

Detecting Pancreatic Ductal Adenocarcinoma in Multi-phase CT Scans via Alignment Ensemble

Yingda Xia^{1,*}, Qihang Yu^{1,*}, Wei Shen¹, Yuyin Zhou¹,
Elliot K. Fishman² and Alan L. Yuille¹

¹ Johns Hopkins University

² Johns Hopkins Medical Institutions

Abstract. Pancreatic ductal adenocarcinoma (PDAC) is one of the most lethal cancers among the population. Screening for PDACs in dynamic contrast-enhanced CT is beneficial for early diagnosis. In this paper, we investigate the problem of automated detecting PDACs in multi-phase (arterial and venous) CT scans. Multiple phases provide more information than single phase, but they are unaligned and inhomogeneous in texture, making it difficult to combine cross-phase information seamlessly. We study multiple phase alignment strategies, *i.e.*, early alignment (image registration), late alignment (high-level feature registration), and slow alignment (multi-level feature registration), and suggest an ensemble of all these alignments as a promising way to boost the performance of PDAC detection. We provide an extensive empirical evaluation on two PDAC datasets and show that the proposed alignment ensemble significantly outperforms previous state-of-the-art approaches, illustrating the strong potential for clinical use.

Keywords: PDAC detection · feature alignment · pancreatic tumor segmentation

1 Introduction

Pancreatic ductal adenocarcinoma (PDAC) is the third most common cause of cancer death in the US with a dismal five-year survival of merely 9% [9]. Computed tomography (CT) is the most widely used imaging modality for the initial evaluation of suspected PDAC. However, due to the subtle early signs of PDACs in CTs, they are easily missed by even experienced radiologists.

Recently, automated PDAC detection in CT scans based on deep learning has received increasing attention [3, 4, 22, 26], which offers great opportunities in assisting radiologists to diagnosis early-stage PDACs. But, most of these methods only utilize one phase of CT scans, and thus fail to achieve satisfying results.

In this paper, we aim to develop a deep learning based PDAC detection system taking multiple phases, *i.e.*, arterial and venous, of CT scans into account. This system consists of multiple encoders, each of which encodes information

* The first two authors equally contributed to the work.

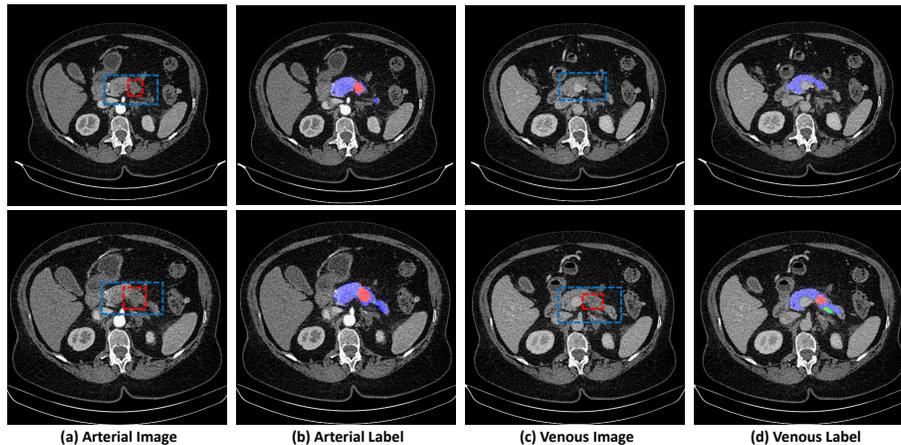


Fig. 1. Visual illustration of opportunity (top row) and challenge (bottom row) for PDAC detection in multi-phase CT scans (normal pancreas tissue - blue, pancreatic duct - green, PDAC mass - red). Top: tumor is barely visible in venous phase alone but more obvious in arterial phase. Bottom: there exist misalignment for images in these two phases given different organ size/shape and image contrast.

for one phase, and a segmentation decoder, which outputs PDAC detection results. Intuitively, multiple phases provide more information than a single phase, which certainly benefits PDAC detection. Nevertheless, how to combine this cross-phase information seamlessly is non-trivial. The challenges lie in two folds: 1) Tumor texture changes are subtle and appear differently across phases; 2) Image contents are not aligned across phases because of inevitable movements of patients during capturing multiple phases of CT scans. Consequently, a sophisticated phase alignment strategy is indispensable for detecting PDAC in multi-phase CT scans. An visual illustration is shown in Fig. 1.

