Clinically applicable histopathological diagnosis system for gastric cancer detection using deep learning

Zhigang Song[1,†], Shuangmei Zou[2,†], Weixun Zhou[3,†], Yong Huang[1], Liwei Shao[1], Jing Yuan[1], Xiangnan Gou[1], Wei Jin[1], Zhanbo Wang[1], Xin Chen[1], Xiaohui Ding[1], Jinhong Liu[1], Chunkai Yu[4], Calvin Ku[5], Cancheng Liu[5], Zhuo Sun[5], Gang Xu[5], Yuefeng Wang[5], Xiaoqing Zhang[5], Dandan Wang[6], Shuhao Wang[5,7,*], Wei Xu[7], Richard C. Davis[8], Huaiyin Shi[1,*]

[1]*Department of Pathology, The Chinese PLA General Hospital, Beijing, China*

[2]*Department of Pathology, National Cancer Center/National Clinical Research Center for Cancer/Cancer Hospital, Chinese Academy of Medical Sciences and Peking Union Medical College, Beijing, China*

[3]*Department of Pathology, Peking Union Medical College Hospital, Beijing, China*

[4]*Department of Pathology, Beijing Shijitan Hospital, Capital Medical University, Beijing, China*

[5]*Thorough Images, Beijing, China*

[6]*Department of Pathology, Third Hospital, School of Basic Medical Sciences, Peking University Health Science Center, Beijing, China*

[7]*Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China*

[8]*Department of Pathology, Duke University Medical Center, Durham, NC, USA*

[†] *These authors contributed equally to this work.*

[*] *Corresponding authors.*

Gastric cancer is among the malignant tumors with the highest incidence and mortality rates.

Early detection and accurate histopathological diagnosis of gastric cancer are essential factors that can help increase the chances of successful treatment. While the worldwide shortage of pathologists has imposed burdens on the current histopathology service, it also offers a unique opportunity for the use of artificial intelligence assistance systems to alleviate the workload and increase diagnostic accuracy. To the best of our knowledge, there has not been a clinically applicable histopathological assistance system with high accuracy, and can generalize to whole slide images created with diverse digital scanner models from different hospitals. Here, we report the clinically applicable artificial intelligence assistance system developed at the Chinese PLA General Hospital, China, using a deep convolutional neural network trained with 2,123 pixel-level annotated whole slide images. The model achieved a sensitivity near 100% and an average specificity of 80.6% on a real world test dataset, which included 3,212 whole slide images digitalized with three scanner models. We showed that the system would aid pathologists in improving diagnostic accuracy and preventing misdiagnosis. Moreover, we demonstrated that our system could perform robustly with 1,582 whole slide images from two other medical centers. Our study proves the feasibility and benefits of using histopathological artificial intelligence assistance systems in routine practice scenarios.

Gastric cancer is the fifth most common cancer worldwide and the third leading cause of cancer death [1]. There is a wide geographic difference in its incidence, with the highest incidence rate in East Asian populations [2]. In China, about 498,000 new cases were diagnosed in 2015, which was the second leading cause of cancer-associated mortality [3]. As early detection, accurate diagnosis and surgical intervention are crucial factors to reduce gastric cancer mortality, robust

and consistently effective pathology service is indispensable. However, there is a critical shortage of anatomical pathologists both nationally and globally, which has created overloaded workforces where diagnostic accuracy being affected [4]. An increasing number of pathology laboratories have adopted digital slides in standard practice in the form of whole slide images (WSIs) in daily routine diagnosis [5–7]. The transformation of practice from microscope to WSI has paved the way for using artificial intelligence (AI) assistance systems in pathology to overcome human limitations and reduce diagnostic error. This has allowed for innovative approaches, such as AI via deep learning [8–18], to be developed. Studies have been in the direction where algorithms are expected to be able to flag the suspicious areas, prompting the pathologists to examine the tissue thoroughly under high magnification or employ immunohistochemical (IHC) studies when necessary, and make an accurate diagnosis [19].

While recent studies have validated the effectiveness of pathology AI for tumor detection in various organ systems, such as lung [20], stomach [21], lymph node metastases in breast cancer [22–24], prostate core needle biopsies [24–26], and mesothelioma [27], we identify many non-trivial challenges should be addressed before considering application in the clinical setting. First, a deep learning model should be able to sustain a thorough test with a substantial number (i.e., thousands) of slides over a continuous time period and with WSIs procured by various brands of digital scanners. The sensitivity should be near 100% without compromising specificity too heavily. Second, with the assistance of the AI system, pathologists should be able to improve their diagnostic accuracy while not drawing out the routine reporting process. To further boost the trust of the pathologists in the AI assistance system, the model predictions should be investigated to know its strengths and

weaknesses. Finally, it is necessary to conduct a multicenter test before system deployment to guarantee the stability of the model performance across different hospitals. Previous studies have addressed some of these challenges, but none met all these criteria.

