PanDiffuNet: Multi-Phase Diffusion Learning for Pancreatic Tumor Detection in Contrast-Enhanced CT

Dr. D. Narmadha [1,2], Prof. Shashi Kant Gupta[3],

[1] Lincoln University College, Malaysia;

[2] Division of AIML, Karunya Institute of Technology and Sciences, Coimbatore, India;

[3] Lincoln University College, Malaysia

pdf.narmadha@lincoln.edu.my
narmadha@karunya.edu
shashigupta@lincoln.edu.my

Abstract

Pancreatic cancer remains one of the most fatal malignancies, and computed tomography (CT) is the main imaging route for its assessment. Early tumors are small and difficult to distinguish from surrounding tissue, and a significant proportion of them are not detected during the routine reading. Although conventional convolutional networks like three-dimensional (3D) U-Net and subsequently transformer models have been able to boost detection, they lack the ability to cope with limited tumor samples and varying contrast phases. This work introduces a phase-conditioned diffusion framework (PanDiffuNet) for pancreatic tumor detection, localization, and segmentation. It combines diffusion-based augmentation with phase-aware denoising, cross-phase consistency and anatomical context guidance. The PanTS dataset consists of 9,901 3D abdominal CT scans and voxel-wise tumor masks in the arterial, portal-venous and delayed phases for training and evaluation. PanDiffuNet achieves an accuracy of 96.4%, a sensitivity of 95.3%, an area under the receiver operating characteristic (ROC) curve (AUROC) of 0.987, a Dice score of 0.842, and 6.9 mm 95th-percentile Hausdorff distance (HD95). The framework demonstrates that a combination of balanced synthetic augmentation and phase conditioning improves the detection accuracy and the accuracy of the boundaries for pancreatic CT analysis.

Keywords: Pancreatic tumor detection, Diffusion models, Computed tomography, Medical image segmentation, Phase conditioning, Synthetic data augmentation

1. Introduction

Pancreatic cancer is one of the most lethal solid tumors, and its prognosis is strongly related to the time of diagnosis [1]. Computed tomography is the mainstay of imaging in the abdomen and is the initial modality for the detection and staging of pancreatic disease. The organ is small, surrounded by large vessels and ducts, and early tumors tend to be attenuated with normal tissue [2]. For this reason, imaging-based detection has real clinical impact and even small improvements in reliability can tip the balance towards earlier stage disease that can be operated.

Despite the advances in imaging, many small pancreatic tumors are not detected at the time of routine imaging. Public detection models are usually trained on small, homogeneous datasets and do not carry over well to varied clinical settings [3]. The difficulty is increased by the contrast phase, as the tumor appears different in arterial, portal-venous and delayed acquisitions [4]. Tumor-positive scans are also rare compared to healthy scans, so detection models don't have many examples of the latter to learn.

Previous methods rely on convolutional segmentation networks like 3D U-Net and nnU-Net, followed by transformer-based segmentation networks such as UNETR and Swin UNETR. These methods have directly learned spatial features from the CT volumes and have shown improvement in the detection accuracy over hand-crafted pipelines [5]. However, they do not distinguish between the different phases of input and rely on the limited number of annotated tumors. Diffusion-based segmentation brought stronger boundary handling, however, phase context and data scarcity are still open problems, such as MedSegDiff.

This work presents a phase-conditioned diffusion framework, named PanDiffuNet, for detecting, localizing, and segmenting pancreatic tumors. The design combines diffusion-based augmentation for rare tumor variation, phase-conditioned denoising to align with the type of acquisition, cross-phase consistency learning, and anatomical context guidance from surrounding structures. The goal is to increase sensitivity while maintaining low false positive rates and to ensure that the predicted boundaries are near the actual tumour boundary.

The key contributions are the following:

- A phase-conditioned diffusion framework that handles arterial, portal-venous, and delayed CT within a single detection and segmentation pipeline.
- A diffusion-based augmentation scheme that supplies synthetic tumor samples, with analysis showing a balanced synthetic-to-real ratio yields the best results.
- Anatomical context guidance and cross-phase consistency that sharpen boundary accuracy and reduce false positives per scan.
- A full evaluation on the PanTS dataset, reporting 96.4% accuracy, 0.842 Dice, and 6.9 mm HD95 against convolutional and transformer baselines.

