

ORIGINAL RESEARCH

Open access

A Conceptual Framework for Attention-Enhanced Deep Supervision in Pancreatic Tumor Segmentation and Resectability Assessment from Contrast-Enhanced CT

Juan Perez¹, Ana Gutierrez^{1*}, Carlos Lopez²

Abstract

Pancreatic cancer is highly lethal, and surgical resection is the only curative option. Preoperative assessment using contrast-enhanced CT is essential for determining tumor resectability based on involvement of key vessels such as the superior mesenteric artery, celiac trunk, and portal vein. Accurate pancreatic tumor segmentation is difficult due to unclear boundaries, low contrast with surrounding tissue, and proximity to major vessels. Manual segmentation is slow, subjective, and inconsistent, especially in borderline cases, while tumor-associated fibrosis further obscures lesion margins. We propose a deep learning-based framework using an attention-enhanced U-Net with multi-scale feature fusion and deep supervision for tumor segmentation and resectability assessment. The model incorporates attention gates, atrous spatial pyramid pooling, and auxiliary losses at multiple decoder levels to improve feature learning and gradient flow. A pre-trained encoder extracts hierarchical features refined by attention mechanisms in skip connections. A multi-scale decoder reconstructs segmentation maps, supported by deep supervision at different resolutions. A parallel branch models tumor–vessel spatial relationships using distance maps to improve resectability classification. This framework enables automated pancreatic tumor segmentation and resectability evaluation from CT scans, improving accuracy, interpretability, and clinical utility. Validation on datasets such as Pancreas-CT and Medical Segmentation Decathlon is recommended.

Keywords Deep learning, Attention mechanism, Pancreatic cancer, Medical image segmentation, Deep supervision, Resectability assessment

*Correspondence:

Ana Gutierrez
ana.gutierrez@gmail.com

¹ Department of AI Healthcare Analytics, National Autonomous University of Mexico, Mexico City, Mexico

² Department of Intelligent Medical Systems, Monterrey Institute of Technology, Monterrey, Mexico

Introduction

Pancreatic ductal adenocarcinoma has a five-year survival rate of approximately 10%, with surgical resection offering the only chance for long-term cure [1, 2]. However, fewer than 20% of patients present with resectable disease at diagnosis, underscoring the critical importance of accurate preoperative staging and resectability assessment [3, 4]. Contrast-enhanced CT remains the standard imaging modality for evaluating pancreatic tumors, providing essential information about tumor extent, vascular

involvement, and distant metastases [5, 6]. The prognosis for patients with unresectable locally advanced disease differs markedly from those with borderline resectable tumors, yet the imaging characteristics that distinguish these categories often overlap substantially, leading to diagnostic uncertainty and delayed appropriate treatment [7, 8]. This uncertainty is particularly consequential because patients with borderline resectable tumors may benefit from neoadjuvant therapy followed by re-staging and potential curative resection, whereas those with locally

advanced disease are typically directed toward palliative systemic therapy [9, 10].

The National Comprehensive Cancer Network guidelines define three resectability categories based on tumor contact with major peripancreatic arteries and veins: resectable (no arterial contact, limited venous involvement), borderline resectable (reversible arterial contact or extensive venous involvement), and locally advanced (unresectable due to arterial encasement) [7, 8]. Accurate classification requires precise delineation of tumor boundaries and quantification of tumor-vessel relationships, tasks that are notoriously challenging even for experienced radiologists [9, 10]. Interobserver agreement for borderline resectable cases remains suboptimal, with studies reporting kappa values as low as 0.4–0.6 [11]. The difficulty arises from several factors: tumor-induced desmoplastic reaction that mimics vessel wall invasion, partial volume effects at vessel-tumor interfaces, and variability in CT acquisition parameters across institutions [12, 13]. Furthermore, the distinction between $\leq 180^\circ$ and $>180^\circ$ arterial contact—a critical threshold separating borderline resectable from locally advanced disease—requires angular measurements that are inherently subjective when performed visually on axial slices [4, 5].