Here, we report the latest operation of AI assistance system at the Chinese PLA General Hospital (PLAGH), China, with careful consideration of the solutions to the challenges that we have discussed above. The deep learning model was trained with 2,123 pixel-level annotated digital slides from 1,500 patients, which included 958 surgical specimens (908 malignancies) and 542 biopsies (102 malignancies) with diverse tumor subtypes, details were illustrated in Figure 1a. The training slides were produced at 40$\times$ magnification (0.238 μm/pixel) by the National Medical Products Administration (NMPA) cleared KFBio KF-PRO-005 digital scanner. We developed an iPad-based annotation system and provided a standard operating procedure for data collection and annotation by 12 senior pathologists (see Supplementary Table S1). We adopted the 4th edition of the WHO Classification of Tumours of the Digestive System as the reference standard [28]. The pathologists circled the precise areas using the Apple Pencil with preset labels including malignant, benign, poor quality, and ignore. We assigned the malignant label to both high-grade intraepithelial neoplasia and carcinoma because both lesions require surgical intervention. Labels of poor quality and ignore were assigned to areas with low preparation or scanning quality and slides difficult to diagnose, respectively.

We utilized a CNN of Deeplab v3 architecture for our binary image segmentation approach, which enabled pixel-level cancer detection. The WSIs and their corresponding annotations were

split into 320×320-pixel patches at 20× magnification (0.476 μm/pixel), then were fed into the network for training. We performed carefully designed data augmentation during training. Since histopathological slides had no specific orientation, we applied random rotations by 90, 180, 270 degrees, and random flips (horizontal and vertical) to the training patches. To boost the model stability on WSIs collected from different hospitals and digitalized from various scanners, we also applied gaussian and motion blurs, color jittering in brightness (0.0-0.2), saturation (0.0-0.25), contrast (0.0-0.2), and hue (0.0-0.04). During training, we considered 'poor quality' as 'benign', and neglected loss coming from the 'ignore' class. In the inference phase, each pixel was assigned by the trained model a probability of being malignant. Slide-level prediction was obtained by sorting the probabilities of all pixel-level predictions. We adopted the top 1,000 probabilities and used the mean to represent the slide-level prediction. Compared with the commonly adopted approaches that utilize patch classification and sliding windows [29,30], the semantic segmentation approach [31–33] gave a more detail-rich prediction at the pixel level.

The AI assistance system was deployed in PLAGH and underwent a three-month (Jun. 2017 to Aug. 2017) trial run with the daily gastric dataset. Overall, 3,212 daily gastric slides from 1,814 patients (1,101/713 males/females with average ages of 54.12/54.66, see Supplementary Figure S1 for detailed distribution), which included 154 surgical specimens (118 malignancies) and 1,660 biopsies (61 malignancies). The slides were grouped biweekly and divided into six consecutive time periods. To test the model performance on data produced by different scanners, the slides were digitalized by three scanner models, including KFBio KF-PRO-005 (403 WSIs, 40×, 0.238 μm/pixel), Ventana DP200 (977 WSIs, 40×, 0.233 μm/pixel), and Hamamatsu NanoZoomer S360

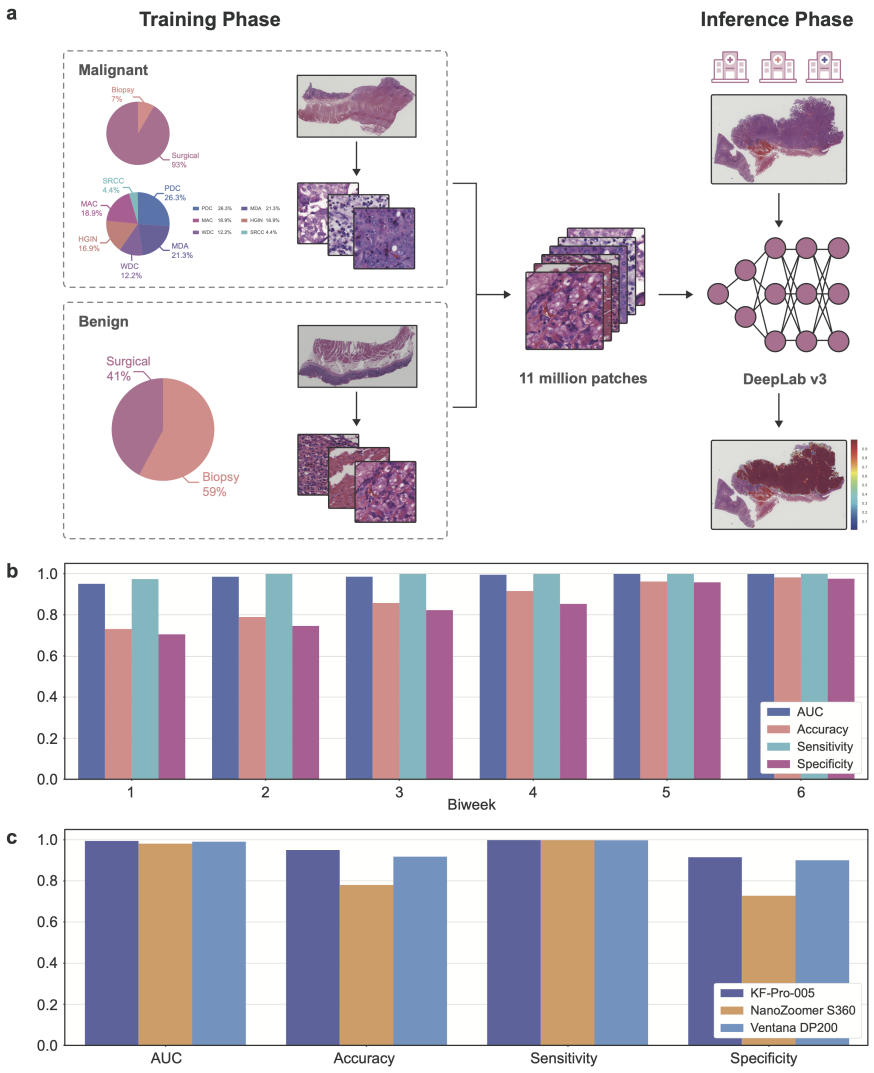

Figure 1: **The framework of our research and model performance on the daily gastric dataset.** **a**, Deep learning model training and inference. We trained the model using WSIs digitalized and annotated at PLAGH. We illustrated the training data distribution on the slide level. The abbreviations were explained in Supplementary Table S2. The trained model was tested by slides collected from PLAGH and two other hospitals. **b**, The plot of model performance histogram of the slides from the daily gastric dataset. **c**, Model performance histogram of the daily gastric slides digitalized by three different scanners.