The rest of the paper is organized as follows. Section II reviews related work on pancreatic detection and diffusion-

based segmentation. Section III describes the PanDiffuNet methodology. Section IV reports the experimental setup and results. Section V discusses the findings, applications, and limits, and Section VI concludes.

2. Related Work

Hussain et al. [6] surveyed 84 studies on hepatopancreatic tumor segmentation and found that pancreatic segmentation was still difficult due to weak boundaries and small targets. Nadeem et al. [7] developed a 4-stage CAD pipeline, which achieved a detection accuracy of 99.64% using a single dataset, and used AlexNet variant. Sethi et al. [8] discussed the uses of AI in pancreatic cancer care and described the existing systems as assisting rather than replacing humans. Moglia et al. [9] reviewed 130 segmentation studies, of which 83.8% employed supervised learning, the highest Dice value was close to 91.37% and validation across centers had not yet been performed. Venkatesh et al. [11] presented PSO-CNN framework for detection and achieved 95.92% accuracy without external validation. Due to the lack of clinical data, Paramzin et al. [12] trained a U-Net with texture indices on 310 patients and achieved an accuracy of 88% and a Dice of 70%. Bi et al. [13] meta-analyzed seven studies, and pooled sensitivity and specificity were 0.92 with AUC of 0.97, but with significant heterogeneity. A pooled sensitivity of 89% and specificity of 88% have been reported in 25 AI studies that distinguished pancreatic cancer from pancreatitis [14]. The study by Oviedo et al. [15] compared pancreas segmentation on 903 CT scans and found that human-in-the-loop revision increased reliability to 0.99; however, there were no pathology cases in the test set. Zhang et al. [16] performed a meta-analysis on 24 radiomics studies for PDAC differentiation and found pooled AUC of 0.94 with low average quality. Genc et al. [17] proposed a segmentation framework that is integrated with a GUI, but is still in the preclinical phase for resource limited settings. Fang et al. [18] presented a semi-supervised CNN on the cytological images from four hospitals with external AUC of 0.99, but the cohort is small (210 patients). Xue et al. [19] presented the PMPD fusion model with internal AUROC of 0.895 and external AUROC of 0.761 for predicting PDAC metastasis.

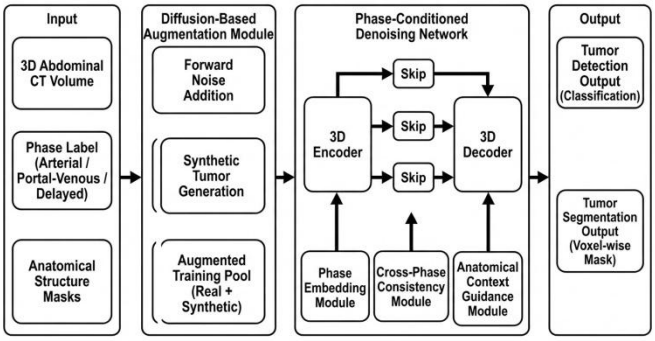

Fig. 1. Overall architecture of the proposed PanDiffuNet framework.

3. Methodology

3.1 Overview of the PanDiffuNet Framework

PanDiffuNet treats pancreatic tumor detection as a conditional denoising task carried out over three-dimensional CT volumes. Each input scan carries a contrast-phase label, and the framework uses that label to steer both augmentation and denoising. As illustrated in Fig. 1, the pipeline has four working parts: a diffusion-based augmentation stage, a phase-conditioned denoising network, a cross-phase consistency term, and anatomical context guidance. The encoder-decoder backbone stays close to a standard 3D design, so the gains come from how the diffusion process and conditioning are wired around it rather than from a heavier network.

