance has an improvement of 6% in classification accuracy, compared to traditional methods. However, one disadvantage of this paper [20], is that four CNN classifiers have to be trained, with one classifier specialized for each of the four magnifications. Seeking to find a better solution to this problem, Bayramoglu *et al.* [10] propose a magnification independent approach with both single-task (malignancy) and multi-task (malignancy and magnification) classification, where they ignore magnification information of the image and train a unique CNN classifier for all magnifications. Although the performance is slightly impaired, it indeed improves the efficiency. Nevertheless, when evaluated on the testing sets, both of the previous work [10, 20] on BreaKHis dataset using CNN fail to determine a diagnosis for a patient based on features from multiple magnification levels at the same time. Instead, they give separate classification accuracy for each individual magnification, independent to other available magnifications.

3 Histopathological case-based classification

In order to build a more reasonable and reliable computer aided diagnosis system, we propose a case-based approach for histopathological malignancy classification, where a case is defined as a sequence of images including one or more images from each of all the available levels of magnification for a certain dataset. For example, for the BreaKHis dataset, a typical case could consist of one or more images at each of the following magnifications in order: 40 \times , 100 \times , 200 \times and 400 \times (Figure 1). A trained classifier should be able to learn all the features from different magnification images, and give a unique and more accurate result based on all information given (e.g. tissue structure at lower magnification, cell phenotype at higher magnification), equivalent to how an histopathological expert would choose to perform analysis at multiple magnification levels.

In this section, we first present our algorithm that constructs a case-based image set from any given histopathological image dataset with multiple magnifications and malignancies (Section 3.1). We then introduce a CNN model to classify our histopathological cases (Section 3.2). Finally, we describe the three performance metrics that will be used for the evaluation of histopathological case-based classification (Section 3.3).




Figure 1: A typical histopathological case of breast tumor with different magnifications.

3.1 Case-based image set initialization

Histopathological image datasets are often given as images in multiple separated magnifications, but not as cases. Therefore, the first step is to build an appropriate number of

histopathological cases based on the given dataset. To limit the size of the input set, the cases will include exactly one image from each magnification level. Algorithm 1 describes the initialization of a case-based image set from the original dataset with multiple magnifications and malignancies. Put simply, for each case build, the algorithm randomly chooses one image from each subset of images that belong to different magnifications, with the restriction that all images in the same case must have the same malignancy label, which will also be the final class label for the resulting case. For simplicity, we illustrate in Algorithm 1, assuming that two types of malignancy (benign and malignant) and four levels of magnification ($40\times$, $100\times$, $200\times$ and $400\times$) are available, which is the case for BreaKHis dataset. However, this algorithm can be applied to any number of malignancy types and magnification levels.

The only parameter passed to the algorithm is the expected number of output cases k (which we assume to be a multiple of the number of types of malignancy). In training set initialization, we want this set size, which we will denote as k_{train} , to be relatively large in order to avoid over-fitting our model later on, but also not too large due to limited computational resources and running time. Therefore, this parameter needs to be fine-tuned for different problem settings as we will show in more detail in Section 4.

Algorithm 1 can be applied to both training and testing sets, depending on the inputs of image sets. Note in the training phase, a case consists of one single image from each magnification level, but not necessary from the same particular patient. The images can be randomly selected from different patients as long as they share the same malignancy. This is why the patient information does not come in Algorithm 1. However, in the testing phase, we may want the cases to be patient specific, which can be achieved by setting patient specific images as the input to Algorithm 1. After the whole process, the initialized case-based image sets are ready for training or evaluation.

3.2 ResNet-based classifier

We choose to use deep residual neural networks (ResNets) to classify the histopathological cases. ResNets are a special kind of convolutional neural networks that have residual units in parallel to regular convolutional layers. The design of residual units are quite flexible such that they can also be further engineered in order to get better performance [17, 8]. We start with a simple 18-layer ResNet model (ResNet-18), as this model can be easily adapted to even deeper models (e.g. 152 layers) if required. The overall architecture of ResNet-18 is shown in Figure 2.

The model contains two types of residual units: the residual unit with an identity shortcut and the residual unit with a projection shortcut. The only difference between these two types is that in the projection shortcut, an additional convolutional layer is required due to the change of dimension from input to output. Each residual unit contains six sequential components: Batch Normalization, Rectified Linear Unit (ReLU), Convolution, Batch Normalization, ReLU and Convolution. An average pooling layer is used before the final fully connected layer.