(1832 WSIs, 40×, 0.220 μm/pixel). On this dataset, the model revealed a stable performance with an average area under the curve (AUC) of 0.986 (accuracy: 0.873, sensitivity: 0.996, specificity: 0.843) and a standard deviation of 0.018 (0.099, 0.011, 0.109) across timeline, as shown in Figure 1b (the detailed AUC curves were given in Supplementary Figure S2). We compared how the model performed on the WSIs produced by the three scanners in Figure 1c, we observed slight model performance drop on AUC (accuracy, sensitivity, specificity) of 0.004 (0.032, 0.005, 0.040) and 0.013 (0.170, 0.0, 0.210) on Ventana DP200 and Hamamatsu NanoZoomer S360, compared with KFBio KF-PRO-005, respectively (detailed results listed in Supplementary Table S3).

To prevent overpass malignancy, the AI assistance system should be able to highlight abnormal areas, prompting pathologists to give a scrutinized reassessment. During the daily gastric slides' examination, we found two missed cases that were overlooked in the initial reports and caught by the AI assistance system. The first case had received disparate diagnoses from the biopsy report and the surgical report, respectively. Cancer was found in the surgical specimen and reported in the surgical report, but because the cancer cells were limited in number, and they were missed in the initial biopsy report. In the slide, the tumor cells were scattered under the normal foveolar epithelium and only better visible under high magnification, as shown in 2a(i). The other case shown in 2a(ii) contained cancer cells deceptively bland-looking cancer cells reflecting another example that can easily be missed. Nevertheless, in the AI predicted heatmap, several red dots clearly marked the positions of the malignant tumor cells. These kinds of misdiagnoses were uncommon but possible, especially when a case was read in haste, such as the last case of the day or slides read while multitasking. The AI assistance system successfully flagged these sub-