We investigate several alignment strategies to combine the information across multiple phases. (1) **Early alignment**: the alignment can be done in image space by performing image registration between multiple phases; (2) **Late alignment**: it can be done late in feature space by performing spatial transformation between the encoded high-level features of multiple phases; (3) **Slow alignment**: it can be also done step-wise in feature space by aggregating multi-level feature transformations between multiple phases. Based on an extensive empirical evaluation on two PDAC datasets [22, 26], we observe that 1) All alignment strategies are beneficial for PDAC detection, 2) alignments in feature space leads to better PDAC (tumor) segmentation performance than image registration, and (3) different alignment strategies are complementary to each other, *i.e.*, an ensemble of them (**Alignment Ensemble**) significantly boosts the results, *e.g.*, approximately 4% tumor DSC score improvements over our best alignment model.

Our contributions can be summarized as follows:

- We provide extensive experimental evaluation of several phase alignment strategies for detecting PDAC in multi-phase CT scans, among which late and slow alignments are novel solutions to multi-modal segmentation problems.
- We highlight early, late and slow alignments are complementary and a simple ensemble is a promising way to boost performance of PDAC detection.
- We validate our approach on two PDAC datasets [22, 26] and achieve state-of-the-art performances on both of them.

2 Related Work

Automated Pancreas and Pancreatic Tumor Segmentation With the recent advances of deep learning, automated pancreas segmentation has achieved tremendous improvements [2, 10, 16, 17, 20, 21, 23, 25], which is an essential prerequisite for pancreatic tumor detection. Meanwhile, researchers are pacing towards automated detection of pancreatic adenocarcinoma (PDAC), the most common type of pancreatic tumor (85%) [18]. Zhu *et al.* [26] investigated using deep networks to detect PDAC in CT scans but only segmented PDAC masses in venous phase. Zhou *et al.* [22] developed the a deep learning based approach for segmenting PDACs in multi-phase CT scans, *i.e.* arterial and venous phase. They used a traditional image registration [19] approach for pre-alignment and then applied a deep network that took both phases as input. Different to their method, we also investigate how to register multiple phases in feature space.

Multi-modal Image Registration and Segmentation Multi-modal image registration [6, 7, 14, 19] is a fundamental task in medical image analysis. Recently, several deep learning based approaches, motivated by Spatial Transformer Networks [8], are proposed to address this task [1, 13, 24]. In terms of multi-modal segmentation, most of the previous works [5, 11, 22] perform segmentation on pre-registered multi-modal images. We also study these strategies for multi-modal segmentation, but we explore more, such as variants of end-to-end frameworks that jointly align multiple phases and segment target organs/tissues.

3 Methodology

3.1 Problem Statement

We aim at detecting PDACs from unaligned two-phase CT scans, *i.e.*, the venous phase and the arterial phase. Following previous works [22, 26], venous phase is our fixed phase and arterial phase is the moving one. For each patient, we have an image \mathbf{X} and its corresponding label \mathbf{Y} in the venous phase, as well as an arterial phase image \mathbf{X}' without label. The whole dataset is denoted as $S = \{(\mathbf{X}_i, \mathbf{X}'_i, \mathbf{Y}_i) | i = 1, 2, \dots, M\}$, where $\mathbf{X}_i \in \mathbb{R}^{H_i \times W_i \times D_i}$, $\mathbf{X}'_i \in \mathbb{R}^{H'_i \times W'_i \times D'_i}$ are 3D volumes representing the two-phase CT scans of the i -th patient. $\mathbf{Y}_i \in \mathcal{L}$

