

Meta-information-aware Dual-path Transformer for Differential Diagnosis of Multi-type Pancreatic Lesions in Multi-phase CT

Bo Zhou^{1,2*}, Yingda Xia¹, Jiawen Yao¹, Le Lu¹, Jingren Zhou¹, Chi Liu^{2,3},
James S. Duncan^{2,3,4}, and Ling Zhang¹

¹ Alibaba Group

² Department of Biomedical Engineering, Yale University

³ Department of Radiology and Biomedical Imaging, Yale University

⁴ Department of Electrical Engineering, Yale University

Abstract. Pancreatic cancer is one of the leading causes of cancer-related death. Accurate detection, segmentation, and differential diagnosis of the full taxonomy of pancreatic lesions, i.e., normal, seven major types of lesions, and “other” lesions, is critical to aid the clinical decision-making of patient management and treatment. However, existing work focus on segmentation and classification for very specific lesion types (PDAC) or groups. Moreover, none of the previous work considers using lesion prevalence-related non-imaging patient information to assist the differential diagnosis. To this end, we develop a meta-information-aware dual-path transformer and exploit the feasibility of classification and segmentation of the full taxonomy of pancreatic lesions. Specifically, the proposed method consists of a CNN-based segmentation path (S-path) and a transformer-based classification path (C-path). The S-path focuses on initial feature extraction by semantic segmentation using a UNet-based network. The C-path utilizes both the extracted features and meta-information for patient-level classification based on stacks of dual-path transformer blocks that enhance the modeling of global contextual information. A large-scale multi-phase CT dataset of 3,096 patients with the pathology-confirmed pancreatic lesion class labels, voxel-wise manual annotations of lesions from radiologists, and patient meta-information, was collected for training and evaluations. Our results show that our method can enable accurate classification and segmentation of the full taxonomy of pancreatic lesions, approaching the accuracy of the radiologist’s report and significantly outperforming previous baselines. Results also show that adding the common meta-information, i.e., gender and age, can boost the model’s performance, thus demonstrating the importance of meta-information for aiding pancreatic disease diagnosis.

Keywords: Pancreatic Lesion · Dual-path Transformer · Meta-information Aware · Differential Diagnosis.

* This work was supported by Alibaba Group through Alibaba Research Intern Program.

1 Introduction

Pancreatic cancer is the third leading cause of death among all cancers in the United States, and has the poorest prognosis among all solid malignancies with a 5-year survival rate of about 10% [4]. Early diagnosis and treatment are crucial, which can potentially increase the 5-year survival rate to about 50% [3]. In clinical practice, pancreatic patient management is based on the pancreatic lesion type and the potential of the lesion to become invasive cancer. However, pancreatic lesions are often hard to reach by biopsy needle because of the deep location in the abdomen and the complex structure of surrounding organs and vessels. To this end, accurate imaging-based differential diagnosis of pancreatic lesion type is critical to aid the clinical decision-making of patient management and treatment, e.g., surgery, monitoring, or discharge [11, 18]. Multi-phase Computed Tomography (CT) is the first-line imaging tool for pancreatic disease diagnosis. However, accurate differential diagnosis of pancreatic lesions is very challenging because 1) the same type of lesion may have different textures, shapes, and contrast patterns across multi-phase CT, and 2) pancreatic ductal adenocarcinoma (PDAC) accounts for the majority of cases, e.g., >60%, in pathology-confirmed patient population, leading to a long-tail problem.

Most related work in automatic pancreatic CT image analysis focus on segmentation of certain types of pancreatic lesions, e.g., PDAC and pancreatic neuroendocrine tumor (PNET). UNet-based detection-by-segmentation approaches have been extensively studied for the detection of PDAC [16, 17, 19, 20, 22] and PNET [21]. Shape-induced information, e.g., tubular structure of dilated duct, is exploited to improve the PDAC detection [9, 13]. Graph-based classification network is proposed for pancreatic patient risk stratification and management [18]. There are also recent attempts in the detection and classification of PDAC and nonPDAC using non-contrast CT [14]. However, none of the previous work has yet attempted to address the key clinical need for detection and classification of full taxonomy of pancreatic lesions, i.e., PDAC, PNET, solid pseudopapillary tumor (SPT), intraductal papillary mucinous lesion (IPMN), mucinous cystic lesion (MCN), chronic pancreatitis (CP), serous cystic lesion (SCN) [11], and other rare types that can be further classified into other benign and other malignant. Furthermore, no methods consider adding lesion prevalence-related non-imaging patient information to aid the diagnosis. For example, based on epidemiological data, the incidence of MCN, SCN, and SPT in women is significantly higher than that in men, and MCN, SCN, and SPT has a higher prevalence in young-age, middle-age, and old-age female, respectively [7]. Integrating easily-accessible clinical patient meta-information, e.g., gender and age in the DICOM head, as classification feature inputs could potentially further improve the diagnosis accuracy without needing radiologists' manual input.