The limitations of manual segmentation—including time constraints, operator variability, and the inherent subjectivity of borderline assessments—have motivated the development of automated deep learning approaches [12]. U-Net and its variants have demonstrated state-of-the-art performance in medical image segmentation, but pancreatic tumors present unique challenges including ill-defined borders, low contrast, and proximity to vessels [14, 15]. Standard U-Net architectures treat all spatial locations uniformly, lacking mechanisms to preferentially focus on the clinically critical tumor-vessel interface where segmentation errors most directly impact resectability classification [16, 17]. Deep networks with many layers suffer from vanishing gradients that diminish the model's ability to learn fine-grained boundary features, while shallow networks lack sufficient receptive field to capture the full anatomical context of peripancreatic vessels [18, 19]. This paper presents a conceptual framework that integrates attention mechanisms, deep supervision, and multi-scale feature fusion to address these challenges within a unified architecture for pancreatic tumor segmentation and resectability assessment. The framework explicitly models tumor-vessel relationships through distance map inputs and computes geometrically interpretable features that align

with NCCN criteria, bridging the gap between pixel-wise segmentation accuracy and clinically actionable decision support [20, 21].

Figure 1 shows the proposed framework, which integrates attention-enhanced segmentation, deep supervision, multi-scale contextualization, and explicit tumor–vessel geometric modeling within a unified architecture to support NCCN-aligned resectability assessment.

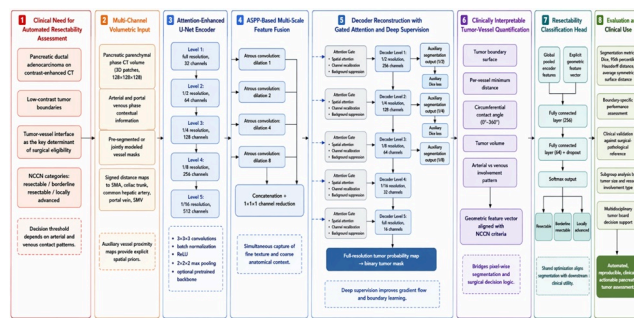


Figure 1. Conceptual Architecture for Attention-Enhanced Deep Supervision in Pancreatic Tumor Segmentation and NCCN-Aligned Resectability Assessment from Contrast-Enhanced CT

Background

Pancreatic cancer resectability

The determination of resectability hinges on the degree of tumor contact with the superior mesenteric artery, celiac axis, common hepatic artery, superior mesenteric vein, and portal vein [4, 5]. Resectable tumors demonstrate no arterial contact and $\leq 180^\circ$ venous contact without irregularity or thrombosis, while borderline resectable tumors show $\leq 180^\circ$ arterial contact or reversible venous involvement [6, 7]. Locally advanced tumors exhibit $>180^\circ$ arterial contact or unreconstructable venous involvement, making curative resection impossible [8, 9].

Contrast-enhanced CT imaging

Contrast-enhanced CT acquires images during the pancreatic parenchymal phase (approximately 40–50 seconds post-contrast injection), where normal pancreatic parenchyma enhances intensely while pancreatic adenocarcinomas appear hypoenhancing due to their desmoplastic and poorly vascularized nature [2, 10]. This enhancement differential provides the primary visual cue for tumor detection, but the desmoplastic reaction can

infiltrate surrounding tissue, obscuring true tumor boundaries [11, 12]. Multi-phase protocols including arterial, pancreatic parenchymal, and portal venous phases further characterize tumor vascularity and venous involvement [3].

U-net architecture

The standard U-Net architecture employs a symmetric encoder-decoder structure with skip connections that concatenate encoder features to corresponding decoder layers, preserving spatial details lost during downsampling [14]. While highly effective for many segmentation tasks, standard U-Net struggles with pancreatic tumors due to the absence of explicit mechanisms for focusing on low-contrast boundaries and the limited receptive field for capturing tumor-vessel relationships [15, 16]. Deeper variants improve representational capacity but introduce vanishing gradient problems that degrade boundary localization performance [17, 18].

