



Hyper-Pairing Network for Multi-phase Pancreatic Ductal Adenocarcinoma Segmentation

Yuyin Zhou¹(✉), Yingwei Li¹, Zhishuai Zhang¹, Yan Wang¹, Angtian Wang², Elliot K. Fishman³, Alan L. Yuille¹, and Seyoun Park³

¹ The Johns Hopkins University, Baltimore, MD 21218, USA

zhouyuyiner@gmail.com

² Huazhong University of Science and Technology, Wuhan 430074, China

³ The Johns Hopkins University School of Medicine, Baltimore, MD 21287, USA

Abstract. Pancreatic ductal adenocarcinoma (PDAC) is one of the most lethal cancers with an overall five-year survival rate of 8%. Due to subtle texture changes of PDAC, pancreatic dual-phase imaging is recommended for better diagnosis of pancreatic disease. In this study, we aim at enhancing PDAC automatic segmentation by integrating multi-phase information (*i.e.*, arterial phase and venous phase). To this end, we present Hyper-Pairing Network (HPN), a 3D fully convolution neural network which effectively integrates information from different phases. The proposed approach consists of a dual path network where the two parallel streams are interconnected with hyper-connections for intensive information exchange. Additionally, a pairing loss is added to encourage the commonality between high-level feature representations of different phases. Compared to prior arts which use single phase data, HPN reports a significant improvement up to 7.73% (from 56.21% to 63.94%) in terms of DSC.

1 Introduction

Pancreatic ductal adenocarcinoma (PDAC) is the 4th most common cancer of death with an overall five-year survival rate of 8%. Currently, detection or segmentation at localized disease stage followed by complete resection can offer the best chance of survival, *i.e.*, with a 5-year survival rate of 32%. The accurate segmentation of PDAC mass is also important for further quantitative analysis, *e.g.*, survival prediction [1]. Computed tomography (CT) is the most commonly used imaging modality for the initial evaluation of PDAC. However, textures of PDAC on CT are very subtle (Fig. 1) and therefore can be easily neglected by even experienced radiologists. To our best knowledge, the state-of-the-art on this matter is [17], which only reports an average Dice of 56.46%. For better detection of PDAC mass, dual-phase pancreas protocol using contrast-enhanced CT imaging, which is comprised of arterial and venous phases with intravenous contrast delay, are recommended.




Fig. 1. Visual comparison of arterial and venous images (after alignment) as well as the manual segmentation of normal pancreas tissues (yellow), pancreatic duct (purple) and PDAC mass (green). Orange arrows indicate the ambiguous boundaries and differences of the abnormal appearances between the two phases. Best viewed in color. (Color figure online)

In recent years, deep learning has largely advanced the field of computer-aided diagnosis (CAD), especially in the field of biomedical image segmentation [4, 10, 11, 16]. However, there are several challenges for applying existing segmentation algorithms to dual-phase images. Firstly, these algorithms are optimized for segmenting only one type of input, and therefore cannot be directly applied to handle multi-phase data. More importantly, how to properly handle the variations between different views requires a smart information exchange strategy between different phases. While how to efficiently integrate information from multi-modalities has been widely studied [3, 6, 15], the direction on learning multi-phase information has been rarely explored, especially for tumor detection and segmentation purposes.

To address these challenges, we propose a multi-phase segmentation algorithm, Hyper-Pairing Network (HPN), to enhance the segmentation performance especially for pancreatic abnormality. Following HyperDenseNet [3] which is effective on multi-modal image segmentation, we construct a dual-path network for handling multi-phase data, where each path is intended for one phase. To enable information exchange between different phases, we apply skip connections across different paths of the network [3], referred as *hyper-connections*. Moreover, by noticing that a standard segmentation loss (cross-entropy loss, Dice loss [8]) only aims at minimizing the differences between the final prediction and the groundtruth thus cannot well handle the variance between different views, we introduce an additional *pairing loss term* to encourage the commonality between high-level features across both phases for better incorporation of multi-phase information. We exploit three structures together in HPN including PDAC mass, normal pancreatic tissues, and pancreatic duct, which serves as an important clue for localizing PDAC. Extensive experiments demonstrate that the proposed HPN significantly outperforms prior arts by a large margin on all 3 targets.