To address these challenges and unmet needs, we propose a meta-information-aware dual-path transformer (MDPFormer) for classification and segmentation of the full taxonomy of pancreatic lesions, including normal, seven major types of pancreatic lesions, other malignant, and other benign. Motivated by the recent dual-path design of Mask Transformers [1, 12], the proposed MDPFormer

consists of a segmentation path (S-path) and a classification path (C-path). The S-path focuses on initial feature extraction by semantic segmentation (normal, PDAC, and nonPDAC) using a CNN-based network. Then, the C-path utilizes both meta-information and the extracted features for individual-level classification (normal, PDAC, PNET, SPT, IPMN, MCN, CP, SCN, other benign, and other malignant) based on stacked dual-path transformer blocks that enhance the modeling of global contextual information. We curated a large-scale multi-phase CT dataset with the pathology-confirmed pancreatic lesion class labels, voxel-wise manual annotations of lesions from radiologists, and patient meta-information. To our knowledge, this model is the most comprehensive to date, and is trained on a labeled dataset (2,372 patients’ multi-phase CT scans) larger than that used in previous studies [10,15,18]. We independently test our method on a test set consisting of one whole year of 724 consecutive patients with pancreatic lesions from a high-volume pancreatic cancer center. The experimental results show that our method enables accurate classification and segmentation of the full taxonomy of pancreatic lesions, approaching the accuracy of radiologists’ reports (by second-line senior readers via referring to current and previous imaging, patient history, and clinical meta-information). Our method without meta-information input demonstrates superior classification and segmentation performance as compared to previous baselines. Adding the meta-information-aware design further boosts the model’s performance, demonstrating the importance of meta-information for improving pancreatic disease diagnosis.

2 Methods

The general pipeline of our method is illustrated in Figure 1. Our pipeline consists of two stages. In the first stage, we use a localization UNet [2] to segment out the pancreas from the whole CT volume. The sub-volume containing the pancreas is then cropped out based on the segmentation mask. In the second stage, the resized sub-volume is inputted into the meta-information-aware dual-path transformer (MDPFormer) to segment and classify the pancreatic lesions. Details are elaborated in the following sections.

Meta-information-aware Dual-path Transformer. For classification, we denote $\mathcal{H}_c = \{0, 1, 2, \dots, 9\}$ for the ten patient/lesion classes, i.e., normal, PDAC, PNET, SPT, IPMN, MCN, CP, SCN, other benign, and other malignant. For segmentation, we group the last eight classes into nonPDAC and denote $\mathcal{H}_s = \{0, 1, 2\}$ for the grouped three patient classes, i.e., normal, PDAC, and nonPDAC. The goal is to enable a more balanced initial class distribution for segmentation, while enabling feature extraction for the full pancreatic lesion taxonomy classification. The training set is thus formulated as $S = \{(X_i, M_i, Y_i, Z_i) | i = 1, 2, \dots, N\}$, where X_i is the cropped pancreas sub-volume of the i -th patients, M_i is the patient meta information (gender and age), $Y_i \in \mathcal{H}_s$ is the 3-class voxel-wise annotation with the same spatial size as X_i , and $Z_i \in \mathcal{H}_c$ is the 10-class volume-wise label that confirmed by pathology or clinical records.

The MDPFormer consists of two paths, including a segmentation path (S-Path) and a classification path (C-Path). The goal of S-path is to extract rich

