
Whole Slide Image Classification of Gastric Cancer using Convolutional Neural Networks

Junni Shou¹, Yan Li¹, Guanzhen Yu^{2,*}, and Guannan Li^{1,*}

¹Awakens Intelligence Technology Co., Ltd, China

²Department of Oncology, Longhua Hospital Affiliated to Shanghai University of Traditional Chinese Medicine, Shanghai, China

*qiaoshanqian@aliyun.com, *guannan.li@awkint.com

Abstract

1 Gastric cancer is one of the main causes of cancer and cancer-related mortality
2 worldwide, and the diagnosis based on histopathology images is a gold standard
3 for gastric cancer detection. However, manual diagnosis is labor-intensive and
4 low in inter-observer agreement. Computer-aided image analysis method were
5 thus developed to alleviate the workload of pathologists and overcome the prob-
6 lem of subjectivity. Histopathology image analysis using deep learning has been
7 proved to give more promising results than traditional methods on many whole
8 slide image cancer detection tasks, including breast cancer detection and prostate
9 cancer detection. In this paper, we further studied a whole slide image classifi-
10 cation method using Convolutional Neural Networks (CNNs) on gastric cancer
11 data. The method classify a whole slide image based on patch-sized classification
12 results. Various experiments for patch-level classification using different existing
13 CNN architectures were conducted. Experiment results show that the architecture
14 gives the state-of-the-art result in natural image classification tasks can also give
15 impressive results in histopathology image classification tasks.

16 1 Introduction

17 Gastric cancer is the second most common cancer in China[3] and the third leading cause of
18 cancer death worldwide[16]. Diagnoses in histopathology images is essential for assessing the
19 tumor response and prognosis of patients to different treatments[11, 2, 5, 20]. Nevertheless, the
20 manual pathological diagnoses are time-consuming, often require tedious and laborious work. Also,
21 manual diagnoses could be subjective and difficult to standardize, leading low level diagnostic
22 concordance[15, 4]. Therefore, computer-aided histopathology image analysis methods are developed
23 to assist pathologists to improve the efficiency, accuracy and consistency of cancer detection[6, 7].

24 Recent works show great success in applying deep learning for histopathology image analysis. Specif-
25 ically, Convolutional Neural Networks (CNNs) are applied to analyze the complicated histopathology
26 images. This technique allows an image analysis method to be designed without specific field-related
27 knowledge, and the model would learn all the features from images itself. Spanhol et al.[19] used
28 a simple CNN architecture, AlexNet[10], to classify hematoxylin and eosin (H&E) stained breast
29 histopathology images into two classes, benign and malignant. Small scaled input patch were used in
30 their work. They then combined the results from patch classification to give the local-region-level
31 classification. Subsequently, Araújo et al.[1] extended the classification problem from 2-class to
32 4-class, and also experimented larger scaled input patches. While these works focused their studies in
33 patch-level and local-region-level classifications, Litjens et al.[12], Wang et al.[21] and Liu et al.[13]
34 further improved the image analysis methods, giving a whole-slide-level classification prediction.

Table 1: Details of annotations given for the gastric cancer datasets

	Training/Validation	Testing
Slide-level labels (WSIs)	150 cancer + 39 normal	110 cancer + 70 normal
Pixel-level annotations	1500 region images (from 150 cancer WSIs, 10 from each slide)	5 cancer WSIs

35 Authors in works mentioned above proposed their classification methods for either breast cancer
 36 data or prostate cancer data, and Sharma et al.[17] later applied deep learning methods to the gastric
 37 cancer data. They proposed an introductory CNN architecture and compared the performance of it
 38 with AlexNet[10] and several other traditional methods. 15 whole slide images (WSIs) were used for
 39 extracting patches for training, validation and testing (11 for cancer classification and 4 for necrosis
 40 detection), and accuracies of 69.90% for cancer classification and 81.44% for necrosis detection were
 41 achieved for patch-level classifications.