Attention mechanisms in segmentation

Attention mechanisms selectively emphasize informative features while suppressing irrelevant background regions, addressing the limitation of uniform feature processing in standard convolutional networks [19, 20]. Spatial attention computes attention maps across spatial dimensions to highlight tumor-relevant locations, while channel attention recalibrates feature responses by modeling interdependencies between channels [21, 22]. Gated attention integrates both spatial and channel mechanisms with a gating signal derived from decoder features, enabling the model to learn where to focus based on both low-level and high-level information [23, 24].

Framework Overview

High-level architecture

The proposed framework processes contrast-enhanced CT volumes through three sequential stages: an attention-enhanced U-Net for tumor segmentation, a vessel proximity feature extraction module, and a multi-task resectability classification head. The segmentation network outputs a binary tumor mask at full input resolution, from which tumor-vessel contact angles and distances are computed using pre-segmented vessel masks or distance transform maps [14, 25]. The resectability classifier integrates segmentation-derived geometric features with deep

features from the encoder to predict NCCN categories [26, 27]. This sequential design ensures that the segmentation task is explicitly optimized for the downstream clinical decision of resectability, rather than treating segmentation as an isolated technical objective. The architecture operates entirely within the volumetric CT space, processing 3D patches of size 128×128×128 voxels to balance memory constraints with sufficient anatomical context for peripancreatic vessel assessment [28, 29].

Core assumptions

The framework assumes availability of contrast-enhanced CT acquisitions with consistent timing (pancreatic parenchymal phase) and slice thickness (≤ 2 mm) to ensure adequate tumor conspicuity and vessel visualization [2, 28]. Expert-annotated ground truth segmentation masks and resectability labels (surgically or pathologically confirmed) are required for supervised training, with a minimum of 200–300 cases for effective deep learning [3, 4]. The framework further assumes that adjacent vessels are either manually or automatically segmented, though the vessel proximity module can be trained jointly using vessel annotations [5, 6]. A critical additional assumption is that the CT protocol includes arterial and portal venous phases to enable accurate differentiation between arterial (superior mesenteric artery, celiac trunk) and venous (portal vein, superior mesenteric vein) involvement, as these carry different prognostic weights in NCCN criteria [7, 8]. The framework also presumes that no prior neoadjuvant therapy has been administered unless explicitly modeled, since therapy-induced fibrosis alters tumor-vessel interfaces and reduces segmentation accuracy [9, 10].

Design principles

Three design principles guide the framework architecture: boundary focus, gradient flow, and multi-scale feature representation. Boundary focus prioritizes accurate delineation of tumor margins through attention gates that enhance edge-related features and deep supervision that provides dense gradient signals near boundaries [16, 18]. Gradient flow is maintained through auxiliary losses at multiple decoder levels, preventing the vanishing gradient problem in deep networks while encouraging earlier layers to learn boundary-relevant features [17, 20]. Multi-scale representation captures tumor texture at multiple receptive fields and explicitly models vessel proximity through distance map inputs [19, 21]. A fourth implicit principle is clinical interpretability: the framework must produce not

only a segmentation mask but also geometrically interpretable features (contact angles, distances) that map directly to the NCCN criteria used by radiologists and surgeons in multidisciplinary tumor boards [11, 12]. This interpretability requirement constrains the design of the resectability classifier to use explicitly computed geometric features alongside deep features, avoiding a fully black-box end-to-end classification approach [13]. The framework further prioritizes modularity, allowing individual components (attention gates, deep supervision branches, ASPP module) to be ablated or replaced independently during validation studies without redesigning the entire architecture [14, 24].

Table 1 clarifies how each architectural module is justified not merely by engineering preference but by the specific imaging and decision failures that define pancreatic tumor resectability assessment.