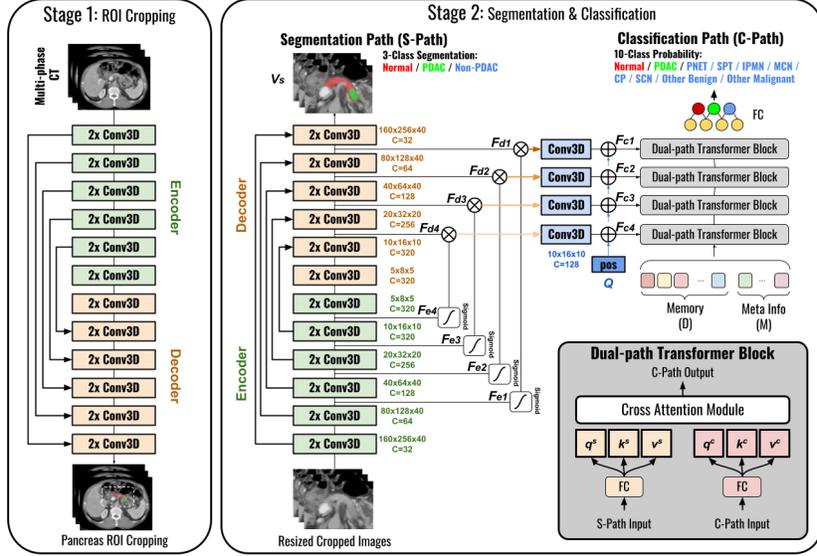


Fig. 1: The overall pipeline and the detailed structure of our MDPFormer. In stage 1, the pancreas sub-volume is cropped based on a coarse pancreas segmentation mask. In stage 2, the resized pancreas sub-volume is inputted into the MDPFormer for segmentation (left path) and classification (right path). The design of dual-path transformer block in the classification path is illustrated on the bottom right (grey box).

feature representations of the lesion and pancreas at multiple scales by first segmenting the image into three general classes. Given an input X and a segmentation network G_s , we have

$$V_s, F_{d1}, F_{d2}, F_{d3}, F_{d4}, F_{e1}, F_{e2}, F_{e3}, F_{e4} = G_s(X) \quad (1)$$

where V_s is the segmentation output, $F_{d1}, F_{d2}, F_{d3}, F_{d4}$ are the multi-scale features from the decoder, $F_{e1}, F_{e2}, F_{e3}, F_{e4}$ are the multi-scale features from the encoder. Here, we deploy a 3D UNet [2] as the S-Path backbone network. Instead of directly using the decoder features as C-path input, we combine the multi-scale encoder and decoder features by

$$F_c = f_c(F_d * \sigma(F_e)) + Q \quad (2)$$

where σ is the sigmoid function for generating attention from the encoder features to guide decoder feature outputs, f_c is a convolution layer that further refines the S-Path feature output, and Q is the learnable position embedding feature that provides position representation to aid the transformer in C-path. F_c is the extracted feature from the S-Path which is used for C-Path input.

The C-Path consists of four consecutive dual-path transformer blocks, where each block takes both the S-Path feature and the global memory feature as inputs. Denote D as the initial 1D memory feature which is randomly initialized learnable parameters [12], we fuse the patient meta-information with the initial

memory feature by

$$F_m = [D, M] \quad (3)$$

where D and M are concatenated in the length dimension and M is the meta-information, i.e., patient gender and age, in this work. In each block, we use a cross-attention module to fuse F_m and F_c . First, we compute S-Path queries q^s , keys k^s , and values v^p , by learnable linear projections of the S-Path feature F_s at each feature location. Similarly, queries q^c , keys k^c , and values v^c are computed from C-path global memory feature F_c with another set of projection matrices. The cross-attention output can then be calculated as follows:

$$y^c = \text{softmax}(q^c \cdot k^{cs})v^{cs}, \quad (4)$$

$$k^{cs} = \begin{bmatrix} k^c \\ k^s \end{bmatrix}, v^{cs} = \begin{bmatrix} v^c \\ v^s \end{bmatrix}, \quad (5)$$

where $[\cdot]$ is the concatenation operator in the channel dimension to fuse the values and keys from both paths. The output y^c is then inputted into the next block as the F_m memory feature input. Using the C-path feature output from the last dual-path transformer block, we predict the final classification P with two fully connected layers and a softmax. The overall training objective can thus be formulated as:

$$\mathcal{L}_{all} = \mathcal{L}_s(V_s, Y) + \mathcal{L}_c(P, Z) \quad (6)$$

where $\mathcal{L}_s(\cdot)$ is the Dice loss function for segmentation training, and $\mathcal{L}_c(\cdot)$ is the cross-entropy loss for classification training.

3 Experimental Results

Data Preparation. We collected a large-scale multi-phase CT dataset consisting of 3,096 patients from a high-volume pancreatic cancer institution. Each multi-phase CT consists of noncontrast, arterial, and venous phase CT. The data were consecutively collected from 2015-2020. All the 724 patients scanned during 2020 were used as the independent test set, and the rest of the 2,372 patients scanned from 2015-2019 were used as the training set. The training set includes 707 normal, 1,088 PDAC, 110 PNET, 68 SPT, 162 IPMN, 32 MCN, 64 CP, 93 SCN, 48 other benign, and 24 other malignant cases. The test set includes 202 normal, 283 PDAC, 34 PNET, 25 SPT, 73 IPMN, 9 MCN, 29 CP, 38 SCN, 14 other benign, and 17 other malignant cases. All patients with lesions were confirmed by surgical pathology, while normal patients were confirmed by radiology reports and at least 2-year follow-ups. The annotation of lesions was performed collaboratively by an experienced radiologist (with 14 years of specialized experience in pancreatic imaging) and an auto-segmentation model on either arterial or venous phase CT, whichever with better lesion visibility. More specifically, the radiologist first annotates some data to train an auto-segmentation model to segment the remaining data, which is then checked/edited by the radiologist. The CT phases were registered using DEEDS [6]. The gender and age information were extracted from the DICOM head as meta information inputs. The

gender is converted to a binary value, i.e., 0 for female and 1 for male. The age is normalized between 0-1 by dividing the value by 100.