42 In our work, with a larger gastric cancer dataset introduced, we evaluated the feasibility of a whole-
 43 slide-image-level classification method for gastric cancer. Additionally, to see would the architecture
 44 with more complicated structures outperform AlexNet[10] for histopathology images, different
 45 existing CNN architectures were assessed in patch-level classifications. The effect of different
 46 patch scales were also experimented. Finally, we achieved an accuracy of 98.698% for patch-level
 47 classification and an accuracy of 97.728% for slide-level classification.

48 2 Dataset

49 The gastric cancer dataset consists of 369 WSIs, each from a distinct patient who underwent curative
 50 surgery at Changhai Hospital in Shanghai, China, from 2001 to 2005. Mean age of these patients
 51 was 59 years old. The slides in the dataset were stained with hematoxylin and eosin (H&E), and
 52 digitized by MAGSCANNER KF-PRO-120¹ at magnification of 20 \times . The use of these slides has
 53 been approved by the Changhai Hospital Institutional Review Board.

54 Annotations of the data are given by expert pathologists, and presented in two different forms, pixel-
 55 level delineation of cancerous regions on images and cancer/normal labels for each slide. 1500 cancer
 56 region images (acquired from 150 WSIs, each with 10 region images), each of size 2048 \times 2048
 57 pixels, and another 5 cancer WSIs are given with pixel-level annotations (the 5 WSIs are exhaustively
 58 annotated). The 1500 cancer region images were used for extracting positive patches used for training
 59 and validating the patch-level classifier, and the 5 WSIs were used for positive patch extraction for
 60 testing the trained classifier. The patch extraction strategies would be further explained in details
 61 in the following sections. Total number of 369 WSIs are given with slide-level labels, and are split
 62 into 189 (150 cancer slides and 39 normal slides) for training of the slide-level classifier and 180
 63 (110 cancer slides and 70 normal ones) for testing. The normal slides were also used for the negative
 64 patch extraction. The first 39 normal slides were used for extracting negative patches for training and
 65 validation of the patch-level classifier, whereas 5 out of the 70 normal slides were used for testing.
 66 Details of the annotations are summarized in Table 1.

67 3 Methods

68 The classification method consists of four steps: (1) image preprocessing to extract the tissue region;
 69 (2) patch-level classification using CNN; (3) cancer likelihood map generated from the patch-level
 70 classification results; (4) slide-level classification based on the likelihood map. Details are explained
 71 in the following sections.

¹<http://www.kfbio.cn/productshow.php?cid=27&id=43>

72 **3.1 Image Preprocessing**

73 Most of the WSI area is non-informative background. These area would lead to unnecessary
74 computational costs. To save computational costs and increase efficiency, we did image preprocessing
75 to extract the tissue regions from the slide first.

76 Common thresholding algorithms were used in [12, 21] to extract the tissue region. These thresholded-
77 ing algorithms differentiate foreground and background objects by setting a threshold intensity,
78 and simply grouping pixels with intensity higher than the threshold and lower than the threshold
79 separately[14]. These methods are capable of separating the foreground objects from the blank
80 background regions, however, it is unable to remove regions of glasses, glues and dirt, which would
81 have pixels with similar intensities to tissue regions but different from blank regions. These useless
82 regions would remain together with the tissue regions, causing computational costs and unnecessarily
83 complicating the cancer detection problem since patches containing different forms of glass textures
84 would also be required as normal patches for the training of the patch-wise classifier. Therefore, a
85 tissue extraction method based on differences between R/G/B color channels was used. The blank
86 background is close to white, whereas the regions of glasses, glues and dirt are generally greyish
87 or close to black. Pixels with color close to black or white, or greyish colors would have relatively
88 uniform values for R/G/B channels. In other words, the difference between the highest channel value
89 and the lowest channel value of the pixel with those colors would be smaller than a certain threshold.
90 Thus, pixels with channel value difference greater than the threshold would be marked as tissue
91 regions, while the remaining would be regarded as the non-informative background. A threshold
92 value of 25 was empirically obtained and was used to get the binary mask. Noises and small holes
93 were later removed by morphological operations. Figure 1 shows the results of applying the tissue
94 extraction method.