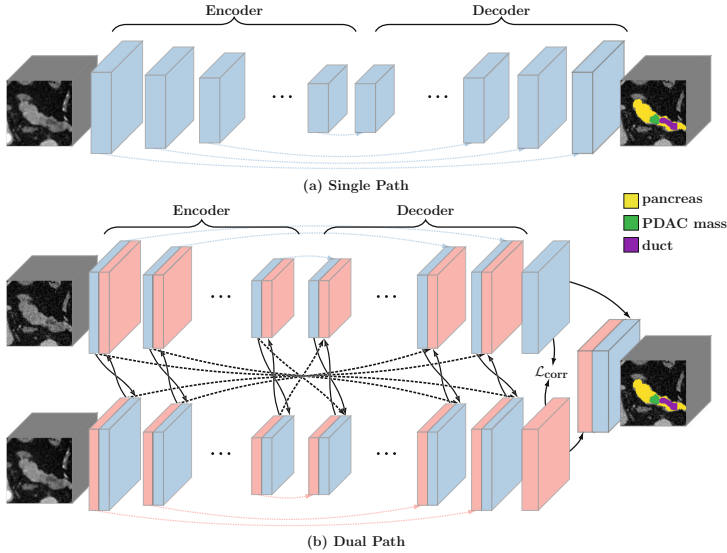


Fig. 2. (a) The single path network where only one phase is used. The dash arrows denote skip connections between low-level features and high-level features. (b) HPN structure where multiple phases are used. The black arrows between the two single path networks indicate hyper-connections between the two streams. An additional pairing loss is employed to regularize view variations, therefore can benefit the integration between different phases. Blue and pink stand for arterial and venous phase, respectively. (Color figure online)

2 Methodology

We hereby focus on dual-phase inputs while our approach can be generalized to multi-phase scans. With phase A and aligned phase B by the deformable registration, we have the set $\mathcal{S} = \{(\mathbf{X}_i^A, \mathbf{X}_i^B, \mathbf{Y}_i) \mid i = 1, \dots, M\}$, where $\mathbf{X}_i^A \in \mathbb{R}^{W_i \times H_i \times L_i}$ is the i -th 3D volumetric CT images of phase A with the dimension $(W_i \times H_i \times L_i) = \mathcal{D}_i$ and $\mathbf{X}_i^B \in \mathbb{R}^{\mathcal{D}_i}$ is the corresponding aligned volume of phase B. $\mathbf{Y}_i = \{y_{ij} \mid j = 1, \dots, \mathcal{D}_i\}$ denotes the corresponding voxel-wise label map of the i -th volume, where $y_{ij} \in \mathcal{L}$ is the label of the j -th voxel in the i -th image, and \mathcal{L} denotes the label of the target structures. In this study, $\mathcal{L} = \{\text{normal pancreatic tissues, PDAC mass, pancreatic duct}\}$. The goal is to learn a model to predict label of each voxel $\hat{\mathbf{Y}} = f(\mathbf{X}^A, \mathbf{X}^B)$ by utilizing multi-phase information.

2.1 Hyper-connections

Segmentation networks (*e.g.*, UNet [2, 10], FCN [7]) usually contain a contracting encoder part and a successive expanding decoder part to produce a full-resolution segmentation result as illustrated in Fig. 2(a). As the layer goes deeper, the output features evolve from low-level detailed representations to high-level abstract

semantic representations. The encoder part and the decoder part share an equal number of resolution steps [2, 10].