Implementation Details. All CT volumes were resampled into $0.68 \times 0.68 \times 3.0$ mm spacing and normalized into zero mean and unit variance. In the training phase of MDPFormer, we cropped the foreground 3D bounding box of the pancreas region, randomly pad small margins on each dimension, and resized the sub-volume into $160 \times 256 \times 40$ (Y×X×Z) for input. We deployed a 5-fold cross-validation strategy using the 2,372 training set to train and validate five models. During inference, the five models’ predictions were ensemble by averaging the prediction results. For each fold, we first pre-trained the S-path network for 1000 epochs, and then trained the whole model in an end-to-end fashion with an SGD optimizer. The initial learning rate was set to 1×10^{-3} with cosine decay, and the batch size was set to 3. The localization UNet in the first stage followed the same training protocol.

Compared Methods and Evaluation Metrics. Our method is compared with two types of baseline approaches. One is the “segmentation for classification (S4C)” method where a segmentation network, i.e., nnUNet [8] or (nn)UNETR [5,8], is first deployed for semantic segmentation of the ten classes on the cropped sub-volume. We then classify the patient based on the class-wise lesion segmentation size. Specifically, if one or multiple lesion classes were presented in the segmentation, we classify the patient to the lesion class with the largest segmentation size; Otherwise, we classify the patient as normal. Note that we implement UNETR [5] in the nnUNet framework [8], called (nn)UNETR, which shows substantially better results than the original UNETR implementation on our data. The other baseline is the CNN-based segmentation-to-classification method. We use the exact same structure of S-path in MDPFormer, and extract all encoder and decoder multi-scale features. Then, we apply global max pooling on each feature map, concatenate them and forward them into two fully connected layers for classification. We also compared our performance with the radiology report, which represents the clinical read performance of second-line senior radiologists (via referring to current and previous imaging, patient history, and clinical information) in the high-volume pancreatic cancer center. The classification performance was evaluated by class-wise accuracy, regular accuracy, and balanced accuracy. The confusion matrices were also reported for detailed evaluation. The segmentation performance was evaluated by the Dice coefficient or score on each class of pancreatic lesion or normal.

3.1 Main Results

The classification results are summarised in Table 1. Comparing DPFormer without meta-information-aware to the previous baseline methods, i.e., UNet-based S4C, UNETR-based S4C, and S-Path+FC, we can see that DPFormer can already outperform all the baselines in 9 out of 10 classes and achieve the highest balanced accuracy of 49.71%. In general, it is challenging to use the conventional segmentation approaches to directly segment out the 10 classes and perform classification based on it. For S4C approaches, we can see the classification

Table 1: Evaluation of classification performance on lesion diagnosis (%). Both averaged accuracy (second last row) and balanced accuracy (last row) are reported.

<i>CLASSIFY</i>	nnUNet	(nn)UNETR	SPath+FC	DPFormer	MDPFormer	Report
Normal	96.0	96.2	97.0	99.0	99.5	100
PDAC	94.3	94.1	94.3	94.3	96.5	93.3
PNET	38.2	37.5	35.3	47.1	47.1	70.6
SPT	64.0	62.8	60.0	64.0	72.0	84.0
IPMN	69.9	68.1	43.8	60.3	65.8	68.5
MCN	0.0	0.0	11.1	11.1	33.3	33.3
CP	6.9	17.2	24.1	31.0	44.8	69.0
SCN	44.7	42.1	50.0	50.0	55.3	42.1
Other-BEN	0.0	0.0	21.4	28.6	35.7	35.7
Other-MLG	0.0	0.0	0.0	11.7	11.7	17.6
Regular Acc	77.4	77.4	76.2	79.8	82.9	84.0
Balance Acc	41.4	41.8	43.6	49.7	56.2	61.4

accuracy of MCN, CP, other benign, and other malignant are all zero or near zero. While S-Path+FC provide slightly better classification result with the additional FC layer for classification, DPFormer with dual path transformer and better feature fusion provides better results. With the meta-information-aware design that incorporates additional gender and age information, our MDPFormer utilizes those easily-accessible tumor-type-related non-imaging information, thus achieving further improved classification balanced accuracy of 56.17%.