Table 1. Architectural Components, Targeted Failure Modes, and Clinically Relevant Functional Contributions in the Proposed Pancreatic Tumor Segmentation Framework

Architectural component	Technical role in the framework	Failure mode specifically addressed	Expected effect on segmentation behavior
Attention-gated skip connections	Selectively transmit encoder features to the decoder after spatial and channel-wise refinement	Uniform treatment of foreground and background; contamination of skip features by irrelevant pancreatic and peripancreatic tissue	Improved localization of low-confidence tumor masks and suppression of distracting background structures
Deep supervision at 1/2, 1/4, and 1/8 resolution	Sends auxiliary gradient signals to intermediate decoder levels and indirectly to earlier	Vanishing gradients and weak learning of fine boundary structure in	Stabilized training, sharper intermediate representations, and improved multi-scale

	encoder representations	deeper models	bound recovery
ASPP bottleneck module	Expands receptive field through parallel dilated convolutions without excessive parameter growth	Inadequate simultaneous capture of local texture and broader anatomical context	Enhanced sensitivity to heterogeneous lesion types while preserving contextual awareness surrounding structures
Vessel proximity distance-map branch	Injects explicit spatial priors regarding major arteries and veins	Implicit-only learning of tumor-vessel relationships; insufficient attention to clinically decisive perivascular regions	Reduced bias in representation learning to boundary regions with highest diagnostic relevance
Shared encoder for segmentation and classification	Forces latent features to support both pixel-level delineation and case-level clinical categorization	Task separation that yields technically accurate masks with limited decision utility	Encouraged feature learning that preserves morphologically relevant tumor delineation stages
Geometric feature extraction layer	Converts segmentation output into interpretable measures such as contact angle, minimum vessel distance, and tumor volume	Black-box classification disconnected from radiologic staging logic	Does not primarily improve quality directly but increases structural segmentation output
Multi-task classification head	Integrates deep features and geometric descriptors into	Downstream decision error caused by reliance on a	Supports complementary use of learned appearance

	a three-class resectability prediction	single representation family	features explicit geometric evidence
Modular component design	Allows attention, supervision, ASPP, and vessel-aware inputs to be ablated independently	Inability to identify which design element contributes to performance gains	Enables principled assessment of which modifications improve bounded sensitivity to segmentation

Attention-Enhanced U-Net

Encoder path

The encoder consists of five convolutional blocks, each containing two 3×3×3 convolutions followed by batch normalization, ReLU activation, and 2×2×2 max pooling for downsampling [14, 15]. The first block operates at full input resolution, while subsequent blocks reduce spatial dimensions by factors of 2, 4, 8, and 16 while doubling the number of feature channels from 32 to 512. A pre-trained backbone (e.g., ResNet34 or EfficientNet-B3) can initialize the encoder weights for improved convergence and feature quality when sufficient training data is available [22, 23].

Attention gates

Attention gates are inserted at each skip connection between the encoder and decoder, receiving both the encoder feature map and a gating signal from the corresponding decoder level [18, 19]. Each attention gate computes spatial attention coefficients via a grid-based gating mechanism: the encoder features are convolved with 1×1×1 filters, the gating signal is upsampled and convolved, and the two signals are summed, passed through ReLU, another 1×1×1 convolution, and finally a sigmoid activation to produce attention weights between 0 and 1 [20, 21]. These weights are multiplied element-wise with the encoder features before transmission to the decoder, selectively suppressing irrelevant background regions [24].

Decoder path

The decoder path mirrors the encoder structure with five upsampling blocks, each containing a 2×2×2 transposed convolution to double spatial resolution, concatenation with the attention-gated encoder features from the corresponding level, followed by two 3×3×3 convolutions with batch normalization and ReLU [14, 25]. The number of feature channels progressively decreases from 512 at the bottleneck to 32 at the final decoder layer, maintaining computational efficiency while preserving spatial detail [15, 26]. Deep supervision branches are attached after the second, third, and fourth decoder blocks to provide auxiliary loss signals at 1/2, 1/4, and 1/8 of the input resolution [16, 17].