3.3 Metrics

Spanhol *et al.* [20] have introduced two ways to report method performances for medical image classification: image recognition rate and patient recognition rate. Here, to accommodate for our case-based approach, however, we use a case level metric instead of an image

Algorithm 1: Case-based image set initialization

Input : image sets $I_{Malignancy \times Magnification}$, where *Malignancy* is the set of malignancy types, e.g. {benign, malignant}, and *Magnification* is the set of magnifications, e.g. {40, 100, 200, 400}

Output : $X \leftarrow$ data, $y \leftarrow$ label

Parameters: k = expected number of output cases

- 1 Initialize $i = 0$;
- 2 **foreach** $mal \in Malignancy$ **do**
- 3 Initialize $counter_{mal} = 0$;
- 4 Initialize new current combination X_i ;
- 5 **repeat**
- 6 **foreach** $mag \in Magnification$ **do**
- 7 randomly pick an image from image set $I_{(mal, mag)}$ and add to X_i ;
- 8 **end**
- 9 **if** current combination X_i not in X **then**
- 10 add X_i to X ;
- 11 $y_i = mal$;
- 12 $i += 1$;
- 13 $counter_{mal} += 1$;
- 14 **end**
- 15 **until** $counter_{mal} \geq k/|Malignancy|$;
- 16 **end**

level metric. The case recognition rate is defined as follows:

$$\text{Case Recognition Rate} = \frac{N_{rec}}{N_{all}} \quad (1)$$

where N_{all} is the total number of all cases constructed for the testing set, and N_{rec} is the number of correctly classified cases.

Unlike the case recognition rate, the patient recognition rate takes patient information into account. For each patient p in the testing set, let $N_{p,all}$ be the total number of cases that belong to patient p , and $N_{p,rec}$ be the number of correctly classified cases for patient p , then the patient recognition rate can be defined as [20]:

$$\text{Patient Recognition Rate} = \frac{\sum_p (N_{p,rec} / N_{p,all})}{\text{Total Number of Patients}} \quad (2)$$

In addition to the above two recognition rates, we also give a new metric defined at the diagnosis level. First, we give a final diagnosis to each patient in the testing set based on a simple voting strategy where we assume that the diagnosis is benign if the ratio of benign to malignant cases for the patient p is above a threshold, *malignancy_threshold*:

$$\text{Patient Diagnosis } p = \begin{cases} \text{benign}, & \text{if } \frac{N_{p,benign}}{N_{p,all}} > \text{malignancy_threshold} \\ \text{malignant}, & \text{otherwise} \end{cases} \quad (3)$$

where $N_{p,benign}$ is the number of cases that are diagnosed as benign for patient p . For example, if *malignancy_threshold* is set to 0.5, the patient p is assigned a diagnosis of benign if more