The classification results compared to the radiology report are also shown in Figure 1 and elaborated in Figure 2. The balanced classification accuracy of the radiology report is 61.41%. Adding meta-information improves our method’s balanced classification performance from 49.71% to 56.17%, approaching the performance of the radiology report. Our method also provides better PDAC (96.5% vs. 93.3%) and SCN (55.4% vs. 42.1%) diagnosis accuracy as compared to the reports, which is critical since PDAC is of the highest priority among all pancreatic abnormalities with a 5-year survival rate of approximately 10% and is the most common type (>60% of all pathology-confirmed pancreatic lesions). In general, the radiology reports that perform diagnosis with more meta-information, e.g., patient history, tumor markers, previous report, etc., provide better classification accuracy. Thus, adding additional meta-information may further improve our method’s performance. In addition, unlike radiology reports that only give the final diagnosis, our MPDFormer provides both classification probabilities and class-wise lesion segmentation outputs with explainability. Examples of our MDPFormer’s classification and segmentation results are shown in Figure 3.

The accuracy of the “Report” for the normal class is 100% (Table 1 and Figure 2). This is because our normal cases were selected based on the radiology reports reporting an absence of pancreatic lesions. Actually, the radiologists’ specificity for the normal pancreas is 93%–96% in a pancreas CT interpretation setting [10]. Our MDPFormer has a higher specificity (99.5%) than radiologists, making it a reliable detection tool for pancreatic lesions in practice.

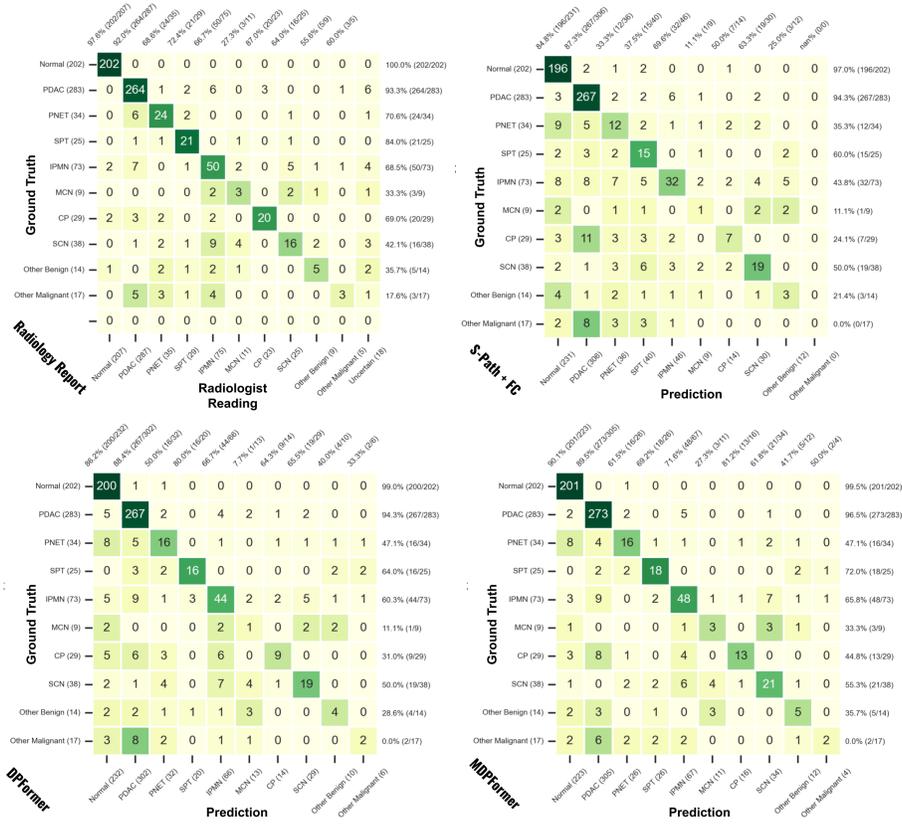


Fig. 2: Comparison of classification performance using confusion matrices.

Ablative studies for the segmentation performance are summarized in Table 2. For MDPFormer, DPFormer, and SPath+FC, please note that the nonPDAC segmentation class is assigned by the final classification prediction. Similar to the observation from classification evaluations, it is difficult for nnUNet and UNETR to directly perform 10-class segmentation with averaged Dice scores of 0.360 and 0.373 reported, respectively. On the other hand, our MDPFormer can provide significantly better segmentation performances for all 10 normal and lesion classes ($p < 0.001$) and achieve an averaged Dice score of 0.604. Comparing MDPFormer to DPFormer, we can also see that adding the meta-information improves the segmentation performance (averaged Dice of 0.604 versus 0.502). Note that the Dice scores reported in Table 2 are generally higher than that reported in previous studies [15, 17, 18]. This is mainly because the ground truth annotations are generated semi-automatically. Nevertheless, the above results clearly demonstrate the superiority of our MDPFormer over compared methods.