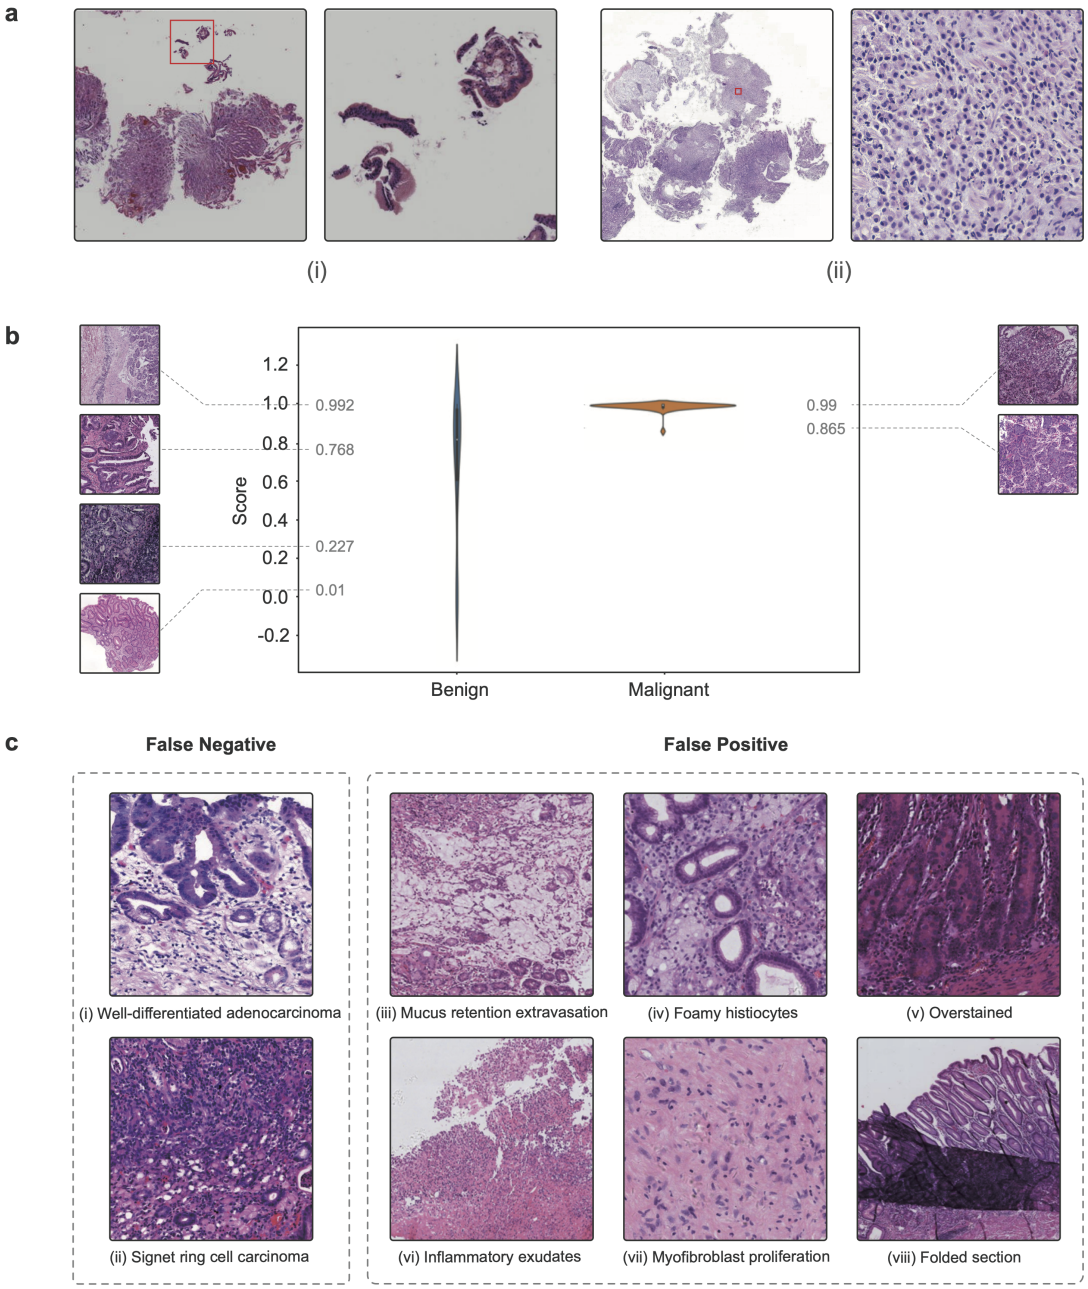

Figure 2: **Highlights of the deep learning model.** **a**, Two cases detected by the AI assistance system that were initially misdiagnosed by pathologists. **b**, Violin plot of the probability distributions for the malignant and benign cases in the IHC dataset. **c**, Eight examples for false negative and false positive cases.

tle regions, which indicated that it may alert pathologists to re-examine the slides and/or perform ancillary tests in a real world scenario.

The current AI assistance system could not only function as a pre-analytic tool to prioritize early attention to suspicious cases for review but also, as an analog to a second opinion from fellow pathologists. For difficult cases, especially for slides advised to have additional IHC stains, the model prediction had a noticeable influence on the final diagnosis. We created an IHC dataset with IHC stained slides included in the daily gastric dataset. The IHC dataset contains 5 surgical specimens (1 malignancy) and 94 biopsies (27 malignancies), with one slide selected from each case. Our model achieved an AUC of 0.923 (accuracy: 0.838, sensitivity: 0.976, specificity: 0.737). In Figure 2b, we saw clear segregation of confidence in the model between malignant and benign cases. The model performance was reasonably accurate with the malignant cases, while it showed less confidence with the benign cases as the distribution spreading out with significant larger variance. While this model only made predictions based on H&E stained WSIs, we demonstrated that our model could provide a useful visual cue using a heatmap along with providing a cancer risk probability. On the left side of Figure 2b, we also put up benign cases sorted by probability for comparison with malignant cases. We observed that the benign cases given low cancer probabilities by the model were those with clearer visual cues and could be diagnosed without resorting to IHC, whereas those with higher cancer probabilities were the more challenging cases, which required scrutinized examination under low and high magnifications, and sometimes ancillary tests.

To quest whether our AI assistance system was able to make an accurate diagnosis in real-

world scenarios, we conducted an examination using 100 slides to assess the performance of 12 junior pathologists who were under-training. As shown in Figure 3a, 100 slides were categorized into four groups depending on the degree of diagnostic difficulties: (I) Easy to diagnose under low magnification (34 WSIs); (II) Easy to diagnose but need high magnification exam (39 WSIs); (III) Difficult to diagnose, ancillary IHC not required (23 WSIs); (IV) Challenging to diagnose, require ancillary IHC (4 WSIs). We randomly divided the pathologists into three groups: a microscopy group, a WSI group, and an AI-assisted group. As the names suggested, the microscopy group worked with microscopes, the WSI group with WSIs, and the AI-assisted group with digital slides plus the AI assistance system. The examination was carried out in duplicates with a one-hour time constraint and without time constraints. In Figure 3b, we compared the performance of the pathologists with the model prediction using the receiver operating characteristic (ROC) curve. We observed that the model performance was on par with the human pathologists, even exceeding the average performance of the 12 pathologists. We discovered the AI assistance system helped the pathologists in achieving better accuracy, as shown in Figure 3c. With the help of the system, the average accuracy increased by 0.008/0.060 and 0.013/0.018 compared with microscopy and WSI groups with/without time constraints, respectively. Besides the improvement in diagnostic accuracy, the AI assistance system was able to assist the pathologists in performing more consistently, even under a time constraint. When comparing the diagnostic accuracy between the same group with/without time constraints, the digital group had a significant performance drop, sensitivity dropped 0.161 and specificity 0.052, when the time constraint was imposed, whereas the AI-assisted group showed less fluctuation, as shown in Figure 3c. The detailed experimental

results were shown in Supplementary Tables S4 and S5.