However, this type of network can only handle single-phase data. We construct a dual path network where each phase has a branch with a U-shape encoder-decoder architecture as mentioned above. These two branches are connected via hyper-connections which enrich feature representations by learning more complex combinations between the two phases. Specifically, hyper-connections are applied between layers which output feature maps of the same resolution across different paths as illustrated in Fig. 2(b). Let $\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_T$ denote the intermediate feature maps of a general segmentation network, where \mathbf{R}_t and \mathbf{R}_{T-t} share the same resolution (\mathbf{R}_t is on the encoder path and \mathbf{R}_{T-t} is on the decoder path). Hyper-connections are applied as follows: $\mathbf{R}_t^A \rightarrow \mathbf{R}_t^B$, $\mathbf{R}_t^B \rightarrow \mathbf{R}_t^A$, $\mathbf{R}_t^A \rightarrow \mathbf{R}_{T-t}^B$, $\mathbf{R}_t^B \rightarrow \mathbf{R}_{T-t}^A$, $\mathbf{R}_{T-t}^A \rightarrow \mathbf{R}_{T-t}^B$, $\mathbf{R}_{T-t}^B \rightarrow \mathbf{R}_{T-t}^A$, while maintaining the original skip connections that already occur within the same path, *i.e.*, $\mathbf{R}_t^A \rightarrow \mathbf{R}_{T-t}^A$, $\mathbf{R}_t^B \rightarrow \mathbf{R}_{T-t}^B$.

2.2 Pairing Loss

The standard loss for segmentation networks only aims at minimizing the difference between the groundtruth and the final estimation, which cannot well handle the variance between different views. Applying this loss alone is inferior in our situation since the training process involves heavy integration of both arterial information and venous information. To this end, we propose to apply an additional pairing loss, which encourages the commonality between the two sets of high-level semantic representations, to reduce view divergence.

We instantiate this additional objective as a correlation loss [13]. Mathematically, for any pair of aligned images (X_i^A, X_i^B) passing through the corresponding view sub-network, the two sets of high-level semantic representations (feature responses in later layers) corresponding to the two phases are denoted as $f_1(X_i^A; \Theta_1)$ and $f_2(X_i^B; \Theta_2)$, where the two sub-networks are parameterized by Θ_1 and Θ_2 respectively. The outputs of two branches will be simultaneously fed to the final classification layer. In order to better integrate the outcomes from the two branches, we propose to use a pairing loss which exploits the consensus of $f_1(X_i^A; \Theta_1)$ and $f_2(X_i^B; \Theta_2)$ during training. The loss is formulated as following:

$$\mathcal{L}_{corr}(X_i^A, X_i^B; \Theta) = -\frac{\sum_{j=1}^N (f_1(X_{ij}^A) - \overline{f_1(X_i^A)}) (f_2(X_{ij}^B) - \overline{f_2(X_i^B)})}{\sqrt{\sum_{j=1}^N (f_1(X_{ij}^A) - \overline{f_1(X_i^A)})^2 \sum_{j=1}^N (f_2(X_{ij}^B) - \overline{f_2(X_i^B)})^2}}, \quad (1)$$

where N denotes the total number of voxels in the i -th sample and Θ denotes the parameters of the entire network. During the training stage, we impose this additional loss to further encourage the commonality between the two intermediate outputs. The overall loss is the weighted sum of this additional penalty term and the standard voxel-wise cross-entropy loss:

$$\mathcal{L}_{total} = -\frac{1}{N} \left[\sum_{j=1}^N \sum_{k=0}^K \mathbb{1}(y_{ij} = k) \log p_{ij}^k \right] + \lambda \mathcal{L}_{corr}(X_i^A, X_i^B; \Theta), \quad (2)$$

where p_{ij}^k denotes the probability of the j -th voxel be classified as label k on the i -th sample and $\mathbf{1}(\cdot)$ is the indicator function. K is the total number of classes. The overall objective function is optimized via stochastic gradient descent.