Next, we provide three patient case studies to show the impact of adding meta-information for classifying the pancreas lesion. The studies are illustrated in Figure 4, including three patients with MCN, SCN, and SPT, respectively.

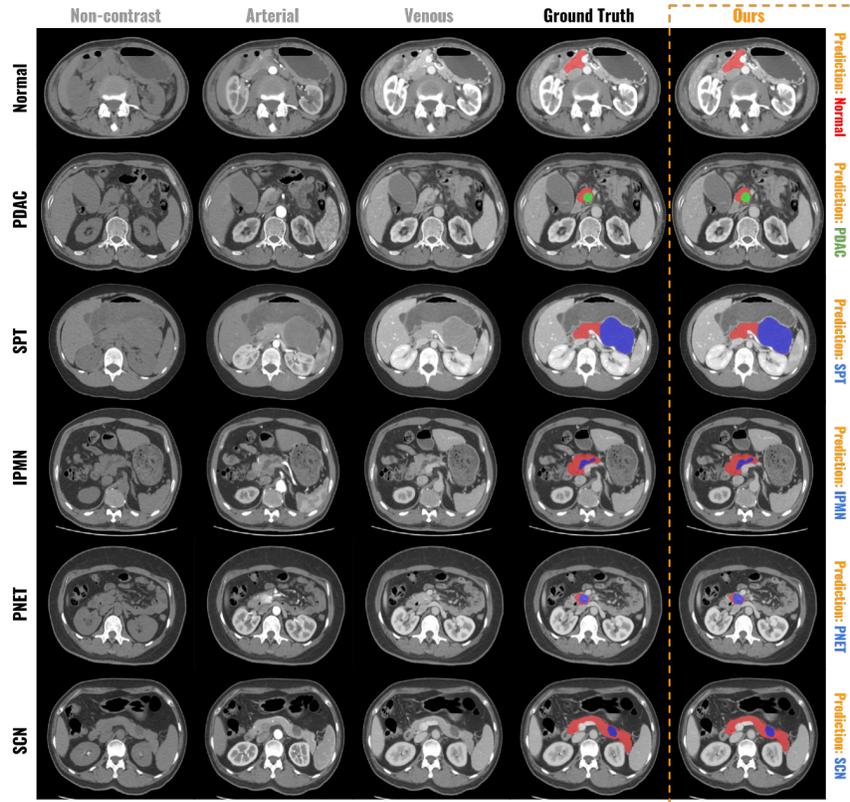


Fig. 3: Examples of classification and segmentation outputs from our MDPFormer. Ground truth lesion classes are annotated on the left and the predicted classes are shown on the right. Segmented pancreas is depicted in Red; lesion in Green or Blue.

Using DPFormer without patient meta-information, the MCN, SCN, and SPT were misclassified as other benign, IPMN, and other malignant, respectively. The MDFormer adding the gender and age information to the imaging information provide more accurate tumor probability predictions. For example, for the female 68-year-old patient with SCN, the maximal probability predicted by DPFormer is 51.63% for IPMN, while MDPFormer with meta-information provides the maximal probability of 82.37% for SCN.

3.2 Discussion

In this work, we present a meta-information-aware dual-path transformer (MDPFormer) for the classification and segmentation of pancreatic lesions in multi-phase CT. The MDFormer consists of an S-path and C-path, where the S-path focuses on initial feature extraction by group-level segmentation and the C-path utilizes both meta-information and the extracted features for individual-level classification. Compared to previous baselines, our method without meta-

Table 2: Evaluation of segmentation performance on normal and lesion (Dice).

SEGMENT	nnUNet	(nn)UNETR	SPath+FC	DPFormer	MDPFormer
Normal	0.950±0.118	0.940±0.109	0.951±0.107	0.953±0.096	0.958±0.069
PDAC	0.863±0.157	0.860±0.149	0.866±0.189	0.865±0.199	0.869±0.196
PNET	0.259±0.302	0.288±0.310	0.352±0.381	0.355±0.391	0.456±0.390
SPT	0.513±0.370	0.537±0.352	0.624±0.326	0.662±0.429	0.766±0.414
IPMN	0.475±0.304	0.468±0.302	0.515±0.340	0.518±0.390	0.598±0.382
MCN	0.071±0.159	0.098±0.189	0.211±0.446	0.312±0.395	0.416±0.441
CP	0.051±0.098	0.112±0.253	0.280±0.323	0.349±0.338	0.382±0.335
SCN	0.431±0.351	0.428±0.348	0.484±0.303	0.587±0.441	0.765±0.438
Other-BEN	0.0±0.0	0.0±0.0	0.227±0.397	0.293±0.364	0.459±0.422
Other-MLG	0.0±0.0	0.0±0.0	0.088±0.247	0.129±0.284	0.373±0.394
Average	0.361	0.373	0.464	0.502	0.604