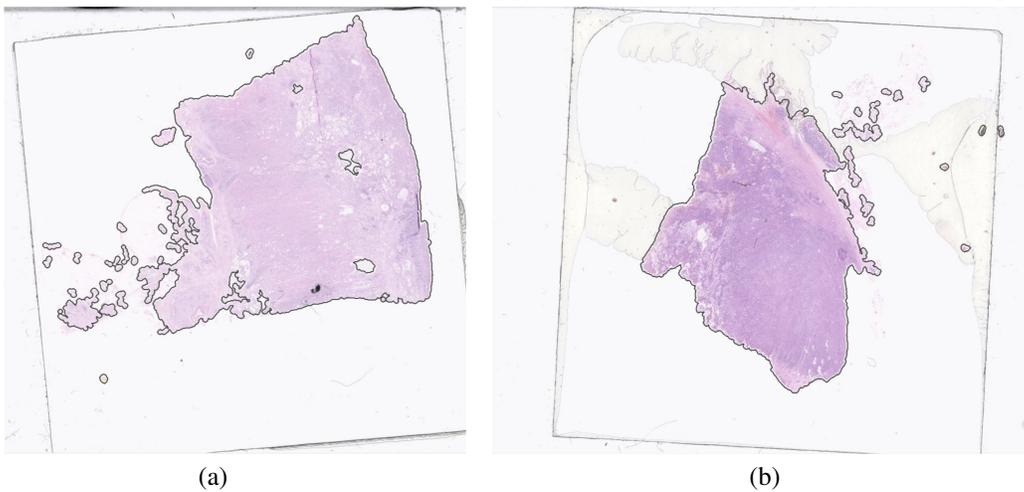


Figure 1: Examples show tissue regions extracted from WSI. The extracted tissue regions are successfully separated from background consists of blank regions, marks of glasses and glues, and dirt, and contoured with black curves.

95 **3.2 Patch-wise Classification**

96 The WSIs are large in size, hence it is impossible to directly put them as the input for the classification.
97 One good way is to divide the tissue regions into small patches and the further slide-level classification
98 could be done by combining all the results of the small patches. Since the objective is to give diagnostic
99 results of cancer/normal for the slide, it is not important to precisely delineate the boundaries of
100 cancerous regions on the slide. As a result, image classification models were considered rather than
101 segmentation ones. For achieving better performance, CNN was used for patch-wise classifications.
102 In the following, patch extraction strategies, data augmentation and detailed explanation of network
103 architectures we used for training the patch-level classifier are presented.

104 3.2.1 Patch extraction

105 Patches for training, validation and testing sets were generated according to the pixel-level annotations
106 given by pathologists. On images with pixel-level annotations, patches were extracted with a stride
107 of 100 pixels in the tissue region and labeled as positive if the patch center located in the annotated
108 cancerous region. Normal patches were generated randomly in the tissue regions of normal slides,
109 and labeled as negative. To avoid bias to the patch dataset, the ratio of total amount of negative
110 patches to total amount of positive patches was controlled to be roughly 1:1. Three patch sizes were
111 extracted for further comparison: 120×120 , 240×240 and 480×480 .

112 3.2.2 Data augmentation

113 Data augmentation was utilized to obtain more robust models. To increase the size of the patch
114 dataset, random cropping of sizes 112×112 , 224×224 and 448×448 was applied to input patches
115 of 120×120 , 240×240 and 480×480 respectively during the training. Since features extracted
116 from histopathology images should be orientation invariant, random flipping and rotation were also
117 used. Vertical flipping and horizontal flipping would be applied to the input patches randomly with a
118 probability of 0.5, and the patches would then be rotated by random multiples of 90° . To combat the
119 variations between different slides caused by, for example, different color staining, problem caused
120 by overexposure during scanning, etc., the brightness, contrast, saturation and hue of the patches
121 were slightly adjusted by a random factor in each training epoch.