3 Experiments

3.1 Experiment Setup

Data Acquisition. This is an institutional review board approved HIPAA compliant retrospective case control study. 239 patients with pathologically proven PDAC were retrospectively identified from the radiology and pathology databases from 2012 to 2017 and the cases with ≤ 4 cm tumor (PDAC mass) diameter were selected for the experiment. PDAC patients were scanned on a 64-slice multidetector CT scanner (Sensation 64, Siemens Healthineers) or a dual-source multidetector CT scanner (FLASH, Siemens Healthineers). PDAC patients were injected with 100–120 mL of iohexol (Omnipaque, GE Healthcare) at an injection rate of 4–5 mL/sec. Scan protocols were customized for each patient to minimize dose. Arterial phase imaging was performed with bolus triggering, usually 30 s post-injection, and venous phase imaging was performed 60 s.

Evaluation. Denote \mathcal{Y} and \mathcal{Z} as the set of foreground voxels in the ground-truth and prediction, *i.e.*, $\mathcal{Y} = \{i \mid y_i = 1\}$ and $\mathcal{Z} = \{i \mid z_i = 1\}$. The accuracy of segmentation is evaluated by the Dice-Sørensen coefficient (DSC): $\text{DSC}(\mathcal{Y}, \mathcal{Z}) = \frac{2 \times |\mathcal{Y} \cap \mathcal{Z}|}{|\mathcal{Y}| + |\mathcal{Z}|}$. We evaluate DSCs of all three targets, *i.e.*, abnormal pancreas, PDAC mass and pancreatic duct. All experiments are conducted by three-fold cross-validation, *i.e.*, training the models on two folds and testing them on the remaining one. Through our experiment, abnormal pancreas stands for the union of normal pancreatic tissues, PDAC mass and pancreatic duct. The average DSC of all cases as well as the standard deviations are reported.

3.2 Implementation Details

Our experiments were performed on the whole CT scan and the implementations are based on PyTorch. We adopt a variation of diffeomorphic demons with direction-dependent regularizations [9, 12] for accurate and efficient deformable registration between the two phases. For data pre-processing, we truncated the raw intensity values within the range $[-100, 240]$ HU and normalized each raw CT case to have zero mean and unit variance. The input sizes of all networks are set as $64 \times 64 \times 64$. The coefficient of the correlation loss λ is set as 0.5. No further post-processing strategies were applied.

We also used data augmentation during training. Different from single-phase segmentation which commonly uses rotation and scaling [5, 17], virtual sets [14] are also utilized in this work. Even though arterial and venous phase scanning are customized for each patient, the level of enhancement can be different from patients by variation of blood circulation, which causes inter-subject

Table 1. DSC (%) comparison of abnormal pancreas, PDAC mass and pancreatic duct. We report results in the format of mean \pm standard deviation.

| Method | Abnormal pancreas | PDAC mass | pancreatic duct |
|------------------------------------|------------------------------------|-------------------------------------|-------------------------------------|
| 3D-UNet-single-phase (Arterial) | 78.35 \pm 11.89 | 52.40 \pm 27.53 | 38.35 \pm 28.98 |
| 3D-UNet-single-phase (Venous) | 79.61 \pm 10.47 | 53.08 \pm 27.06 | 40.25 \pm 27.89 |
| 3D-UNet-multi-phase (fusion) | 80.05 \pm 10.56 | 52.88 \pm 26.97 | 39.06 \pm 27.33 |
| 3D-UNet-multi-phase-HyperNet | 82.45 \pm 9.98 | 54.36 \pm 26.34 | 43.27 \pm 26.33 |
| 3D-UNet-multi-phase-HyperNet-aug | 83.67 \pm 8.92 | 55.72 \pm 26.01 | 43.53 \pm 25.94 |
| 3D-UNet-multi-phase-HPN (Ours) | 84.32 \pm 8.59 | 57.10 \pm 24.76 | 44.93 \pm 24.88 |
| 3D-ResDSN-single-phase (Arterial) | 83.85 \pm 9.43 | 56.21 \pm 26.33 | 47.04 \pm 26.42 |
| 3D-ResDSN-single-phase (Venous) | 84.92 \pm 7.70 | 56.86 \pm 26.67 | 49.81 \pm 26.23 |
| 3D-ResDSN-multi-phase (fusion) | 85.52 \pm 7.84 | 57.59 \pm 26.63 | 48.49 \pm 26.37 |
| 3D-ResDSN-multi-phase-HyperNet | 85.79 \pm 8.86 | 60.87 \pm 24.95 | 54.18 \pm 24.74 |
| 3D-ResDSN-multi-phase-HyperNet-aug | 85.87 \pm 7.91 | 61.69 \pm 23.24 | 54.07 \pm 24.06 |
| 3D-ResDSN-multi-HPN (Ours) | 86.65 \pm 7.46 | 63.94 \pm 22.74 | 56.77 \pm 23.33 |