We have performed a thorough analysis of the deep learning model to further improve the pathologists' confidence in AI. As given in Figure 2c, we listed eight common failure patterns in the daily gastric dataset. The false negative (missed) cases included a well-differentiated adenocarcinomas (Figure 2c(i)) and an early atypical signet ring cell carcinoma involving only the mucosa (Figure 2c(ii)). The intramucosal well-differentiated adenocarcinoma is morphologically similar to dysplasia and had not yet caused structural disturbances and stromal desmoplasia. For the signet ring cell carcinoma case, the cancer cells were much limited. Apparently, malignancies with minimal structural disturbance in the stroma ran the risk of being overpassed. In addition, there were two situations where overdiagnosis might occur (more false positive cases were illustrated in Supplementary Figure S3). One of them was due to poor image quality, which was related to poor slide preparation, such as section folds (Figure 2c(viii)), knife marks, overstaining (Figure 2c(v)). Poor images also occurred during the digitization stage, for example, poor focus caused by the scanner. These issues may be alleviated through a better data augmentation technique or slide normalization. The second issue was that some lesions were cancer mimickers. For example, mucus extravasation resembled mucinous adenocarcinoma (Figure 2c(iii)). Correct diagnosis was easier for human pathologists when slides could be reviewed repeatedly by switching from low to high magnifications. The other case with aggregates of foamy histiocytes in lamina propria resembled signet ring cell carcinoma (Figure 2c(iv)), which again better appreciated by human pathologists under 40$\times$ magnification. Inflammatory necrotic exudates and florid granulation tissue, when there were bizarre endothelial cells and proliferated fibroblasts, could be mistaken as poorly differenti-

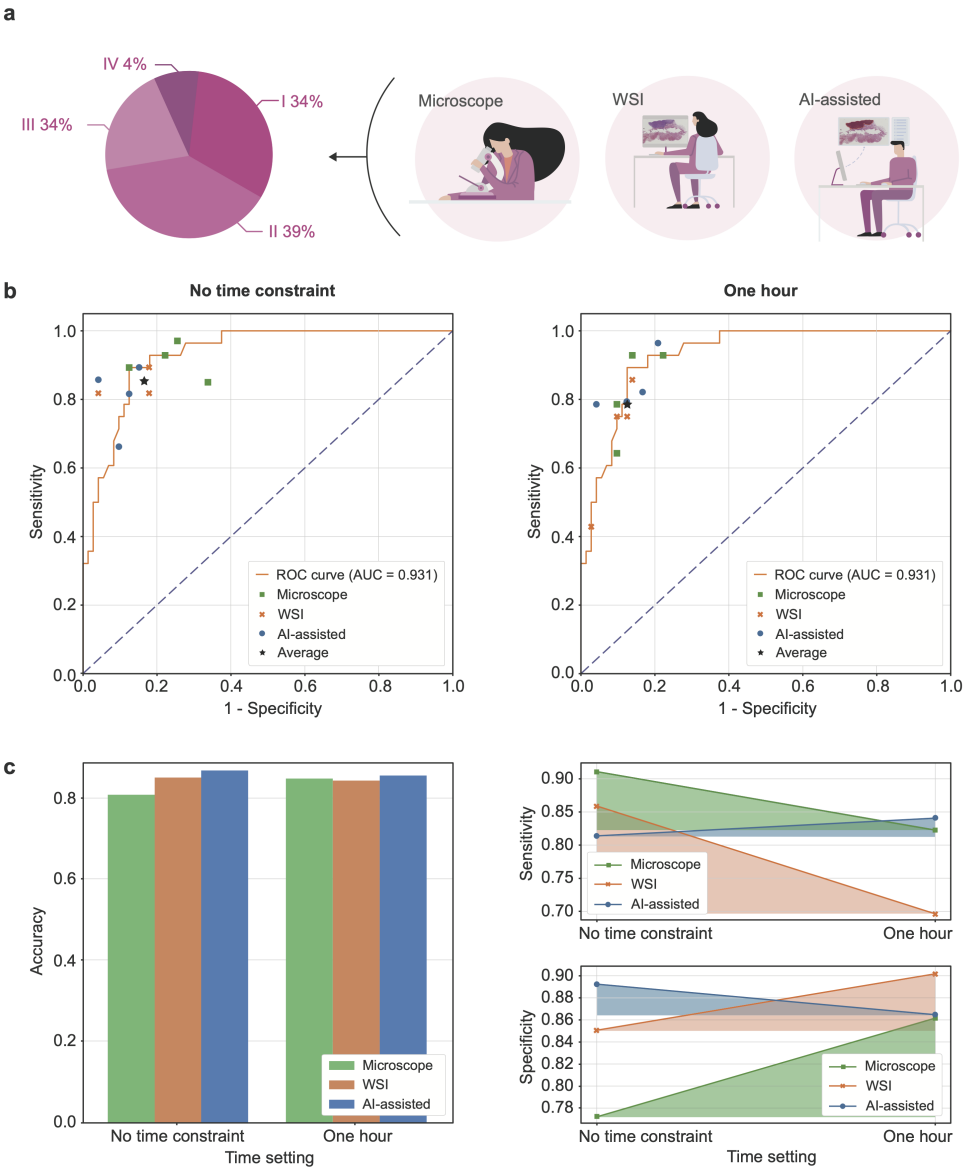

Figure 3: **Experiment settings and examination results on the performance of the pathology trainees.** **a**, Trainee pathologists were divided into three groups to make diagnosis on 100 slides of class I to VI. **b**, The model prediction ROC curve and 12 pathologists' performance in the examination. **c**, The average diagnostic accuracy of the three groups under different time settings. **d**, The diagnostic consistency among different groups.

ated adenocarcinoma (Figure 2c(vi) and (vii)). For these cases, human pathologists often needed IHC to help them make a correct diagnosis.