122 3.2.3 Using existing network architectures

123 We experimented with various previously existing CNN architectures to find the network architecture
124 that suits the gastric cancer classification problem best. We started the evaluation with AlexNet[10],
125 a network architecture simply composed of layers of convolution and pooling sequentially, followed
126 by fully-connected layers. Next, VGG-16[18] was evaluated. VGG-16 has very similar "plain"
127 network structures to AlexNet, but with more layers. We then experimented on more complex
128 models, ResNets[8] and DenseNets[9]. These models have much deeper networks. ResNets[8] use
129 identity-based shortcut connections to bypass the signal from previous layers to the next, alleviating
130 degradation problems during the training for very deep networks. DenseNets[9] provide with a
131 reformulation of the connection, which helps to train a deeper network but also substantially improves
132 the parameter efficiency and better the generality of the trained model.

133 3.3 Cancer Likelihood Map

134 Tissue regions were extracted from the slide first, and patch-level inference was then carried out in a
135 sliding window manner with strides of 28, 56, 112 and 224 pixels in the tissue regions. Smaller strides
136 would lead to finer results but with more computation time. We did experiments to compare the
137 results using different strides, and the experiment suggests that the results with stride of 56 was good
138 enough to give visually smooth likelihood map and deliver enough information for later slide-level
139 classification, while not being too time-consuming. Therefore, stride of 56 pixels was used for later
140 experiments in this paper.

141 Classification results for small patches were then merged into the cancer likelihood map. Pixels in
142 the patches predicted to be positive would be added by one on the map while the ones predicted
143 to be negative would be remained as the original value. The final cancer likelihood map were then
144 normalized to values in range [0,1] by dividing by

$$factor = \left(\frac{patch\ size}{stride} \right)^2 \quad (1)$$

145 3.4 Slide-level Classification Based on Cancer Likelihood Maps

146 After getting the cancer likelihood map for the slide, N_t binary masks according to N_t different
147 thresholds of likelihood were obtained from the map, where N_t is the number of thresholds used. For
148 each thresholded binary mask, we collect 9 features, including area, solidity, eccentricity and extent
149 of the largest component and the second largest component of the cancer area, and ratio between the
150 total cancer area and the tissue area. Thus each slide could get $N_t \times 9$ features.

Table 2: Details of datasets for training the patch-level classifier

	Training	Validation	Testing
positive/negative patches	135k/135k	34k/34k	50k/50k
Total	270k	68k	100k

151 189 slides (150 cancer slides and 39 normal slides) provided with slide-level labels were used for
 152 training the slide-level classifier. Cancer likelihood maps were first generated for these slides. With
 153 the features extracted from the map as the inputs and the slide-level label as the output label, a random
 154 forest classifier was trained to determine whether the slide should be predicted as cancer or normal.
 155 The performance of the slide-level classifier was further tested by the testing set consisting of 110
 156 cancer slides and 70 normal slides.

157 4 Experiments

158 4.1 Patch-level Classification

159 1500 annotated 2048×2048 region images were used to generate a total amount of 169,275 positive
 160 patches for training and validation. Negative patches for training and validation were then extracted
 161 from 39 normal slides. In order to keep the ratio of the total amount of positives to negatives to be
 162 roughly 1:1, each normal slide was used for generating 4340 negative patches randomly positioned
 163 in the tissue regions of the normal slide. Accordingly, total amount of 338k positive and negative
 164 patches were generated. Then, these patches are equally divided into 5 groups. One of the groups
 165 was used as validation set whereas the remaining was used as the training set, leading to a training set
 166 of 270k patches and a validation set of 68k patches. Then, 5 annotated cancer slides and 5 normal
 167 slides from the additional 180 slides were used to produce the testing set. For each slide, 10k patches
 168 were extracted, making a testing set of 100k patches in total, consisting of 50k positives and 50k
 169 negatives. Details for each dataset are summarized in Table 2.