enhancement variations on each phase. Therefore we construct virtual examples by interpolating between venous and arterial data, similar to [14]. The i -th augmented training sample pair can be written as: $\tilde{X}_i^A = \lambda X_i^A + (1 - \lambda) X_i^B$, $\tilde{X}_i^B = \lambda X_i^B + (1 - \lambda) X_i^A$, where $\lambda \sim \text{Beta}(\alpha, \alpha) \in [0, 1]$. The final outcome of HPN is obtained by taking the union of predicted regions from models trained with the original paired sets and the virtual paired sets. We set the hyper-parameter $\alpha = 0.4$ following [14].

3.3 Results and Discussions

All results are summarized in Table 1. We compare the proposed HPN with the following algorithms: (1) single-phase algorithms which are trained exclusively on one phase (denoted as “single-phase”); (2) multi-phase algorithm where both arterial and venous data are trained using a dual path network bridged with hyper connections (denoted as “HyperNet”). In general, compared with single-phase algorithms, multi-phase algorithms (*i.e.*, HyperNet, HPN) observe significant improvements for all target structures. It is no surprise to observe such a phenomenon as more useful information is distilled for multi-phase algorithms.

Efficacy of Hyper-connections. To show the effectiveness of hyper-connections, output from different phases (using single-phase algorithms) are fused by taking at each position the average probability (denoted as “fusion”). However, we observe that simply fusing the outcomes from the different phases usually yield either similar or slightly better performances compared with single-phase algorithms. This indicates that simply fusing the estimations during the inference stage cannot effectively integrate multi-phase information. By contrast, hyper-connections enable the training process to be communicative between the




Fig. 3. Qualitative comparison of different methods, where HPN enhances PDAC mass segmentation (green) significantly compared with other methods. (Best viewed in color) (Color figure online)




Fig. 4. Qualitative example where HPN detects the PDAC mass (green) while single-phase methods for both phases fail. From left to right: venous and arterial images (aligned), groundtruth, predictions of single-phase algorithms, HyperNet prediction, HPN prediction (overlaid with venous and arterial images). (Best viewed in color) (Color figure online)

two phase branches and thus can efficiently elevate the performance. Note that directly applying [3] yield unsatisfactory results. Our hyper-connections are not densely connected but are carefully designed based on previous state-of-the-art on PDAC segmentation [17] for better segmentation of PDAC. Meanwhile, we show much better performance of 63.94% compared to 56.46% reported in [17].

Efficacy of Data Augmentation. From Table 1, compared with HyperNet, HyperNet-aug witnesses performance gain especially for PDAC mass (*i.e.*, from 60.87% to 61.69% for 3D-ResDSN; from 54.36% to 55.72% for 3D-UNet), which validates the usefulness of using virtual paired sets as data augmentation.

Efficacy of HPN. We can observe additional benefit of our HPN over hyperNet-aug (*e.g.*, abnormal pancreas: 85.87% to 86.65%, PDAC mass: 61.69% to 63.94%,

pancreatic duct: 54.07% to 56.77%, 3D-ResDSN). Overall, HPN observes an evident improvement compared with HyperNet, *i.e.*, abnormal pancreas: 85.79% to 86.65%, PDAC mass: 61.69% to 63.94%, pancreatic duct: 54.07% to 56.77% (3D-ResDSN). The *p-values* for testing significant difference between hyperNet and our HPN of all 3 targets are $p < 0.0001$, which suggests a general statistical improvement. We also show two qualitative examples in Fig. 3, where HPN shows much better segmentation accuracy especially for PDAC mass.