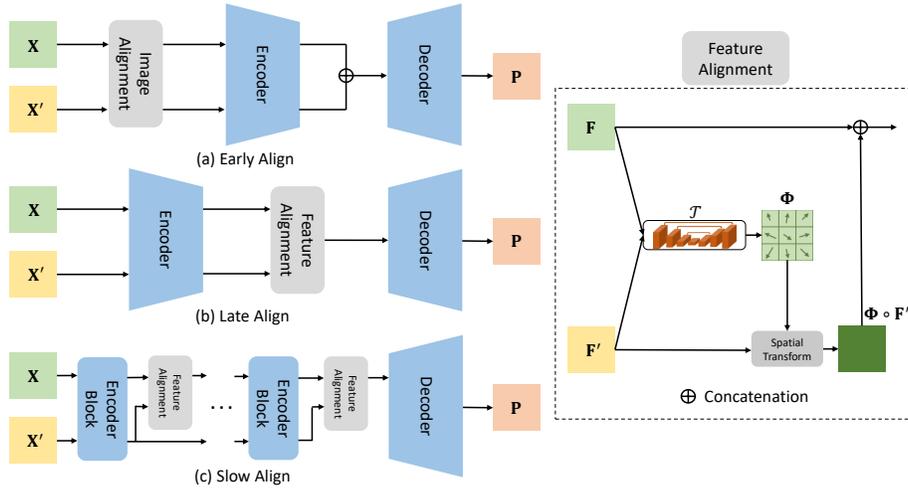


Fig. 2. An illustration of (a) early alignment (image registration) (b) late alignment and (c) slow alignment. Right: feature alignment block.

is a voxel-wise annotated label map, which have the same (H_i, W_i, D_i) three dimensional size as \mathbf{X}_i . Here, $\mathcal{L} = \{0, 1, 2, 3\}$ represents our segmentation targets, *i.e.*, background, healthy pancreas tissue, pancreatic duct (crucial for PDAC clinical diagnoses) and PDAC mass, following previous literature [22, 26]. Our goal is to find a mapping function \mathcal{M} whose inputs and outputs are a pair of two-phase images \mathbf{X}, \mathbf{X}' and segmentation results \mathbf{P} , respectively: $\mathbf{P} = \mathcal{M}(\mathbf{X}, \mathbf{X}')$. The key problem here is how to align \mathbf{X} and \mathbf{X}' , either in image space or feature space.

3.2 Cross-phase Alignment and Segmentation

As shown in Fig 2, we propose and explore three types of alignment strategies, *i.e.*, early alignment, late alignment and slow alignment, for accurate segmentation.

Early (image) alignment Early alignment, or image alignment strategy is adopted in [22] and some other multi-modal segmentation tasks such as BraTS challenge [11], where multiple phases (modalities) are first aligned by image registration algorithms and then fed forward into deep networks for segmentation. Here, we utilize a well-known registration algorithm, DEEDS [7], to estimate the registration field Φ from an arterial image \mathbf{X}' to its corresponding venous image \mathbf{X} . After registration, we use a network, consisting of two separate encoders $\mathcal{F}, \mathcal{F}'$ and a decoder \mathcal{G} , to realize the mapping function \mathcal{M} :

$$\mathbf{P} = \mathcal{M}(\mathbf{X}, \mathbf{X}') = \mathcal{G}(\mathcal{F}(\mathbf{X}) \oplus \mathcal{F}'(\Phi \circ \mathbf{X}')), \quad (1)$$

where \oplus and \circ denote the concatenation of two tensors and the element-wise deformation operations on a tensor, respectively.

This strategy relies on the accuracy of image registration algorithms for information alignment. If such algorithms produce errors, especially possible on subtle texture changes of PDACs, these errors will propagate and there will be no way to rescue (since alignment is only done on image level). Also, it remains a question that how much performance gain a segmentation algorithm will achieve through this separate registration procedure.

Late alignment An alternative way is late alignment, *i.e.*, alignment in feature space. We first encode the pair of unaligned images $(\mathbf{X}, \mathbf{X}')$ with two phase-specific encoders $(\mathcal{F}, \mathcal{F}')$, respectively. The encoded features of the two images, *i.e.*, $\mathbf{F} = \mathcal{F}(\mathbf{X})$ and $\mathbf{F}' = \mathcal{F}'(\mathbf{X}')$, are presumably in a shared feature space. We then use a network \mathcal{T} to estimate the deformable transformation field Φ from arterial (moving) to venous (fixed) in the feature space by $\Phi = \mathcal{T}(\mathbf{F}, \mathbf{F}')$. We apply the estimated transformation field Φ to feature map \mathbf{F}' , then concatenate this transformed feature map $\Phi \circ \mathbf{F}'$ to \mathbf{F} . The segmentation result \mathbf{P} is obtained by feeding the concatenation to a decoder \mathcal{G} :