170 Because of the tremendous differences between histopathology images and natural images, we did not
 171 use any pretrained models and all networks were trained from scratch. The network was trained for
 172 20 epochs, and the one with the highest accuracy on validation set was saved. Then the best model
 173 from last 20-epoch-training would be trained for another 20 epochs with a learning rate ten times
 174 smaller than before. Repeating for three times and the final network was obtained. All networks for
 175 comparison were acquired in the same way.

176 Table 3 shows the accuracy of different network architectures. Results in Table 3 indicate that the
 177 accuracy of a testing set is always lower than that of a validation set. This is because that the validation
 178 set was draw out from the same dataset where the training set was from, which means these two
 179 sets have though slightly different but similar patches. As for the testing set, patches were extracted
 180 from another 10 slides that had never been seen in the training process. Hence, patches in the testing
 181 set should be more different from the training patches, and that difference leads to the decrease in
 182 the accuracy. However, although there is a slight drop, the accuracy on the testing set is still very
 183 high. This may be due to the large amount of patches we used for training, and proper extensive data
 184 augmentation encourages the generality of the model and avoids over-fitting problems.

185 We first evaluated the performance of different network architectures with the same patch scale,
 186 224×224 . Although VGG-16[18] has a very simple and straightforward architecture, it still achieved
 187 surprisingly good result. This may be due to large amount of parameters in the VGG-16 architecture.
 188 Still, the highest accuracy was achieved by DenseNet-201[9]. DenseNet-201 has much fewer
 189 parameters than VGG-16, but its structure utilize features in an efficient way to avoid feature
 190 redundancy and help to generate a more compact network delivering better results. Regarding the
 191 fact that DenseNet-201 gave the best results for both validation and testing sets, we performed the
 192 following experiment using DenseNet-201.

193 Next, we compared the performance of DenseNet-201[9] using different input scales. Input size of
 194 112×112 pixels gave much lower accuracy as expected, whereas model with 448×448 sized inputs
 195 gave slightly better results than 224×224 . Since the improvement of model using 448×448 sized
 196 inputs was not very significant, about 0.1% in accuracy on training set and 0.2% on testing set, and it

Table 3: Patch-wise classification accuracy (%)

Network	Input patch size	Validation	Testing
AlexNet	224×224	98.156	96.722
VGG-16	224×224	99.565	98.413
ResNet-101	224×224	98.879	98.353
ResNet-152	224×224	99.290	97.244
DenseNet-121	224×224	98.444	98.153
DenseNet-201	112×112	97.845	96.821
	224×224	99.655	98.698
	448×448	99.758	98.973

Table 4: Slide-level classification accuracy (%)

Random forest classifier	Training	Testing
accuracy	100.000	97.728
sensitivity	100.000	95.454
specificity	100.000	100.000

197 would cause substantial increase in computational costs. Considering the time constraint, we chose
 198 DenseNet-201 with 224×224 as the input size to finish the following experiments.

199 4.2 Slide-level Results

200 Once the trained network is obtained, it can be applied to the tissue region of the slide in a sliding
 201 window manner. The cancer likelihood map can be generated afterwards. An example^{2,3} of cancer
 202 likelihood map for gastric cancer detection is shown in Figure 2(c). Figure 2(b) presents the
 203 corresponding ground truth annotation given by pathologists. Regions predicted with high likelihood
 204 of being cancerous are shown in red or yellow, whereas regions with low likelihood of cancer are
 205 shown in green or blue. Transparent areas indicate normal tissue regions. Most of the cancerous
 206 regions are correctly detected. Few false positives exist.