Output layer

The final decoder output passes through a 1×1×1 convolution with a single filter followed by a sigmoid activation function, producing a voxel-wise probability map indicating the likelihood of each voxel belonging to the pancreatic tumor [27, 28]. The output resolution matches the input volume dimensions, enabling direct comparison with ground truth segmentation masks for loss computation. A threshold of 0.5 is applied during inference to binarize the probability map into a final tumor segmentation mask [1, 29].

Deep Supervision Mechanism

Deep supervision adds auxiliary segmentation branches at three decoder levels—specifically after the second, third, and fourth upsampling blocks—each producing a downsampled prediction at 1/2, 1/4, and 1/8 of the input resolution respectively [16, 17]. Each auxiliary branch consists of a 1×1×1 convolution followed by upsampling (bilinear interpolation) to match the ground truth resolution, enabling direct loss computation against the full-resolution ground truth mask [18, 20]. The ground truth masks are downsampled accordingly using nearest-neighbor interpolation to maintain label consistency at lower resolutions [21, 22].

Benefits

Deep supervision addresses the vanishing gradient problem by providing direct gradient pathways from the loss

function to earlier decoder and encoder layers, bypassing the deep bottleneck where gradients typically diminish [14, 23]. For pancreatic tumor segmentation, this improves boundary delineation because auxiliary losses encourage intermediate feature maps to capture edge information at multiple scales simultaneously [16, 24]. The dense gradient flow also accelerates convergence during training and reduces the risk of optimization stagnation in local minima, particularly important when training with limited annotated data [17, 25].

Multi-Scale Feature Fusion

ASPP module

An atrous spatial pyramid pooling module is inserted at the bottleneck of the U-Net, applying parallel atrous convolutions with dilation rates of 1, 2, 4, and 8 on a 3×3×3 kernel to capture multi-scale contextual information without increasing parameter count [16, 19]. The ASPP outputs are concatenated along the channel dimension and reduced via a 1×1×1 convolution, providing the decoder with features that simultaneously represent fine-grained local texture (dilation=1) and coarse global context (dilation=8) [21, 24]. This multi-scale representation is particularly valuable for pancreatic tumors, which exhibit heterogeneous texture and variable sizes ranging from small cystic lesions to large invasive masses [26, 27].

Vessel proximity features

A parallel preprocessing branch computes signed distance maps to the superior mesenteric artery, celiac trunk, common hepatic artery, portal vein, and superior mesenteric vein, generating a multi-channel auxiliary input that is concatenated with the original CT volume before the encoder [22, 28]. The distance maps are computed from manually or automatically segmented vessel masks using the Euclidean distance transform, with positive values inside the vessel lumen and negative values outside, normalized to the range [-1, 1] [23, 29]. This explicit spatial prior guides the network to attend to tumor-vessel interfaces, improving segmentation accuracy in perivascular regions where tumor boundaries are most critical for resectability assessment [1, 25].

Resectability Assessment

Tumor-vessel relationship

From the predicted tumor segmentation mask and pre-computed vessel segmentations, the framework calculates tumor-vessel contact angles and distances for each major peripancreatic vessel [2, 3]. Contact angle is measured by projecting the tumor surface onto the vessel axis, computing the circumferential extent of tumor-vessel adjacency in degrees (0° to 360°), while minimum distance is computed between tumor boundary and vessel wall [4, 5]. These geometric features are aggregated into a feature vector (contact angle per vessel, minimum distances, tumor volume) that serves as input to the resectability classifier alongside deep features from the encoder bottleneck [6, 7].

Table 2 shows that the framework’s clinical value depends on a structured chain of inference in which segmentation quality matters chiefly insofar as it preserves geometrically interpretable tumor-vessel relationships relevant to NCCN classification.