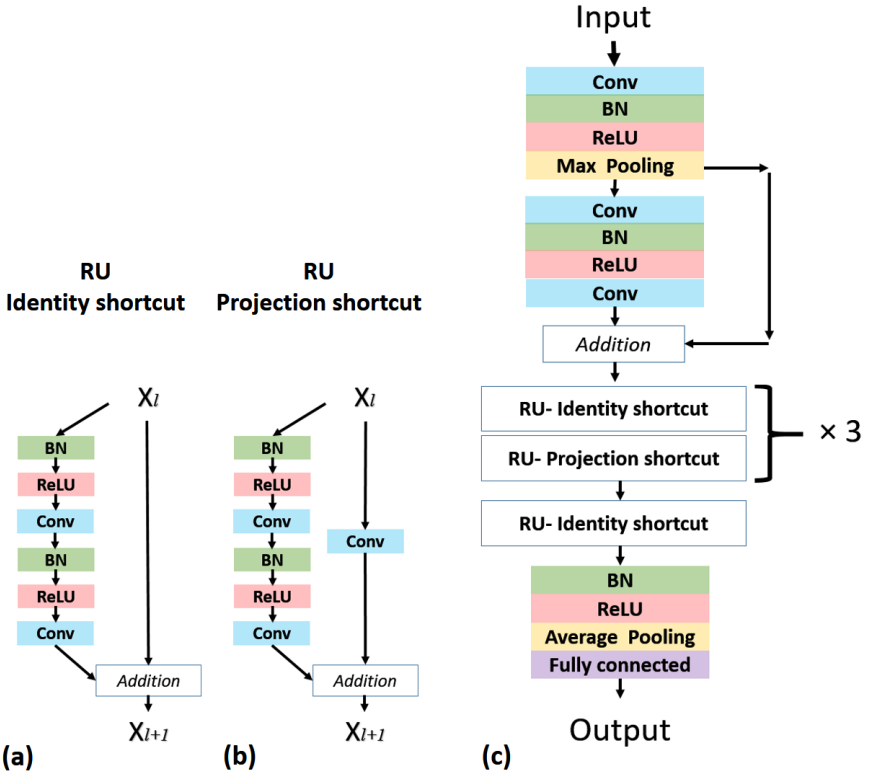


Figure 2: (a) Residual unit with identity shortcut; (b) Residual unit with projection shortcut; (c) Overall architecture of the ResNet used in the paper.

than half of the cases for patient p are classified as benign. Based on the diagnoses assigned to the patients, diagnosis accuracy for the classification is defined as the follows:

$$\text{Diagnosis Accuracy} = \frac{\text{Number of Correctly Diagnosed Patients}}{\text{Total Number of Patients}}. \quad (4)$$

We believe the diagnosis accuracy metric should be emphasized more for future research on histopathological diagnosis problems, as it is of utmost clinical importance that a computer-aided diagnosis system be able to give a final diagnosis for a patient, and based on the accuracy at which the diagnosis is correct or not, we can judge its performance.

4 Experiments and Results

This section evaluates our case-based approach for histopathological diagnosis that is proposed in Section 3.

4.1 Dataset

To test the proposed case-based approach for histopathological diagnosis, we use the BreakHis database [19], a recently released dataset of breast tumor histopathological images. BreakHis

contains both benign and malignant breast tumor images, which were collected from 82 patients at multiple magnification levels ($40\times$, $100\times$, $200\times$ and $400\times$). Each patient may have a different number of images for each magnification. In total, there are 2480 benign and 5429 malignant images, with each image acquired in three channels (RGB).

Besides the histopathological images, BreaKHis also provides a five-fold protocol for testing. We use the same testing protocol as previous work [14, 24], where the whole dataset is split into training (70%) and testing (30%) set for five trials, such that none of the images associated with the patients in the training set are used in the testing set. In the end, 54 out of the total of 82 patients are grouped into the training set, and the rest of the 28 patients are used as evaluation samples for all the five folds.

The BreaKHis images are originally of size $700\times 460\times 3$. To speed up the processing times and lessen the memory requirements, the images are resized to $100\times 100\times 3$ for both the training and testing sets.

4.2 Implementation

With all images from BreaKHis, we implement Algorithm 1 to build histopathological cases for both the training and testing sets. To find the best parameter k_{train} for Algorithm 1 when initializing the training sets, we utilize fold 1 for a series of experiments by setting the number of output cases over a range of values from 100 to 40,000 as shown in Figure 3. After comparing the case-level accuracies for the different sizes of training sets, we choose the smallest size of the training set that gives the best accuracy as our final k_{train} . Note that for some of the smaller sizes of the training sets, we repeat some of the experiments independently three or five times since the performance of trained model can vary a lot for these sizes. The final chosen parameter k_{train} for training set initialization achieves a balance between computational resource requirement and performance. On the other hand, for testing set initialization, we simply use $k_{test} = 30,000$ for the size of the testing sets as evaluation on thirty thousands cases gives a quite stable estimation of model performance according to our trial experiments.

For all experiments, we implement our classifier ResNet using Keras, a deep learning library written in python with either TensorFlow or Theano as a backend [2]. We use Theano as the backend in this paper. To optimize the weights, we use stochastic gradient descent, with a batch size of 100 to compute the gradients using back propagation. The initial learning rate is set to 0.001, decay by $1e-6$ over each update, and Nestrov momentum is set to 0.9. We train our neural network for 100 epoches. All experiments are done on 4 Intel Xeon(R) E3-1271 v3 processors with a NVIDIA Quadro K2000/PCIe/SSE2 GPU with CUDA 7.5 installed in a Ubuntu 16.04 LTS.