3.2 Phase-Conditioned Diffusion Process

The forward process follows the usual diffusion setup, where a clean tumor map x_0 is gradually corrupted by Gaussian noise across T steps. The noisy sample at step t is written in (1).

$$q(x_t \mid x_0) = \mathcal{N}\big(x_t;\ \sqrt{\bar{\alpha}_t}\, x_0,\ (1-\bar{\alpha}_t)\mathbf{I}\big) \quad (1)$$

Here $\bar{\alpha}_t$ is the cumulative noise schedule up to step t. The reverse process learns to remove that noise, but unlike a plain diffusion model the prediction is conditioned on the phase embedding p and the CT volume c. The denoising network ϵ_θ estimates the added noise as shown in (2).

$$\hat{\epsilon} = \epsilon_\theta(x_t,\ t,\ c,\ p) \quad (2)$$

Feeding the phase embedding into every denoising step lets the model adapt to how the tumor appears under arterial, portal-venous, or delayed acquisition instead of assuming one fixed contrast pattern.

3.3 Training Objective and Cross-Phase Consistency

The base training loss minimizes the error between the true and predicted noise, given in (3).

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{x_0,t,\epsilon}[\ \| \epsilon - \epsilon_\theta(x_t, t, c, p) \|^2\ ] \quad (3)$$

A single tumor should read the same regardless of which phase captured it. To push the model toward that behavior, a cross-phase consistency term penalizes disagreement between predictions made under two different phase conditions p_i and p_j for the same scan, as in (4).

$$\mathcal{L}_{\text{cp}} = \big\| f_\theta(x_t, t, c, p_i) - f_\theta\big(x_t, t, c, p_j\big) \big\|^2 \quad (4)$$

where f_θ is the predicted tumor representation. The total loss combines both terms with a weight λ, shown in (5).

$$\mathcal{L} = \mathcal{L}_{\text{diff}} + \lambda\, \mathcal{L}_{\text{cp}} \quad (5)$$

3.4 Anatomical Context Guidance

The location of the pancreatic tumor in relation to vessels and nearby organs is important in distinguishing a true tumor from a false tumor. PanDiffuNet incorporates the surrounding

structures that are annotated as auxiliary signals in the denoising process, which helps to bias the model towards anatomically plausible tumor regions. This guidance is what minimises the false positives, because predictions which do not take into account the local anatomy are discouraged. These components, along with the diffusion-based augmentation, which produces additional tumor examples to compensate for the lack of positive scans, comprise the entire framework assessed in the subsequent section.

4. Results And Discussion

Table 1: Dataset Characteristics and Composition Used for the PanDiffuNet Study

Characteristic	Description / Value
Dataset	PanTS
Imaging modality	3D abdominal CT
Publicly released training scans	9,901
Arterial-phase scans	2,638
Portal-venous-phase scans	6,422
Delayed-phase scans	60
Contrast-enhanced scans used	9,120
Tumor-positive scans	1,076
Tumor annotation	Voxel-wise pancreatic tumor masks
Additional annotations	Pancreatic regions and surrounding anatomical structures
Main task	Pancreatic tumor detection
Auxiliary task	Tumor localization and segmentation
Data split strategy	Patient-aware and phase-stratified

Table 2: Comparative Pancreatic Tumor Detection Performance

Method	Accuracy (%)	Sensitivity (%)	Specificity (%)	Precision (%)	F1-Score (%)	AUROC	AUPRC
3D U-Net	89.8	83.2	90.6	78.9	81.0	0.918	0.641
nnU-Net	92.1	87.9	92.6	84.1	85.9	0.947	0.708
UNETR	92.8	88.8	93.3	85.3	87.0	0.953	0.734
Swin UNETR	93.6	90.7	94.0	87.6	89.1	0.962	0.771
MedSegDiff	94.5	92.5	94.7	89.8	91.1	0.971	0.812
PanDiffuNet	**96.4**	**95.3**	**96.5**	**93.2**	**94.2**	**0.987**	**0.874**

Table 3: Tumor Localization and Segmentation Performance

Method	Dice Score	IoU	Lesion-wise Recall (%)	Precision (%)	HD95 (mm) ↓	ASD (mm) ↓
3D U-Net	0.682	0.518	79.1	76.8	15.4	3.62
nnU-Net	0.741	0.589	84.8	82.6	12.6	2.94
UNETR	0.756	0.608	86.2	84.1	11.8	2.71
Swin UNETR	0.778	0.637	88.4	86.3	10.3	2.38
MedSegDiff	0.804	0.672	91.2	88.9	8.7	2.03
PanDiffuNet	**0.842**	**0.727**	**93.4**	**91.8**	**6.9**	**1.74**