207 After getting the likelihood map, features were extracted from the map and fed as the training inputs
 208 to a Random Forest classifier. The slide-level label was used as the training ground truth. Dataset of
 209 189 labeled slides (comprised of 150 cancer slides and 39 normal slides) was used as the training
 210 set and additional 180 slides (comprised of 110 cancer slides and 70 normal ones) were used as
 211 the testing set. We evaluated the performance of slide-level classifier with accuracy, sensitivity and
 212 specificity. Results for the slide-level classifier are summarized in Table 4. It can be seen from Table 4
 213 that the classifier classified all normal slides in the testing set correctly but mis-classified several
 214 cancer slides. Most of the mis-classified cancer slides contains very few amount of cancerous area,
 215 like 0.23% of the tissue area. In the training set, the cancer slides contain an average cancerous region
 216 of 8.77% of the tissue regions (the least amount is 1.32% of the tissue region). Hence, even though
 217 the patch-level classifier is able to detect the cancerous regions in the slide, the post-processing
 218 slide-level classifier was trained to "assume" those detected cancerous regions to be false positives
 219 and gave incorrect slide-level classification results. In spite of this, the slide-level classifier is still
 220 able to give 100% of accuracy for gastric cancer detection if the slide contains cancerous area more
 221 than 1.5% of the tissue region.

²Full sized example image and the corresponding cancer likelihood map can be viewed by link:
<http://box.histogram.cn/s/i7Aune>. The link is generated by HISTOGRAMTM for data sharing. Likelihood
 map can be viewed by clicking the "heatmap" icon, and the pixel-level annotation can be viewed by clicking the
 eye icon.

³More full sized examples and corresponding maps can be viewed here: <http://box.histogram.cn/s/OP0eNa>