4.3 Results

First, regarding the choice of k_{train} , as is shown in Figure 3, with the increase of the total number of cases used for training, the testing accuracy also increases and finally reaches the plateau. When the model is trained on only 100 cases, there is a large variation in model performance based on five independently conducted experiments. In the worst case, for 100 cases, the performance is not much better than a random guess. However, the model performance is significantly improved when trained on a large number of cases. In addition, the variation becomes smaller as well. From the bottom plot in Figure 3, we can see that the curve starts to converge when the number of cases is increased to 5,000, and reaches a

maximum at around 10,000. To understand the effect of the choice of k_{train} on the running time, when setting k_{train} equal to 40,000, the total training time required for a single fold is around 900 seconds per epoch, while it is 2 seconds per epoch for k_{train} set to 100. By setting our parameter k_{train} equal to 10,000, we significantly reduce the running time by around four times, from 900 seconds to 226 seconds, when compared to k_{train} equal to 40,000, without sacrificing accuracy. Therefore, we choose to set our parameter k_{train} equal to 10,000 in training set initialization algorithm for all the following experiments.

Experiments	No. of cases									
	100	1k	5k	10k	15k	20k	25k	30k	35k	40k
Exp. 1	0.8565	0.9156	0.9243	0.9317	0.9199	0.9124	0.9192	0.9276	0.9208	0.9294
Exp. 2	0.8133	0.9028	0.9382	0.9316						
Exp. 3	0.7758	0.9291	0.9222	0.9106						
Exp. 4	0.5008	0.9119								
Exp. 5	0.8501	0.9325								




Figure 3: Performance in terms of case-level accuracy versus number of histopathological cases k_{train} used for training. Top table shows the testing accuracies in each experiment; The bottom plot is the visualization of the table.

With the parameters k_{train} and k_{test} set, we can then thoroughly evaluate our case-based approach based on five-fold testing protocol, using the metrics that we described in Section 3.3. For each fold, we first build the case-based training and testing sets using Algorithm 1, by setting $k_{train} = 10,000$ for the training set and $k_{test} = 30,000$ for the testing set. Note that both the training and testing sets are balanced in terms of the different malignancy types. After the models are trained, we then evaluate the model performance using the following three metrics: case recognition rate, patient recognition rate, and diagnosis accuracy. For the diagnosis of benign or malignant, we set *malignancy_threshold* to 0.5. Table 1 shows the final results.

Our case-based approach gives average accuracies of 91.48% (case-level), 86.36% (patient-level) and 88.57% (diagnosis-level) on the testing sets. As we are the first to use case-level and diagnosis-level accuracies, we can't compare the results for these metrics to previous results. However, based on patient-level accuracy, our case-based approach (86.36%) outperforms the multi-task CNN method (82.13%, average of four magnifications) [14] and the magnification independent single-task CNN method (83.25%, average of four magnifications) [15], and achieves a comparable performance to the best results obtained from the

Table 1: Performance of case-based approach for histopathological malignancy diagnosis based on case-level, patient-level and diagnosis-level accuracy.

Accuracy Type	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Average
Case Recogn. Rate	0.9246	0.8596	0.9355	0.9220	0.9323	0.9148
Patient Recogn. Rate	0.8731	0.8424	0.8753	0.8090	0.9182	0.8636
Diagnosis Accuracy	25/28	23/28	26/28	23/28	27/28	0.8857

combination of four patch image extraction strategies and three fusion rules using a patch-based method for specific magnifications ($40\times$: 90.0%; $100\times$: 88.4%; $200\times$: 84.6%; $400\times$: 86.1%) [20].

We further investigate the misclassified patients in terms of malignancy diagnosis for all five folds, and summarize the results as confusion matrices in Figure 4. In total, 16 out of 140 patient samples over the five folds are misclassified, with a false positive rate of 5.0% and a false negative rate of 6.43%.




Figure 4: The confusion matrices of case-based approach for histopathological malignancy diagnosis in five folds.

5 Conclusion

In this paper, we propose a case-based approach for histopathological malignancy diagnosis using deep residual neural networks. We first introduce an algorithm for case-based image set initialization for both training and testing based on histopathological images at multiple magnification levels, and then present a ResNet-based classifier and three metrics to report method performances for medical image classification. Finally, we evaluate our proposed approach using the breast tumor histopathological image dataset BreKHis. Our results show that the case-based approach achieves better performance than the state-of-the-art methods. Moreover, we believe our case-based approach is a more reasonable way for histopathological malignancy classification since it makes diagnosis decision based on features learned at multiple magnifications. Another principle advantage of our work over the previous work [10, 20] is that our method gives a single diagnosis for the patient, whereas in the previous work four potentially differing diagnoses are given for the same patient, one for each of four magnification levels. To be clinically applicable, these latter approaches would then require a final voting step or similar diagnosis selection step which are not discussed in their papers [10, 20]. For future work, more complex deep CNN architectures will be investigated.