A mature clinically applicable AI assistance system should have robust performance on slides collected from different hospitals. To prove clinical utility with reproducible sensitivity and specificity of our deep learning model, we tested the performance of our model with slides collected from two other hospitals. We built a multicenter dataset, which included 355 cases (595 slides) from Peking Union Medical College Hospital (PUMCH) and 541 cases (987 slides) from Cancer Hospital, Chinese Academy of Medical Sciences (CHCAMS), to examine whether our model can cope with the variances created by different laboratories, such as different sectioning and staining configurations. The WSIs in the multicenter dataset were produced by the same KF-PRO-005 digital scanner with 40$\times$ magnification. To compare the model performance on the multicenter dataset and the daily gastric dataset, in this comparison, we included 403 WSIs which were produced by the KF-PRO-005 digital scanner from the daily gastric dataset to control the confounding factors. In Figure 4, we observed that the AUC (accuracy, sensitivity, specificity) on data collected from PUMCH and CHCAMS were 0.990 (0.943, 0.986, 0.937) and 0.996 (0.976, 1.0, 0.968), confirming consistent performance.

To conclude, we have shown that there is a clinical utility for using deep learning model to improve diagnostic accuracy and consistency on WSIs of gastric cancers. In our practice, to successfully build a clinically applicable histopathological assistant AI system, two factors are essential. The first and foremost one is to recruit a large number of WSIs in the training phase

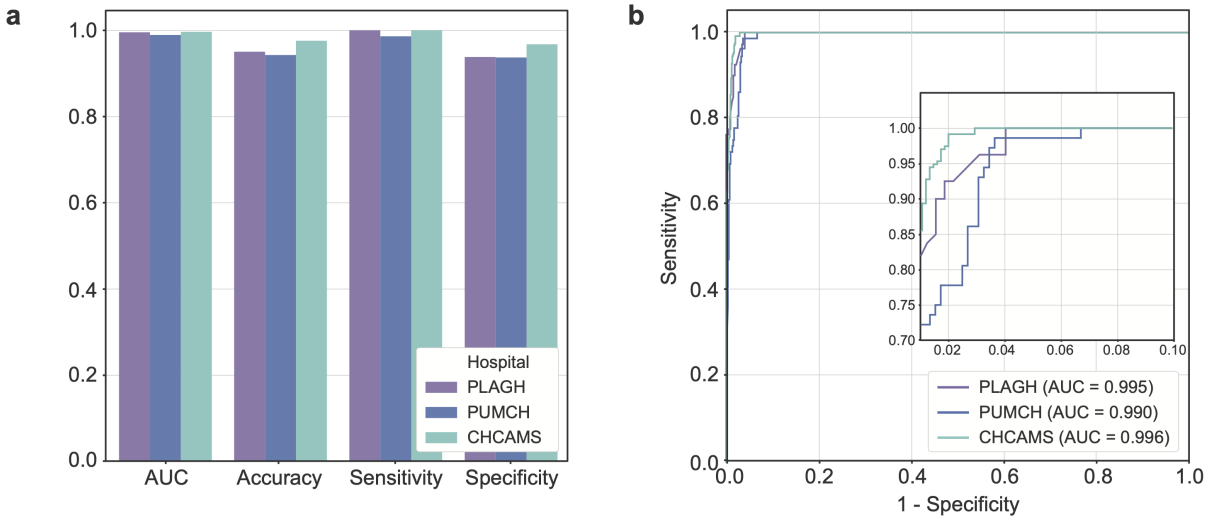

Figure 4: **Model performance on the multicenter dataset.** **a**, The AUC, accuracy, sensitivity, specificity of the deep learning model on data collected from three hospitals. **b**, ROC curves of the model on the multicenter dataset.

covering diverse tumor subtypes with accurate pixel-level annotations under a carefully designed SOP. The annotation process has to be monitored constantly by repeated reviews of model predictions to reduce the rates of false negatives and false positives. The second factor is the ability of AI model to perform pixel-level prediction based on deep convolutional neural network trained with augmented data generated from domain-specific features of histopathology. Our model-building approach can be applied in the development of histopathological AI assistance systems for a variety of cancers of different organ systems.

Methods

Datasets.

- PLAGH:
  - Training dataset: contains 2,123 WSIs (1,391 malignant tumors) as previously described.
  - Validation dataset: contains 300 WSIs for use for model hyperparameter tuning.
  - Internal examination dataset: contains 100 WSIs that were used in the collaboration test.
  - Daily gastric dataset: contains 3,212 WSIs used in the trial-run.
  - IHC dataset: a subset of the daily gastric dataset (99 WSIs) which contains difficult cases that required an immunohistochemistry.

- Multicenter: contains 595 WSIs from the Peking Union Medical College Hospital and 987 WSIs from the Cancer Hospital, Chinese Academy of Medical Sciences and Peking Union Medical College.

See Supplementary Tables S6 and S7 for a detailed description and data distribution.

Annotation procedure. Pixel-level annotations were performed by pathologists from PLAGH on 1,391 WSIs. The denotation of malignant tumors for model training was conducted using an iPad-based annotation system. The system interface was shown in Supplementary Figure S4.

The annotation procedure of a WSI comprised three steps, the initial labeling stage, the verification stage, and the final check stage. A slide was first randomly assigned to a pathologist. Once the labeling was finished, the slide and annotations were then passed on to another pathologist for review. In the final step, a senior pathologist would spot-check 30% of the slides that had passed the first two steps. The algorithm was developed gradually along with the progress of the annotation. To aid the annotation process, we also incorporated a review routine where difficult cases found during the training phase would be sent back for a second-round review.