$$\mathbf{P} = \mathcal{M}(\mathbf{X}, \mathbf{X}') = \mathcal{G}(\mathbf{F} \oplus \Phi \circ \mathbf{F}') = \mathcal{G}(\mathcal{F}(\mathbf{X}) \oplus \mathcal{T}(\mathbf{F}, \mathbf{F}') \circ \mathcal{F}'(\mathbf{X}')). \quad (2)$$

We name such operation as “late alignment” since the alignment is performed at the last block of feature encoders.

Slow alignment Late alignment performs one-off registration between two phases by only using high level features. However, it is known that the low level features of the deep network contain more image details, which motivates us to gradually align and propagate the features from multiple levels of the deep network. Following this spirit, we propose slow alignment, which leverages a stack of convolutional encoders and feature alignment blocks to iteratively align feature maps of two phases.

Let k be an integer which is not less than 1 and $(\mathbf{F}_{k-1}, \mathbf{F}'_{k-1})$ are the fused (aligned to the venous phase) feature map and the arterial feature map outputted by the $(k-1)^{th}$ convolutional encoder, respectively. First, they are encoded by a pair of convolutional encoders $(\mathcal{F}_k, \mathcal{F}'_k)$, respectively, which results in the venous feature map $\mathbf{F}_k = \mathcal{F}_k(\mathbf{F}_{k-1})$ and the arterial feature map $\mathbf{F}'_k = \mathcal{F}'_k(\mathbf{F}'_{k-1})$ at the k -th layer. Then a feature alignment block estimates a transformation field from the arterial (moving) phase to the venous (fixed) phase by

$$\Phi_k = \mathcal{T}_k(\mathcal{F}_k(\mathbf{F}_{k-1}), \mathcal{F}'_k(\mathbf{F}'_{k-1})), \quad (3)$$

where \mathcal{T}_k is a small U-Net. We apply the transformation field Φ_k to the arterial (moving) phase, resulting in transformed arterial feature map $\Phi_k \circ \mathcal{F}'_k(\mathbf{F}'_{k-1})$. Finally, the transformed arterial feature map is concatenated with the venous feature map $\mathcal{F}_k(\mathbf{F}_{k-1})$, resulting in the fused feature map at the k^{th} layer:

$$\mathbf{F}_k = \mathcal{F}_k(\mathbf{F}_{k-1}) \oplus \Phi_k \circ \mathcal{F}'_k(\mathbf{F}'_{k-1}). \quad (4)$$

Let us rewrite the above process by a function $\mathcal{R}_k: \mathbf{F}_k = \mathcal{R}_k(\mathbf{F}_{k-1}, \mathbf{F}'_{k-1})$ and define $\mathbf{F}_0 = \mathbf{X}$ and $\mathbf{F}'_0 = \mathbf{X}'$, then we can iteratively derive the fused feature map at n -th convolutional encoder:

$$\mathbf{F}_n = \mathcal{R}_n \left(\mathcal{R}_{n-1} \left(\cdots \left(\mathcal{R}_1(\mathbf{F}_0, \mathbf{F}'_0), \mathbf{F}'_1 \right), \cdots \right), \mathbf{F}'_{n-1} \right), \quad (5)$$

where $\mathbf{F}'_{n-1} = \mathcal{F}'_{n-1}(\mathcal{F}'_{n-2}(\cdots(\mathcal{F}'_1(\mathbf{F}'_0)))$. The final fused feature map \mathbf{F}_n is fed to the decoder \mathcal{G} to compute the segmentation result \mathbf{P} :

$$\mathbf{P} = \mathcal{M}(\mathbf{X}, \mathbf{X}') = \mathcal{G}(\mathbf{F}_n). \quad (6)$$

Alignment Ensemble We ensemble the three proposed alignment variants by simple majority voting of the predictions. The goal of the ensemble are in two folds, where the first is to improve overall performance and the second is to see whether these three alignment methods are complementary. Usually, an ensemble of complementary approaches can lead to large improvements.