References

- [1] Neslihan Bayramoglu, Juho Kannala, and Janne Heikkilä. Deep learning for magnification independent breast cancer histopathology image classification. In *23th International Conference on Pattern Recognition, ICPR*, 2016.
- [2] François Chollet. Keras, 2015. <https://github.com/fchollet/keras>.
- [3] Dan C Cireşan, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber. Mitosis detection in breast cancer histology images with deep neural networks. In *International Conference on Medical Image Computing and Computer-assisted Intervention*, pages 411–418. Springer, 2013.
- [4] Angel Cruz-Roa, Ajay Basavanahally, Fabio González, Hannah Gilmore, Michael Feldman, Shridar Ganesan, Natalie Shih, John Tomaszewski, and Anant Madabhushi. Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks. In *SPIE medical imaging*, pages 904103–904103. International Society for Optics and Photonics, 2014.
- [5] Cigdem Demir and Bülent Yener. Automated cancer diagnosis based on histopathological images: a systematic survey. *Rensselaer Polytechnic Institute, Tech. Rep*, 2005.
- [6] Metin N Gurcan, Laura E Boucheron, Ali Can, Anant Madabhushi, Nasir M Rajpoot, and Bulent Yener. Histopathological image analysis: A review. *IEEE reviews in biomedical engineering*, 2:147–171, 2009.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016.
- [9] Lei He, L Rodney Long, Sameer Antani, and George R Thoma. Histology image analysis for carcinoma detection and grading. *Computer methods and programs in biomedicine*, 107(3):538–556, 2012.
- [10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [11] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- [12] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [13] Geert Litjens, Clara I Sánchez, Nadya Timofeeva, Meyke Hermsen, Iris Nagtegaal, Iringo Kovacs, Christina Hulsbergen-Van De Kaa, Peter Bult, Bram Van Ginneken, and Jeroen Van Der Laak. Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis. *Scientific reports*, 6, 2016.

- [14] Michael T McCann, John A Ozolek, Carlos A Castro, Bahram Parvin, and Jelena Kovacevic. Automated histology analysis: Opportunities for signal processing. *IEEE Signal Processing Magazine*, 32(1):78–87, 2015.
- [15] Lucia Roa-Peña, Francisco Gómez, and Eduardo Romero. An experimental study of pathologist’s navigation patterns in virtual microscopy. *Diagnostic pathology*, 5(1):71, 2010.
- [16] David Romo, Juan D García-Arteaga, Pablo Arbeláez, and Eduardo Romero. A discriminant multi-scale histopathology descriptor using dictionary learning. In *SPIE Medical Imaging*, pages 90410Q–90410Q. International Society for Optics and Photonics, 2014.
- [17] Raphael Rubin, David S Strayer, Emanuel Rubin, et al. *Rubin’s pathology: clinico-pathologic foundations of medicine*. Lippincott Williams & Wilkins, 2008.
- [18] Korsuk Sirinukunwattana, Shan E Ahmed Raza, Yee-Wah Tsang, David RJ Snead, Ian A Cree, and Nasir M Rajpoot. Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. *IEEE transactions on medical imaging*, 35(5):1196–1206, 2016.
- [19] Fabio A Spanhol, Luiz S Oliveira, Caroline Petitjean, and Laurent Heutte. A dataset for breast cancer histopathological image classification. *IEEE Transactions on Biomedical Engineering*, 63(7):1455–1462, 2016.
- [20] Fabio Alexandre Spanhol, Luiz S Oliveira, Caroline Petitjean, and Laurent Heutte. Breast cancer histopathological image classification using convolutional neural networks. In *2016 International Joint Conference on Neural Networks (IJCNN)*, pages 2560–2567. IEEE, 2016.
- [21] Mitko Veta, Josien PW Pluim, Paul J Van Diest, and Max A Viergever. Breast cancer histopathology image analysis: A review. *IEEE Transactions on Biomedical Engineering*, 61(5):1400–1411, 2014.