Pre-processing. The annotations we obtained were curves with no specific stroke orders. In the data pre-processing stage, we selected the closed curves and filled in the enclosed areas to obtain pixel-level labels. Outer curves were filled first in the case of nested curves. Otsu's method was applied to the thumbnail of each WSI to obtain the tissue coordinates in the foreground. The coordinates were then rescaled to the original zoom level to obtain the WSI-level coordinates. We

only extracted training patches from coordinates that cover a tissue. During training, the WSIs were split into tiles of 320×320 pixels in size. We obtained 11,013,286 (malignant: 6,887,275, benign: 4,126,011) training patches with pixel-level annotations.

Algorithm development. We built our deep learning model based on DeepLab v3 with the ResNet-34 architecture as its backbone [33].

All models were implemented in TensorFlow [34] and trained using Adam optimizer with a fixed learning rate of 1×10^{-4}. The batch size was set to 128 (32 on each GPU). In the inference stage, we instead used larger tiles of $2,000 \times 2,000$ pixels and a 10 percent overlap ratio, by feeding $2,200 \times 2,200$-pixel tiles into the network while only using the $2,000 \times 2,000$-pixel central area for the final prediction, to further retain the environment information.

AI assistance system design. The system architecture was illustrated in Supplementary Figure S5, where we split different system components into microservices. The trained model was served by the containerized TensorFlow Serving [35]. Each worker and TensorFlow Serving pair were bound to a GPU, providing the inference service for the scheduler. Once a client initialized a prediction request, the message was passed to the pre-processing module by the message queue (MQ). Then the effective area of the whole slide image was cut into tiles and fed into the scheduler. The scheduler managed all the tasks and monitored the workers. When the predictions of all the slide tiles were complete, the post-processing module merged the tile predictions into one single slide prediction and returned it to the client through the MQ. The client could always send a message to the MQ to query the job progress. Since the communications between the microservices

were decoupled by the MQ, and the scheduler manages the tasks independently, our system was designed to be distributable with high scalability.

Collaboration test. The experiments were carried out in two conditions on the same day. In the morning, each group was asked to finish the 100 test slides within one hour. After a 3-hour break, the pathologists would be reassigned to a different group, and hence, not working under the same setting. In contrast to the morning test, the afternoon test did not have a time constraint. The pathologists were allowed to work at a self-controlled pace. The average years of experience of the attended pathologists were 4.5. For the AI-assisted group, heatmap overlay was displayed over abnormal areas, along with a probability score in the AI assistance system. The heatmap could be turned on and off with a tap on the keyboard space bar.

Evaluation metrics. We used slide-level AUC (area under the ROC curve), accuracy, sensitivity, specificity to measure model performance, and accuracy, sensitivity, specificity when comparing with human pathologists. These metrics were defined as follows

$$\begin{aligned} \text{Accuracy} &= \frac{\text{TP}+\text{TN}}{\text{TP}+\text{TN}+\text{FP}+\text{FN}}, \\ \text{Sensitivity} &= \frac{\text{TP}}{\text{TP}+\text{FN}}, \\ \text{Specificity} &= \frac{\text{TN}}{\text{TN}+\text{FP}}, \end{aligned} \tag{1}$$

where TP, TN, FP, TN represented true positive, true negative, false positive, false negative, respectively.

Plots and charts. All the plots were made using the matplotlib package in Python. The

model performance was revealed with both the ROC curve with 1 - specificity as the x-axis and sensitivity as the y-axis. We adopted bar plots showing the variance of the predictions on time-consecutive data and WSIs from different digital scanners and hospitals. We used line plots in the collaboration test result analysis to compare performance between different groups. The color fill below the lines serves the purpose of making the visual variation clearer. To study the IHC dataset, we gave a violin plot. The violin plot combined the traditional boxplot with a kernel density estimate (KDE). The KDE gave a rough estimation of the underlying data distribution. We used it to show the prediction distribution from the model, grouped by two classes (malignant and benign).

Acknowledgements The authors would like to thank Hainan Lu, Xiang Gao, Lang Wang, Lichao Pan, Fangjun Ding, Yao Lu, Li Chen, Daling Liu and Yuxi Wang at Thorough Images for data processing and helpful discussions. This work is supported by National Natural Science Foundation of China (NSFC) No. 61532001, CAMS Innovation Fund for Medical Sciences (CIFMS) No. 2018-I2M-AI-008, Medical Big Data and Artificial Intelligence Project of the Chinese PLA General Hospital, Tsinghua Initiative Research Program Grant No. 20151080475, and Beijing Hope Run Special Fund of Cancer Foundation of China No. LC2017A07.

Author Contributions Statement Z.S., S.Z., W.Z., S.W., W.X. and H.S. proposed the research, Y.H., L.S., J.Y., X.G., W.J., Z.W., X.C., X.D., J.L., C.Y., Z.S. and H.S. performed the WSI annotation, S.Z. and W.Z. led the multicenter study, Z.S. and S.W. conducted the experiment, C.K., C.L., Z.S., G.X. and Y.W. wrote the deep learning code and performed the experiment, Z.S., C.K., S.W., X.Z., D.W. and R.C.D. wrote the manuscript, H.S. and W.X. reviewed the manuscript.