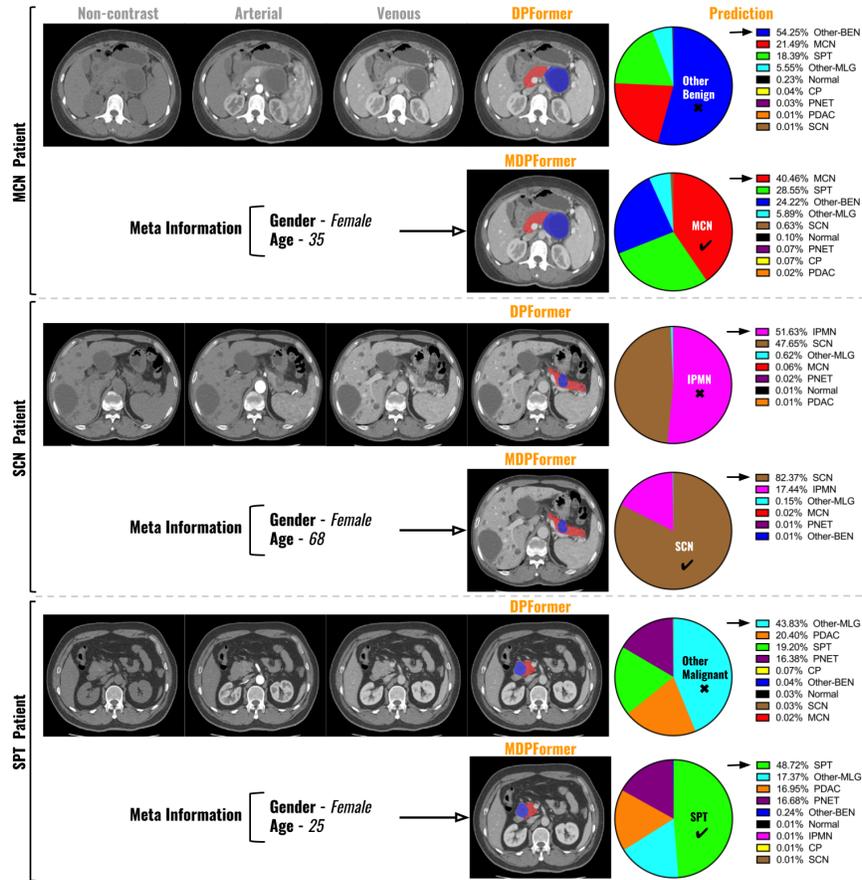


Fig. 4: Case studies of three patients with MCN, SCN, and SPT. The classification probability predictions of DPFormer and MDPFormer models are shown on the right.

information input already shows superior classification and segmentation per-

formance. Adding the meta-information-aware design further boost these performances, demonstrating the importance of meta-information when diagnosing specific pancreatic lesion type. Our MDPFormer is an open framework with several key components adjustable, which could potentially further improve our future performances. First, we used a simple UNet architecture with two consecutive convolution layers at each scale level for feature extraction. Using more advanced segmentation network blocks maybe can provide richer feature representations for better classification and segmentation performances. Second, we only used meta-information of patient gender and age as inputs, which can be automatically extracted from every DICOM data in practice. Adding additional non-imaging information, e.g., family history, symptoms (weight loss, jaundice), and other patient records (CA 19-9 blood test), may further potentially improve MDPFormer to better match the performance of the radiologists who have access to those non-imaging information for diagnosis. Those are important research directions for our future work.

4 Conclusion

This paper presents a new meta-information-aware dual-path transformer for classification and segmentation of the full taxonomy of pancreatic lesions. Our experimental results show that the proposed dual-path transformer can efficiently incorporate the patient meta-information and the extracted image features from the CNN-based segmentation path to make accurate pancreatic lesion classification and segmentation. We demonstrate that our method achieves better performance than previous baselines and approaches the accuracy of radiology reports. Our system could be a useful assistant tool for pancreatic lesion detection, segmentation, and diagnosis in the clinical reading environment.