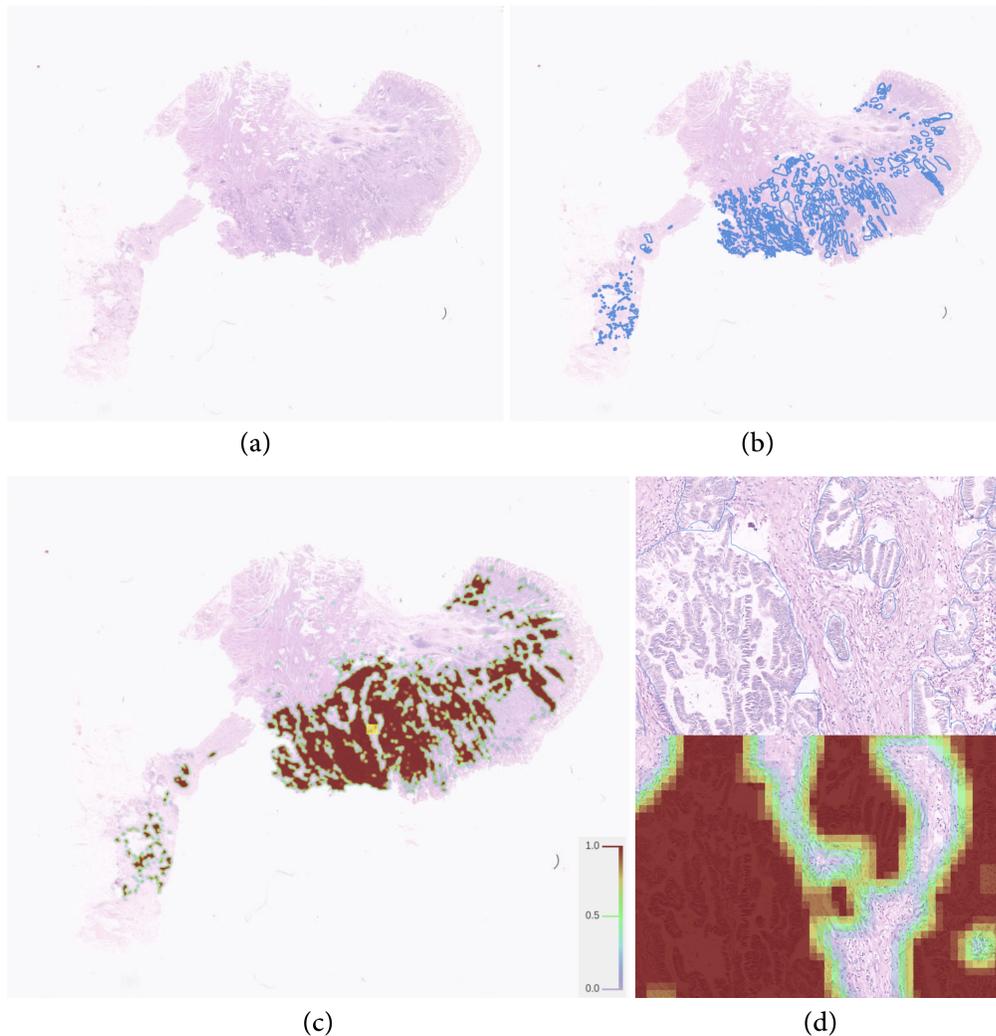


Figure 2: (a) Whole slide image of a tissue sample of gastric cancer. (b) The ground truth annotation of cancerous regions given by expert pathologists. (c) The predicted cancer likelihood map of the slide. (d) Zoom in on area indicated by the yellow square on (c) to a higher magnification ($10\times$). The top image is the annotation of the cancerous region, and the bottom image is the cancer likelihood map.

222 5 Conclusions

223 In this paper, the feasibility of a whole slide image classification method of gastric cancer using
 224 CNN is studied. An image preprocessing method is introduced to extract tissue regions from non-
 225 informative background, including blank regions, marks of glassess, glues and dirt. The whole slide
 226 classification is acquired by combining patch-level classification results. Patch extraction strategies
 227 are shown, and data augmentation is applied to increase the size of the training dataset for the patch-
 228 level classifier. Different existing CNN architectures are evaluated for patch-wise classification, and
 229 DenseNet-201 is reported to be the best network architecture for histopathology image classification
 230 of gastric cancer, giving an accuracy of 98.698% for the testing set. This leads to the conclusion that
 231 the best-in-class network architecture for natural image classification tasks can also give promising
 232 results in gastric cancer histopathology image classification. The cancer likelihood map for whole
 233 slide image of gastric cancer is produced by aggregating the patch-wise classification results. Final
 234 slide-level classifier is trained based on Random Forest classifier, using features extracted from the

235 corresponding cancer likelihood map as the inputs. Experiment demonstrates that the slide-level
236 classifier achieves an accuracy of 97.728% for the testing set. We thus conclude that the whole slide
237 image classification method is useful for gastric cancer detection.

238 Future work can extend the binary classification of cancer/normal to multi-class classification, to
239 distinguish between various sub-types of the gastric cancer. Moreover, data used for this work are
240 from the same center using the same scanner. Further studies can explore the method with data from
241 multiple centers and different digitization equipments.