Competing Interests Xiaoqing Zhang is the founder of Thorough Images. Shuhao Wang is the co-founder and chief technology officer (CTO) of Thorough Images. Calvin Ku, Cancheng Liu, Zhuo Sun, Gang Xu, Yuefeng Wang are algorithm researchers of Thorough Images. All remaining authors have declared no conflicts of interest.

Correspondence Correspondence and requests for materials should be addressed to H.S. (email: shihuaiyin@sina.com) and S.W. (email: ericwang@tsinghua.edu.cn).

1. Siegel, R. L., Miller, K. D. & Jemal, A. Cancer statistics, 2019. *CA: a Cancer Journal for Clinicians* **69**, 7–34 (2019).

2. Chen, J., Cheng, L., Xie, Z. & Li, Z. Impact of preoperative oral liquid carbohydrate on postoperative insulin resistance in gastric cancer patients and its associated study. *Chinese Journal of Gastrointestinal Surgery* **18**, 1256–1260 (2015).

3. Chen, W. *et al.* Cancer statistics in China, 2015. *CA: a Cancer Journal for Clinicians* **66**, 115–132 (2016).

4. Metter, D. M., Colgan, T. J., Leung, S. T., Timmons, C. F. & Park, J. Y. Trends in the US and Canadian pathologist workforces from 2007 to 2017. *JAMA Network Open* **2**, 194337–194337 (2019).

5. Thorstenson, S., Molin, J. & Lundström, C. Implementation of large-scale routine diagnostics using whole slide imaging in Sweden: Digital pathology experiences (2006-2013). *Journal of Pathology Informatics* **5** (2014).

6. Mukhopadhyay, S. *et al.* Whole slide imaging versus microscopy for primary diagnosis in surgical pathology: A multicenter blinded randomized noninferiority study of 1992 cases. *The American Journal of Surgical Pathology* **42**, 39 (2018).

7. Retamero, J. A., Aneiros-Fernandez, J. & Del Moral, R. G. Complete digital pathology for routine histopathology diagnosis in a multicenter hospital network. *Archives of pathology & laboratory medicine* (2019).

8. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436 (2015).

9. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* **25**, 2012 (2012).

10. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *Computer Science* (2015).

11. Szegedy, C., Liu, W., Jia, Y. & et al. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1–9 (2014).

12. He, K., Zhang, X., Ren, S. & et al. Deep residual learning for image recognition. *Computer Science* (2015).

13. Esteva, A. *et al.* Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**, 115 (2017).

14. Wang, P. *et al.* Development and validation of a deep-learning algorithm for the detection of polyps during colonoscopy. *Nature Biomedical Engineering* **2**, 741 (2018).

15. Gulshan, V. *et al.* Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama* **316**, 2402–2410 (2016).

16. Kermany, D. S. *et al.* Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell* **172**, 1122–1131 (2018).

17. De Fauw, J. *et al.* Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine* 1 (2018).

18. Litjens, G. *et al.* A survey on deep learning in medical image analysis. *Medical Image Analysis* **42**, 60–88 (2017).

19. Bera, K., Schalper, K. A., Rimm, D. L., Velcheti, V. & Madabhushi, A. Artificial intelligence in digital pathology: New tools for diagnosis and precision oncology. *Nature Reviews Clinical Oncology* **16**, 703–715 (2019).

20. Coudray, N. *et al.* Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nature Medicine* **24**, 1559 (2018).

21. Yoshida, H. *et al.* Automated histological classification of whole-slide images of gastric biopsy specimens. *Gastric Cancer* **21**, 249–257 (2018).

22. Bejnordi, B. E. *et al.* Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama* **318**, 2199–2210 (2017).

23. Bandi, P. *et al.* From detection of individual metastases to classification of lymph node status at the patient level: The CAMELYON17 challenge. *IEEE Transactions on Medical Imaging* **38**, 550–560 (2018).

24. Campanella, G. *et al.* Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine* **25**, 1301–1309 (2019).

25. Arvaniti, E. *et al.* Automated gleason grading of prostate cancer tissue microarrays via deep learning. *Scientific Reports* **8** (2018).

26. Bulten, W. *et al.* Epithelium segmentation using deep learning in H&E-stained prostate specimens with immunohistochemistry as reference standard. *Scientific Reports* **9**, 864 (2019).

27. Courtiol, P. *et al.* Deep learning-based classification of mesothelioma improves prediction of patient outcome. *Nature Medicine* **25**, 1519–1525 (2019).

28. Bosman, F. T., Carneiro, F., Hruban, R. H., Theise, N. D. *et al. WHO classification of tumours of the digestive system*. Ed. 4 (World Health Organization, 2010).

29. Liu, Y. *et al.* Detecting cancer metastases on gigapixel pathology images. *arXiv preprint arXiv:1703.02442* (2017).

30. Janowczyk, A. & Madabhushi, A. Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases. *Journal of Pathology Informatics* **7** (2016).

31. Chen, L. C., Papandreou, G., Kokkinos, I. & et al. Semantic image segmentation with deep convolutional nets and fully connected CRFs. *Computer Science* 357–361 (2014).

32. Long, J., Shelhamer, E. & Darrell, T. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3431–3440 (2015).

33. Chen, L.-C., Papandreou, G., Schroff, F. & Adam, H. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587* (2017).

34. Abadi, M. *et al.* Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation*, 265–283 (2016).

35. Olston, C. *et al.* TensorflowServing: Flexible, high-performance ML serving. *arXiv preprint arXiv:1712.06139* (2017).