Table 4: Phase-Specific Detection Performance of PanDiffuNet

Input Configuration	Accuracy (%)	Sensitivity (%)	Specificity (%)	Precision (%)	F1-Score (%)	AUROC	AUPRC
Arterial phase only	94.7	93.6	94.9	89.7	91.6	0.978	0.823
Portal-venous phase only	95.1	94.1	95.3	90.5	92.3	0.981	0.839
Delayed phase only	92.8	90.4	93.1	86.8	88.6	0.956	0.774
Mixed phases without phase conditioning	95.4	94.3	95.6	91.1	92.7	0.979	0.831
Mixed phases with phase embedding	95.9	94.8	96.1	92.3	93.5	0.983	0.852
Full PanDiffuNet phase-conditioned framework	**96.4**	**95.3**	**96.5**	**93.2**	**94.2**	**0.987**	**0.874**

Table 5: Ablation Study of the Proposed PanDiffuNet Framework

Model Configuration	Accuracy (%)	Sensitivity (%)	AUROC	AUPRC	Dice Score	HD95 (mm) ↓
Baseline 3D encoder-decoder	93.6	90.7	0.962	0.771	0.778	10.3
+ Diffusion-based augmentation	94.5	92.1	0.972	0.801	0.805	8.9
+ Phase-conditioned denoising	95.1	93.4	0.978	0.823	0.821	8.1
+ Cross-phase consistency learning	95.6	94.2	0.982	0.844	0.831	7.5
+ Anatomical context guidance	96.0	94.8	0.985	0.859	0.837	7.2
Full PanDiffuNet	**96.4**	**95.3**	**0.987**	**0.874**	**0.842**	**6.9**

Table 6: Effect of Diffusion-Generated Synthetic Tumor Samples

Synthetic-to-Real Augmentation Ratio	Accuracy (%)	Sensitivity (%)	AUROC	AUPRC	Dice Score	False Positives per Scan ↓
0% synthetic data	93.6	90.7	0.962	0.771	0.778	0.42
25% synthetic data	95.0	92.9	0.976	0.818	0.812	0.35
50% synthetic data	**96.4**	**95.3**	**0.987**	**0.874**	**0.842**	**0.24**
75% synthetic data	96.0	94.7	0.984	0.861	0.837	0.27
100% synthetic data	95.5	93.9	0.981	0.846	0.829	0.31

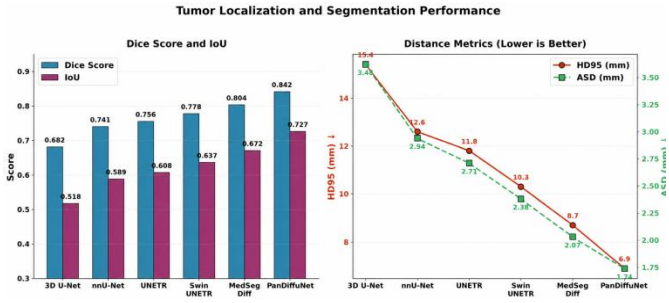

Fig. 2. Segmentation performance showing Dice/IoU scores and distance metrics (HD95, ASD) across methods.

4.1 Comparative Detection Performance

The results of the detection of the six models are summarized in Table 2. The convolutional baseline 3D U-Net achieves an accuracy of 89.8% and sensitivity of 83.2%, significantly lower than the transformer variants. nnU-Net and UNETR close the gap with 92.1% and 92.8% accuracy, respectively. Swin UNETR and MedSegDiff take it one step further and achieve 94.5% accuracy and an AUROC of 0.971, respectively. PanDiffuNet provides the best values in the table. It records 96.4% accuracy, 95.3% sensitivity, and 93.2% precision, along with an AUROC of 0.987 and AUPRC of 0.874. Sensitivity gain is important here because there is a high clinical cost of missed pancreatic tumors. This is a precision improvement of 3.4 points over MedSegDiff, indicating that the phase-conditioned design is not just detecting more voxels, but also reducing false detections.