242 References

- 243 [1] Teresa Araújo, Guilherme Aresta, Eduardo Castro, José Rouco, Paulo Aguiar, Catarina Eloy, António
244 Polónia, and Aurélio Campilho. Classification of breast cancer histology images using convolutional neural
245 networks. *PLoS one*, 12(6):e0177544, 2017.
- 246 [2] Gregorio Chejfec. Robbins pathologic basis of disease. *Archives of Pathology & Laboratory Medicine*,
247 123(9):859, 1999.
- 248 [3] WQ Chen, H Li, KX Sun, RS Zheng, SW Zhang, HM Zeng, XN Zou, XY Gu, and J He. Report of
249 cancer incidence and mortality in china, 2014. *Zhonghua zhong liu za zhi [Chinese journal of oncology]*,
250 40(1):5–13, 2018.
- 251 [4] Joann G Elmore, Gary M Longton, Patricia A Carney, Berta M Geller, Tracy Onega, Anna NA Tosteson,
252 Heidi D Nelson, Margaret S Pepe, Kimberly H Allison, Stuart J Schnitt, et al. Diagnostic concordance
253 among pathologists interpreting breast biopsy specimens. *Jama*, 313(11):1122–1132, 2015.
- 254 [5] Paula Jiménez Fonseca, Alberto Carmona-Bayonas, Raquel Hernández, Ana Custodio, Juana Maria Cano,
255 Alejandra Lacalle, Isabel Echavarría, Ismael Macías, Monserrat Mangas, Laura Visa, et al. Lauren subtypes
256 of advanced gastric cancer influence survival and response to chemotherapy: real-world data from the
257 agamenon national cancer registry. *British journal of cancer*, 117(6):775, 2017.
- 258 [6] Farzad Ghaznavi, Andrew Evans, Anant Madabhushi, and Michael Feldman. Digital imaging in pathology:
259 whole-slide imaging and beyond. *Annual Review of Pathology: Mechanisms of Disease*, 8:331–359, 2013.
- 260 [7] Metin N Gurcan, Laura E Boucheron, Ali Can, Anant Madabhushi, Nasir M Rajpoot, and Bulent Yener.
261 Histopathological image analysis: A review. *IEEE reviews in biomedical engineering*, 2:147–171, 2009.
- 262 [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition.
263 In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- 264 [9] Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. Densely connected convo-
265 lutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
266 volume 1, page 3, 2017.
- 267 [10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional
268 neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- 269 [11] P Laukén. Session 5: histogenesis of intestinal and diffuse types of gastric carcinoma. *Scandinavian*
270 *Journal of Gastroenterology*, 26(sup180):160–164, 1991.
- 271 [12] Geert Litjens, Clara I Sánchez, Nadya Timofeeva, Meyke Hermesen, Iris Nagtegaal, Iringo Kovacs, Christina
272 Hulsbergen-Van De Kaa, Peter Bult, Bram Van Ginneken, and Jeroen Van Der Laak. Deep learning as a
273 tool for increased accuracy and efficiency of histopathological diagnosis. *Scientific reports*, 6:26286, 2016.
- 274 [13] Yun Liu, Krishna Gadepalli, Mohammad Norouzi, George E Dahl, Timo Kohlberger, Aleksey Boyko,
275 Subhashini Venugopalan, Aleksei Timofeev, Philip Q Nelson, Greg S Corrado, et al. Detecting cancer
276 metastases on gigapixel pathology images. *arXiv preprint arXiv:1703.02442*, 2017.
- 277 [14] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE transactions on systems,*
278 *man, and cybernetics*, 9(1):62–66, 1979.
- 279 [15] Stephen S Raab, Dana Marie Grzybicki, Janine E Janosky, Richard J Zarbo, Frederick A Meier, Chris
280 Jensen, and Stanley J Geyer. Clinical impact and frequency of anatomic pathology errors in cancer
281 diagnoses. *Cancer*, 104(10):2205–2213, 2005.
- 282 [16] Massimo Rugge, Matteo Fassan, and David Y Graham. Epidemiology of gastric cancer. In *Gastric Cancer*,
283 pages 23–34. Springer, 2015.

- 284 [17] Harshita Sharma, Norman Zerbe, Iris Klempert, Olaf Hellwich, and Peter Hufnagl. Deep convolutional
285 neural networks for automatic classification of gastric carcinoma using whole slide images in digital
286 histopathology. *Computerized Medical Imaging and Graphics*, 61:2–13, 2017.
- 287 [18] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recogni-
288 tion. *arXiv preprint arXiv:1409.1556*, 2014.
- 289 [19] Fabio Alexandre Spanhol, Luiz S Oliveira, Caroline Petitjean, and Laurent Heutte. Breast cancer histopatho-
290 logical image classification using convolutional neural networks. In *Neural Networks (IJCNN), 2016*
291 *International Joint Conference on*, pages 2560–2567. IEEE, 2016.
- 292 [20] Rosa T van der Kaaij, Petur Snaebjornsson, Francine EM Voncken, Jolanda M van Dieren, Edwin PM
293 Jansen, Karolina Sikorska, Annemieke Cats, and Johanna W van Sandick. The prognostic and potentially
294 predictive value of the laurén classification in oesophageal adenocarcinoma. *European Journal of Cancer*,
295 76:27–35, 2017.
- 296 [21] Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H Beck. Deep learning for
297 identifying metastatic breast cancer. *arXiv preprint arXiv:1606.05718*, 2016.