4 Experiments and discussion

4.1 Dataset and evaluation

We evaluate our approach on two PDAC datasets, proposed in [26] and [22] respectively. For the ease of presentation, we regard the former as PDAC dataset I and the latter as PDAC dataset II. PDAC dataset I contains 439 CT scans in total, in which 136 cases are diagnosed with PDAC and 303 cases are normal. Annotation contains voxel-wise labeled pancreas and PDAC mass. Evaluation is done by 4 fold cross-validation on these cases following [26]. PDAC dataset II contains 239 CT scans, all from PDAC patients, with pancreas, pancreatic duct (crucial for PDAC detection) and PDAC mass annotated. Evaluation are done by 3 fold cross-validation following [22].

All cases contain two phases: arterial phase and venous phase, with a spacing of 0.5mm in axial view and all annotations are verified by experienced board certified radiologists. The segmentation accuracy is evaluated using the Dice-Sørensen coefficient (DSC): $DSC(\mathcal{Y}, \mathcal{Z}) = \frac{2 \times |\mathcal{Y} \cap \mathcal{Z}|}{|\mathcal{Y}| + |\mathcal{Z}|}$, which has a range of $[0, 1]$ with 1 implying a perfect prediction for each class. On dataset I, we also evaluate classification accuracy by sensitivity and specificity following a “segmentation for classification” strategy proposed in [26].

4.2 Implementation details

We implemented our network with PyTorch. The CT scans are first truncated within a range of HU value $[-100, 240]$ and normalized with zero mean and unit

Method	N.Pancreas	A.Pancreas	Tumor	Misses	Sens.	Spec.
U-Net [15]	86.9±8.6	81.0±10.8	57.3±28.1	10/136	92.7	99.0
V-Net [12]	87.0±8.4	81.6±10.2	57.6±27.8	11/136	91.9	99.0
MS C2F [26]	84.5± 11.1	78.6 ± 13.3	56.5± 27.2	8/136	94.1	98.5
Baseline - NA	85.8±8.0	79.5±11.2	58.4±27.4	11/136	91.9	96.0
Ours - EA	86.7±9.7	81.8±10.0	60.9±26.5	4/136	97.1	94.5
Ours - LA	87.5±7.6	82.0±10.3	62.0±27.0	7/136	94.9	96.0
Ours - SA	87.0±7.8	82.8±9.4	60.4±27.4	4/136	97.1	96.5
Ours - Ensemble	87.6±7.8	83.3±8.2	64.4±25.6	4/136	97.1	96.0

Table 1. Results on PDAC dataset I with both healthy and pathological cases. We compare our variants of alignment methods with the state-of-the-art method [26] as well as our baseline - no align (NA) version. “Misses” represents the number of cases failed in tumor detection. We also report healthy vs. pathological case classification (sensitivity and specificity) based on segmentation results. EA/LA/SA are short for early align/late align/slow align. The last row is the ensemble of the three alignments.

variance. In training stage, we randomly crop a patch size of 96^3 in roughly the same position from both arterial and venous phases. The optimization objective is Dice loss [12]. We use SGD optimizer with initial learning 0.005 and a cosine learning rate schedule for 40k iterations. For all our experiments, we implement the encoder and decoder architecture as U-Net [15] with 4 downsampling layers, making a total alignments of $n = 4$ in Eq 6. The transformation fields are estimated by light-weighted U-Nets in late align and slow align, each is $\sim 8\times$ smaller than the large U-Net for segmentation, since the inputs of the small U-Nets are already the compact encoded features. The computation of EA/LA/SA is approximately 1.5/1.7/1.9 times of the computation of a single-phase U-Net. The image registration algorithm for our early alignment is DEEDS [7].

4.3 Results

Results on dataset I and II are summarized in Table 1 and Table 2 respectively, where our approach achieves the state-of-the-art performance on both datasets. Based on the results, we have three observations which leads to three findings.

Dual-phase alignments are beneficial for detecting PDACs in multi-phase CT scans. On both datasets, our approaches, *i.e.* early align (EA), late align (LA) and slow align (AA), outperform single phase algorithms, *i.e.* U-Net [15], V-Net [12], ResDSN [25] and MS C2F [26], as well as our non-alignment dual-phase version (Baseline-NA).