4.2 Localization and Segmentation Quality

Overlap and boundary accuracy for the same set of models is reported in Table 3 and Fig. 2. The Dice score gradually improves from 0.682 for 3D U-Net to 0.842 for PanDiffuNet, with IoU following a similar pattern, culminating in 0.727. Lesion-wise recall is a helpful check since it counts tumors found not voxels matched. PanDiffuNet achieves a recovery rate of 93.4% and a precision of 91.8% and rarely compromises on one for the other. Boundary metrics are similar. The weakest model of the HD95 has a thickness of 15.4 mm while the ASD is reduced to 1.74 mm. The closer the boundary is to the true tumor boundary, the closer the predicted mask will be to the true boundary, which is useful when the tumor is lying adjacent to vessels. This improvement over MedSegDiff (8.7 mm HD95) is not insignificant for small pancreatic lesions.

4.3 Effect of Contrast Phase

The phase-specific behavior is shown in Table 4. Single phase is a restriction on performance. The accuracy of portal-venous input alone is 95.1%, which is slightly better than the accuracy of arterial-only (94.7%) and the sensitivity is 94.1% compared with 94.7% for arterial-only. With the lowest proportion of the data set, delayed-phase input is the weakest, with 92.8% accuracy and 90.4% sensitivity. Even in a naive setup, mixing phases without any conditioning improves accuracy to 95.4%, indicating that more input improves accuracy. Adding a phase embedding lifts it again to 95.9%. The full phase-conditioned framework reaches 96.4% accuracy and 0.987 AUROC. The gap between naive mixing and full conditioning, about one accuracy point, is where the phase-aware handling earns its place. It tells the model how each scan was acquired instead of leaving it to guess.

4.4 Component Contribution

The ablation in Table 5 and Fig. 3 traces how each part adds value. Starting from a plain 3D encoder-decoder at 93.6% accuracy and 0.778 Dice, diffusion-based augmentation is the largest single jump, raising accuracy to 94.5% and Dice to 0.805. Phase-conditioned denoising then brings accuracy to 95.1% and trims HD95 from 8.9 to 8.1 mm. Cross-phase consistency learning and anatomical context guidance each add smaller but steady gains, reaching 96.0% accuracy and 0.985 AUROC before the final assembly. The full PanDiffuNet closes at 96.4% accuracy, 95.3% sensitivity, 0.842 Dice, and 6.9 mm HD95. No component looks redundant. The pattern suggests augmentation drives the coarse gain while the conditioning and guidance modules refine boundaries and cut error at the margins.

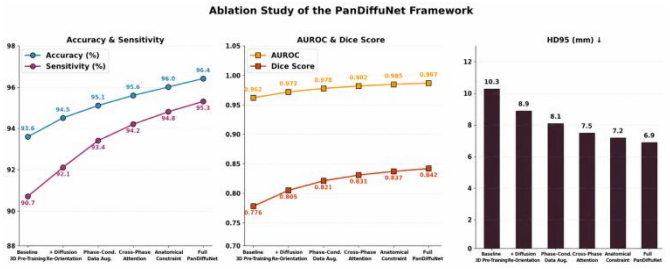

Fig. 3. Ablation analysis showing incremental contributions of each PanDiffuNet component.

4.5 Synthetic Augmentation Ratio

The last study, in Table 6 and Fig. 4, looks at how much synthetic tumor data helps. With no synthetic samples the model gives 93.6% accuracy and 0.42 false positives per scan. Synthetic data adds value up to a certain extent. The framework peaks at a 50% synthetic-to-real ratio with 96.4% accuracy, 95.3% sensitivity, 0.842 Dice, and the lowest false-positive rate of 0.24 per scan. After that, the curve bends. At 75% accuracy slips to 96.0%, and at 100% it falls to 95.5% with false positives rising back to 0.31. The pattern is easy to read. A balanced mix supplies rare tumor variation, but over-reliance on generated data starts to pull the model away from real distributions and hurts both accuracy and specificity.

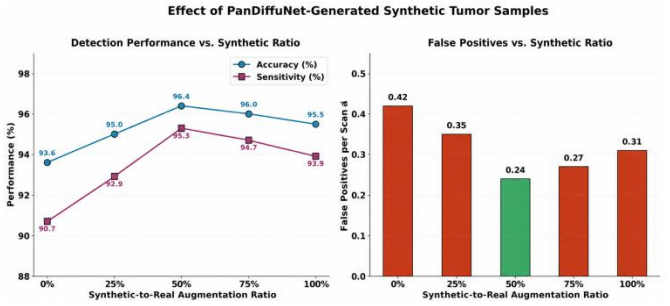

Fig. 4. Impact of synthetic-to-real augmentation ratio on detection performance and false positive rate.