References

1. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems* **34**, 17864–17875 (2021)
2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 424–432. Springer (2016)
3. Conroy, T., Hammel, P., Hebbbar, M., Ben Abdelghani, M., Wei, A.C., Raoul, J.L., Choné, L., Francois, E., Artru, P., Biagi, J.J., et al.: Folfirinox or gemcitabine as adjuvant therapy for pancreatic cancer. *New England Journal of Medicine* **379**(25), 2395–2406 (2018)
4. Grossberg, A.J., Chu, L.C., Deig, C.R., Fishman, E.K., Hwang, W.L., Maitra, A., Marks, D.L., Mehta, A., Nabavizadeh, N., Simeone, D.M., et al.: Multidisciplinary standards of care and recent progress in pancreatic ductal adenocarcinoma. *CA: a cancer journal for clinicians* **70**(5), 375–403 (2020)
5. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: *International MICCAI Brainlesion Workshop*. pp. 272–284. Springer (2022)

6. Heinrich, M.P., Jenkinson, M., Brady, M., Schnabel, J.A.: Mrf-based deformable registration and ventilation estimation of lung ct. *IEEE transactions on medical imaging* **32**(7), 1239–1248 (2013)
7. Hu, F., Hu, Y., Wang, D., Ma, X., Yue, Y., Tang, W., Liu, W., Wu, P., Peng, W., Tong, T.: Cystic neoplasms of the pancreas: Differential diagnosis and radiology correlation. *Frontiers in Oncology* **12** (2022)
8. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods* **18**(2), 203–211 (2021)
9. Liu, F., Xie, L., Xia, Y., Fishman, E., Yuille, A.: Joint shape representation and classification for detecting pdac. In: *International Workshop on Machine Learning in Medical Imaging*. pp. 212–220. Springer (2019)
10. Park, H.J., Shin, K., You, M.W., Kyung, S.G., Kim, S.Y., Park, S.H., Byun, J.H., Kim, N., Kim, H.J.: Deep learning-based detection of solid and cystic pancreatic neoplasms at contrast-enhanced CT. *Radiology* p. 220171 (2022)
11. Springer, S., Masica, D.L., Dal Molin, M., Douville, C., Thoburn, C.J., Afsari, B., Li, L., Cohen, J.D., Thompson, E., Allen, P.J., et al.: A multimodality test to guide the management of patients with a pancreatic cyst. *Science translational medicine* **11**(501), eaav4772 (2019)
12. Wang, H., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: Max-deeplab: End-to-end panoptic segmentation with mask transformers. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5463–5474 (2021)
13. Wang, Y., Wei, X., Liu, F., Chen, J., Zhou, Y., Shen, W., Fishman, E.K., Yuille, A.L.: Deep distance transform for tubular structure segmentation in ct scans. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3833–3842 (2020)
14. Xia, Y., Yao, J., Lu, L., Huang, L., Xie, G., Xiao, J., Yuille, A., Cao, K., Zhang, L.: Effective pancreatic cancer screening on non-contrast ct scans via anatomy-aware transformers. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 259–269. Springer (2021)
15. Xia, Y., Yu, Q., Chu, L., Kawamoto, S., Park, S., Liu, F., Chen, J., Zhu, Z., Li, B., Zhou, Z., et al.: The felix project: Deep networks to detect pancreatic neoplasms. *medRxiv* (2022)
16. Xia, Y., Yu, Q., Shen, W., Zhou, Y., Fishman, E.K., Yuille, A.L.: Detecting pancreatic ductal adenocarcinoma in multi-phase ct scans via alignment ensemble. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 285–295. Springer (2020)
17. Zhang, L., Shi, Y., Yao, J., Bian, Y., Cao, K., Jin, D., Xiao, J., Lu, L.: Robust pancreatic ductal adenocarcinoma segmentation with multi-institutional multi-phase partially-annotated ct scans. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 491–500. Springer (2020)
18. Zhao, T., Cao, K., Yao, J., Nogues, I., Lu, L., Huang, L., Xiao, J., Yin, Z., Zhang, L.: 3d graph anatomy geometry-integrated network for pancreatic mass segmentation, diagnosis, and quantitative patient management. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13743–13752 (2021)
19. Zhou, Y., Li, Y., Zhang, Z., Wang, Y., Wang, A., Fishman, E.K., Yuille, A.L., Park, S.: Hyper-pairing network for multi-phase pancreatic ductal adenocarcinoma segmentation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 155–163. Springer (2019)

20. Zhou, Y., Xie, L., Fishman, E.K., Yuille, A.L.: Deep supervision for pancreatic cyst segmentation in abdominal ct scans. In: International conference on medical image computing and computer-assisted intervention. pp. 222–230. Springer (2017)
21. Zhu, Z., Lu, Y., Shen, W., Fishman, E.K., Yuille, A.L.: Segmentation for classification of screening pancreatic neuroendocrine tumors. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3402–3408 (2021)
22. Zhu, Z., Xia, Y., Xie, L., Fishman, E.K., Yuille, A.L.: Multi-scale coarse-to-fine segmentation for screening pancreatic ductal adenocarcinoma. In: International conference on medical image computing and computer-assisted intervention. pp. 3–12. Springer (2019)