Feature space alignments have larger improvements on segmentation performances than early alignments. Generally speaking for both datasets, our feature space alignment models (LA, SA) outperform image registration based approaches, *i.e.* HPN, Ours-EA, in terms of segmentation performance. Since early alignment methods apply image registration in advance, they do not guarantee a final improvement on segmentation performance. In contrast, feature space alignment methods jointly align and segment the targets

Method	A.Pancreas	Tumor	Panc. duct	Misses
U-Net [15]	79.61±10.47	53.08±27.06	40.25±27.89	11/239
ResDSN [25]	84.92±7.70	56.86±26.67	49.81±26.23	11/239
HPN-U-Net [22]	82.45±9.98	54.36±26.34	43.27±26.33	-/239
HPN-ResDSN [22]	85.79±8.86	60.87±24.95	54.18±24.74	7/239
Ours - EA	83.65±9.22	60.87±22.15	55.38±29.47	5/239
Ours - LA	86.82±6.13	62.02±24.53	64.35±29.94	9/239
Ours - SA	87.13±5.85	61.24±24.26	64.19±29.46	8/239
Ours - Ensemble	87.37±5.67	64.14±21.16	64.38±29.67	6/239

Table 2. Results on PDAC dataset II with pathological cases only. We compare our variants of alignment methods with the state-of-the-art method [22]. “Misses” represents the number of cases failed in tumor detection. EA/LA/SA are short for early align/late align/slow align. The last row is the ensemble of the three alignments.

in an end-to-end fashion by optimizing the final segmentation objective function, which leads to a larger improvements compared with single phase or naive dual phase methods without alignment. However, we indeed observe that early alignment leads to relatively less false negatives (misses).

An ensemble of the three alignment strategies significantly improve the performances. For both dataset, Ours-Ensemble achieves the best performances, illustrating that the three alignment strategies are **complementary** to each other. An ensemble leads to significant performance gain (relatively 4% improvements on tumor segmentation DSC score compared to the best alignment model from 62.0% to 64.4%) and achieves the state-of-the-art performances on both datasets. A qualitative analysis is also shown in Fig 3.

Last but not least, our alignment approaches also improve the sensitivity of healthy vs. pathological classification. In dataset I, we adopt the same “segmentation for classification” strategy as in [22], which classifies a case as pathological if we are able to detect any tumor mass larger than 50 voxels. Our approach can

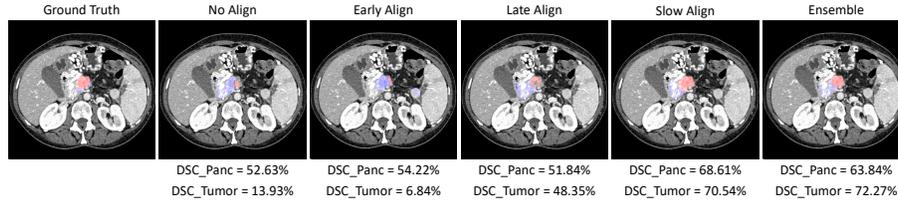


Fig. 3. An example of PDAC dataset I on venous phase. From left to right, we display ground-truth, prediction of our baseline without alignment, prediction of our early align, late align, slow align and alignment ensemble. Our feature space alignments (LA, SA) outperform no-align baseline and image registration (EA). Ensemble of the three alignment predictions also improves tumor segmentation DSC score.

improve the overall sensitivity from 94.1% to 97.1% by reducing misses from 8 to 4, which is beneficial for the early detection of PDAC. Our approach thus has valuable potential of winning precious time for early treatments for patients.

5 Conclusion

In this paper, we study three types of alignment approaches for detecting pancreatic adenocarcinoma (PDACs) in multi-phase CT scans. Early alignment first applies registration in image space and then segment with a deep network. Late alignment and slow alignment jointly align and segment with an end-to-end deep network. The former aligns in the final encoded feature space while the latter aligns multi-stage features and propagate slowly. An ensemble of the three approaches improve the performances significantly illustrating these alignment variants are complementary to each other. We achieve the state-of-the-art performances on two PDAC datasets.