5. Conclusion

This study set out to improve pancreatic tumor detection and segmentation on contrast-enhanced CT, and the phase-conditioned PanDiffuNet framework meets that aim across the reported experiments. On the PanTS scans it records 96.4% accuracy, 95.3% sensitivity, 0.987 AUROC, 0.842 Dice, and 6.9 mm HD95, ahead of the transformer and diffusion baselines tested. The ablation shows each component adds value, with diffusion-based augmentation giving the largest single gain and the conditioning modules tightening boundaries. A balanced synthetic-to-real ratio of 50% works best, cutting false positives to 0.24 per scan, while heavier synthetic use erodes accuracy. The most direct application is decision support for early screening, where stronger sensitivity may help catch tumors that routine reading misses. The work carries clear limits. Evaluation stays within the PanTS training scans, delayed-phase data is sparse, and no external clinical test is included. The method also requires accurate phase information. Future research could aim to validate across institutions, investigate sub-2 cm lesions, and determine if synthetic-ratio effect holds true in other organs and imaging applications.

References

[1] S. Mukherjee, A. Antony, N. G. Patnam, K. H. Trivedi, A. Karbhari, M. Nagaraj, M. Murlidhar, and A. H. Goenka, "Pancreas segmentation using AI developed on the largest CT dataset with multi-institutional validation and implications for early cancer detection," *Scientific Reports*, vol. 15, Art. no. 17096, 2025, doi: 10.1038/s41598-025-01802-9.

[2] S. Suneja, R. Talwar, and M. Sharma, "A systematic review on leveraging artificial intelligence for pancreatic cancer diagnosis," *Premier Journal of Science*, vol. 20, Art. no. 100268, 2026, doi: 10.70389/PJS.100268.

[3] A. C. Anameje and I. Mgbeaafulike, "Advance deep learning model for digital segmentation of pancreatic cancer," *Int. J. Real-Time Appl. Comput. Syst.*, vol. 4, no. 56, pp. 727–735, Feb. 2025.

[4] S. Nalliah *et al.*, "Deep-learning based automated pancreas segmentation on CT scans of chronic pancreatitis patients," *Eur. J. Radiol.*, vol. 189, Art. no. 112175, Aug. 2025, doi: 10.1016/j.ejrad.2025.112175.

[5] K. Dodda and G. Muneeswari, "Automatic pancreatic cancer segmentation and classification in CT images using an integrated deep-learning approach," *PeerJ Comput. Sci.*, vol. 11, Art. no. e3263, 2025, doi: 10.7717/peerj-cs.3263.

[6] R. Hussain, M. Mohsin, D. Khan, and M. Zohaib, "A systematic review of deep learning approaches for hepatopancreatic tumor segmentation," *J. Imaging*, vol. 12, no. 4, Art. no. 147, 2026, doi: 10.3390/jimaging12040147.

[7] A. Nadeem, R. Ashraf, T. Mahmood, and S. Parveen, "Automated CAD system for early detection and classification of pancreatic cancer using deep learning model," *PLOS ONE*, vol. 20, no. 1, Art. no. e0307900, 2025, doi: 10.1371/journal.pone.0307900.

[8] N. Sethi, C. V. Varma, and S. S. Reddy, "Artificial intelligence in pancreatic cancer: Applications in early detection, tumor staging, and survival prediction—a comprehensive review," *Clin. Transl. Oncol.*, online first, May 2026, doi: 10.1007/s12094-026-04381-x.

[9] A. Moglia, M. Cavicchioli, L. Mainardi, and P. Cerveri, "Deep learning for pancreas segmentation on computed tomography: A systematic review," *Artif. Intell. Rev.*, vol. 58, Art. no. 220, 2025, doi: 10.1007/s10462-024-11050-4.