Another noteworthy fact is that 11/239 cases are false negatives which failed to detect any PDAC mass using either phase (Dice = 0%). Out of these 11 cases, 7 cases are successfully detected by HPN. An example is shown in Fig. 4—the PDAC mass is missing from both single phases and almost missing in the original HyperNet (DSC = 0.27%), but our HPN can detect a reasonable portion of the PDAC mass (DSC = 61.5%).

The deformable registration error by computing pancreas surface distances between two phases is 1.01 ± 0.52 mm (mean \pm standard deviations) which can be considered as acceptable for this study. However, the effects between different alignments can be described as a further study.

4 Conclusions

Motivated by the fact that radiologists usually rely on analyzing multi-phase data for better image interpretations, we develop an end-to-end framework, HPN, for multi-phase image segmentation. Specifically, HPN consists of a dual path network where different paths are connected for multi-phase information exchange, and an additional loss is added for removing view divergence. Extensive experiment results demonstrate that the proposed HPN can substantially and significantly improve the segmentation performance, *i.e.*, HPN reports an improvement up to 7.73% in terms of DSC compared to prior arts which use single phase data. In the future, we plan to examine the behaviour of HPN when using different alignment strategies and try to extend the current approach to other multi-phase learning problems.

Acknowledgements. This work was supported by the Lustgarten Foundation for Pancreatic Cancer Research.

References

1. Attiyeh, M.A., Chakraborty, J., Doussot, A., Langdon-Embry, L., Mainarich, S., et al.: Survival prediction in pancreatic ductal adenocarcinoma by quantitative computed tomography image analysis. *Ann. Surg. Oncol.* **25**, 1034–1042 (2018)
2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49

3. Dolz, J., Gopinath, K., Yuan, J., Lombaert, H., Desrosiers, C., Ayed, I.B.: HyperDense-Net: a hyper-densely connected CNN for multi-modal image segmentation. *TMI* **38**, 1116–1126 (2018)
4. Dou, Q., et al.: Automatic detection of cerebral microbleeds from MR images via 3D convolutional neural networks. *TMI* **35**(5), 1182–1195 (2016)
5. Kamnitsas, K., et al.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *arXiv*
6. Li, Y., et al.: Multimodal hyper-connectivity of functional networks using functionally-weighted LASSO for MCI classification. *Med. Image Anal.* **52**, 80–96 (2019)
7. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *CVPR* (2015)
8. Milletari, F., Navab, N., Ahmadi, S.: V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: *3DV* (2016)
9. Reaungamornrat, S., et al.: MIND demons: symmetric diffeomorphic deformable registration of MR and CT for image-guided spine surgery. *TMI* **35**(11), 2413–2424 (2016)
10. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
11. Roth, H.R., Lu, L., Farag, A., Sohn, A., Summers, R.M.: Spatial aggregation of holistically-nested networks for automated pancreas segmentation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9901, pp. 451–459. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_52
12. Vercauteren, T., Pennec, X., Perchange, A., Ayache, N.: Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage* **45**(1), S61–S82 (2009)
13. Yao, J., Zhu, X., Zhu, F., Huang, J.: Deep correlational learning for survival prediction from multi-modality data. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10434, pp. 406–414. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_46
14. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: beyond empirical risk minimization. In: *ICLR* (2018)
15. Zhang, W., et al.: Deep convolutional neural networks for multi-modality isointense infant brain image segmentation. *NeuroImage* **108**, 214–224 (2015)
16. Zhu, W., et al.: AnatomyNet: deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy. *Med. Phys.* **46**(2), 576–589 (2019)
17. Zhu, Z., Xia, Y., Xie, L., Fishman, E.K., Yuille, A.L.: Multi-scale coarse-to-fine segmentation for screening pancreatic ductal adenocarcinoma. *arXiv* (2018)