References

1. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* (2019)
2. Cai, J., Lu, L., Zhang, Z., Xing, F., Yang, L., Yin, Q.: Pancreas segmentation in mri using graph-based decision fusion on convolutional neural networks. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 442–450. Springer (2016)
3. Chu, L.C., Park, S., Kawamoto, S., Fouladi, D.F., Shayesteh, S., Zinreich, E.S., Graves, J.S., Horton, K.M., Hruban, R.H., Yuille, A.L., et al.: Utility of ct radiomics features in differentiation of pancreatic ductal adenocarcinoma from normal pancreatic tissue. *American Journal of Roentgenology* **213**(2), 349–357 (2019)
4. Chu, L.C., Park, S., Kawamoto, S., Wang, Y., Zhou, Y., Shen, W., Zhu, Z., Xia, Y., Xie, L., Liu, F., et al.: Application of deep learning to pancreatic cancer detection: Lessons learned from our initial experience. *Journal of the American College of Radiology* **16**(9), 1338–1342 (2019)
5. 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. *IEEE transactions on medical imaging* **38**(5), 1116–1126 (2018)
6. Gaens, T., Maes, F., Vandermeulen, D., Suetens, P.: Non-rigid multimodal image registration using mutual information. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 1099–1106. Springer (1998)
7. 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)
8. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: *Advances in neural information processing systems*. pp. 2017–2025 (2015)
9. Lucas, A.L., Kastrinos, F.: Screening for pancreatic cancer. *Jama* **322**(5), 407–408 (2019)

10. Man, Y., Huang, Y., Feng, J., Li, X., Wu, F.: Deep q learning driven ct pancreas segmentation with geometry-aware u-net. *IEEE transactions on medical imaging* **38**(8), 1971–1980 (2019)
11. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging* **34**(10), 1993–2024 (2014)
12. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 565–571. IEEE (2016)
13. Qin, C., Shi, B., Liao, R., Mansi, T., Rueckert, D., Kamen, A.: Unsupervised deformable registration for multi-modal images via disentangled representations. In: International Conference on Information Processing in Medical Imaging. pp. 249–261. Springer (2019)
14. Roche, A., Malandain, G., Pennec, X., Ayache, N.: The correlation ratio as a new similarity measure for multimodal image registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 1115–1124. Springer (1998)
15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
16. Roth, H.R., Lu, L., Farag, A., Shin, H.C., Liu, J., Turkbey, E.B., Summers, R.M.: Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In: International conference on medical image computing and computer-assisted intervention. pp. 556–564. Springer (2015)
17. Roth, H.R., Lu, L., Farag, A., Sohn, A., Summers, R.M.: Spatial aggregation of holistically-nested networks for automated pancreas segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 451–459. Springer (2016)
18. Ryan, D.P., Hong, T.S., Bardeesy, N.: Pancreatic adenocarcinoma. *New England Journal of Medicine* **371**(11), 1039–1049 (2014)
19. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage* **45**(1), S61–S72 (2009)
20. Xia, Y., Xie, L., Liu, F., Zhu, Z., Fishman, E.K., Yuille, A.L.: Bridging the gap between 2d and 3d organ segmentation with volumetric fusion net. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 445–453. Springer (2018)
21. Yu, Q., Xie, L., Wang, Y., Zhou, Y., Fishman, E.K., Yuille, A.L.: Recurrent saliency transformation network: Incorporating multi-stage visual cues for small organ segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8280–8289 (2018)
22. 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)
23. Zhou, Y., Xie, L., Shen, W., Wang, Y., Fishman, E.K., Yuille, A.L.: A fixed-point model for pancreas segmentation in abdominal ct scans. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 693–701. Springer (2017)

24. Zhu, W., Myronenko, A., Xu, Z., Li, W., Roth, H., Huang, Y., Milletari, F., Xu, D.: Neurreg: Neural registration and its application to image segmentation. In: The IEEE Winter Conference on Applications of Computer Vision. pp. 3617–3626 (2020)
25. Zhu, Z., Xia, Y., Shen, W., Fishman, E., Yuille, A.: A 3d coarse-to-fine framework for volumetric medical image segmentation. In: 2018 International Conference on 3D Vision (3DV). pp. 682–690. IEEE (2018)
26. 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)