[10] F. Lopez-Ramirez, E. A. Syailendra, F. Tixier, S. Kawamoto, E. K. Fishman, and L. C. Chu, "Early detection of pancreatic cancer on computed tomography: Advancements with deep learning," *Radiol. Adv.*, vol. 2, no. 5, Art. no. umaf028, Sep. 2025, doi: 10.1093/radadv/umaf028.

[11] C. Venkatesh, F. N. Al-Wesabi, H. M. Alshahrani, A. E. Yahya, A. Kiran, J. Chinna Babu, L. Sivayamini, and A. Alam, "An optimal diagnosis model for pancreatic cancer prediction based on deep learning method," *PeerJ Comput. Sci.*, vol. 12, Art. no. e3541, 2026, doi: 10.7717/peerj-cs.3541.

[12] F. N. Paramzin, V. V. Kakotkin, D. A. Burkin, Z. A. Ponimash, M. A. Nikitin, and M. A. Agapov, "Development and validation of a diagnostic model based on artificial intelligence for diagnosing pancreatic cancer," *Chin. Clin. Oncol.*, vol. 15, no. 1, Art. no. 3, 2026, doi: 10.21037/cco-25-17.

[13] Y. Bi, D. Li, R. Pang, C. Du, D. Li, X. Zhao, and H. Lv, "Diagnosis methods for pancreatic cancer with the technique of deep learning: A review and a meta-analysis," *Front. Oncol.*, vol. 15, Art. no. 1597969, 2025, doi: 10.3389/fonc.2025.1597969.

[14] J. Lu, H. Zhang, Z. Yuan, J. Yue, Q. Yao, Y. Liu, P. Jie, M. Fan, and J. Zhao, "Image-based artificial intelligence for preoperative differentiation of pancreatic cancer from pancreatitis: A systematic review and meta-analysis," *Front. Oncol.*, vol. 15, Art. no. 1660271, 2026, doi: 10.3389/fonc.2025.1660271.

[15] F. Oviedo, F. Lopez-Ramirez, F. Tixier, S. Kawamoto, A. Blanco, R. Dodhia, R. H. Hruban, W. B. Weeks, J. M. Lavista Ferres, E. K. Fishman, and L. Chu, "Benchmarking robustness of automated CT pancreas segmentation: Achieving human-level reliability through human-in-the-loop optimization," *Radiol. Adv.*, vol. 2, no. 6, Art. no. umaf040, 2025, doi: 10.1093/radadv/umaf040.

[16] L. Zhang, D. Li, T. Su, T. Xiao, and S. Zhao, "Effectiveness of radiomics-based machine learning models in differentiating pancreatitis and pancreatic ductal adenocarcinoma: Systematic review and meta-analysis," *J. Med. Internet Res.*, vol. 27, Art. no. e72420, 2025, doi: 10.2196/72420.

[17] M. Z. Genc, Y. Dalveren, G. G. M. Dalveren, A. Kara, M. Derawi, J. Kubicek, and M. Penhaker, "From simulation to clinical translation: A deep learning framework for pancreatic tumor segmentation with GUI integration," *IEEE Access*, vol. 14, pp. 26767–26783, 2026, doi: 10.1109/ACCESS.2026.3665109.

[18] D. Fang, Y. Huang, S. Li, C. Shi, J. Bao, D. Du, L. Xuan, L. Ye, Y. Zhang, C. Zhu, H. Zheng, Z. Shi, Q. Mei, and H. Wang, "A semi-supervised convolutional neural network for diagnosis of pancreatic ductal adenocarcinoma based on EUS-FNA cytological images," *BMC Cancer*, vol. 25, Art. no. 495, 2025, doi: 10.1186/s12885-025-13910-w.

[19] N. Xue *et al.*, "A fusion-based deep-learning algorithm predicts PDAC metastasis based on primary tumour CT images: A multinational study," *Gut*, vol. 74, no. 12, pp. 2024–2034, 2025, doi: 10.1136/gutjnl-2024-334237.

[20] V. Apostolou, G. Papageorgiou, and C. Tjortjis, "Artificial intelligence for early detection of pancreatic cancer: Pre-diagnostic detection across imaging, biomarkers, and EHRs: A systematic review," *Evol. Intell.*, vol. 19, Art. no. 89, 2026, doi: 10.1007/s12065-026-01202-6.