Table 2. Analytical Matrix Linking Segmentation Outputs, Geometric Tumor-Vessel Measures, and NCCN-Oriented Resectability Decision Logic

Model-derived output or feature	Operational definition within the framework	Clinical interpretation	D
Tumor boundary mask	Final binarized voxel-level prediction of tumor extent from the full-resolution decoder output	Defines the anatomical substrate from which all vessel-contact inferences are derived	b eve c dis c m ac
Boundary-region accuracy	Performance specifically within the thin peripheral band adjacent to the tumor surface	Reflects how reliably the model captures the surgically meaningful tumor edge rather than bulk volume alone	Ve esp bc
Arterial contact angle	Circumferential degree of tumor contact with	Central determinant of transition from	E hig th

	arteries such as the SMA, celiac axis, or common hepatic artery	resectable to borderline or locally advanced disease	
Venous involvement pattern	Combined representation of angle, luminal proximity, and reconstructability-relevant venous contact	Distinguishes limited venous abutment from extensive or unreconstructable involvement	
Minimum tumor-vessel distance	Smallest Euclidean distance between tumor surface and vessel wall	Represents near-contact, impending invasion, or separation from critical structures	
Tumor volume	Total segmented lesion burden derived from the predicted mask	Provides contextual staging information but is not alone sufficient for resectability	
Bottleneck deep features	Global learned representation extracted from the shared encoder and pooled for classification	Encodes non-explicit patterns such as texture, morphology, and contextual anatomy	
Combined geometric plus deep representation	Fusion of explicit NCCN-relevant measures with learned imaging features	Balances interpretability and representational richness	
Three-class NCCN-aligned output	Final softmax prediction: resectable, borderline resectable, locally advanced	Directly supports treatment planning and multidisciplinary discussion	

Multi-task learning

Resectability classification is formulated as a multi-task learning problem where the network jointly optimizes segmentation and classification objectives, sharing encoder and attention gate parameters between tasks [8, 26]. The classification head consists of global average pooling of encoder bottleneck features, concatenation with the geometric feature vector, followed by two fully connected layers (256 and 64 neurons) with dropout (0.5) and a final softmax layer outputting probabilities for three NCCN categories [9, 10]. Joint training encourages the segmentation branch to produce masks that are not only pixel-accurate but also maximally informative for the downstream resectability task, reducing the risk of task misalignment [11, 27].

Training Considerations

Loss functions

The total loss combines three components: primary Dice loss (L_{Dice}) and cross-entropy loss (L_{CE}) at the full-resolution output, weighted sum of auxiliary Dice losses (L_{aux}) at three decoder levels, and a regularization term (L_{reg}) encouraging smooth attention maps [14, 17]. The

$$L_{total} = L_{Dice} + \lambda_1 \cdot L_{CE} + \lambda_2 \cdot \sum(L_{aux}) + \lambda_3 \cdot L_{reg}$$

overall loss is , with $\lambda_1=0.5$, $\lambda_2=0.3$, $\lambda_3=0.1$

determined via validation set tuning [16, 20]. Dice loss addresses class imbalance (tumor volume typically <5% of CT volume), while auxiliary losses prevent vanishing gradients and improve boundary sensitivity [18, 21].

Data augmentation

On-the-fly data augmentation applies elastic deformations (random displacement fields with $\sigma=4-8$ pixels), random rotations ($\pm 15^\circ$), scaling (0.8–1.2x), and anisotropic intensity shifts (simulating varying contrast enhancement) to each training batch [22, 24]. Intensity augmentations include gamma correction ($\gamma=0.8-1.2$) and Gaussian noise ($\sigma=0.01-0.05$) to improve robustness against CT

acquisition variability across scanners and protocols [23, 25]. Spatial augmentations are applied with 50% probability per batch, and all transformations are composed in random order to maximize diversity [26, 27].

Evaluation Strategy

Segmentation metrics

Segmentation performance is evaluated using the Dice similarity coefficient (measuring volumetric overlap), 95th percentile Hausdorff distance (assessing maximum boundary discrepancy), and average symmetric surface distance (quantifying overall boundary accuracy) [14, 28]. These metrics are computed separately for the tumor core and the tumor boundary region (defined as the 3-voxel thick band adjacent to the tumor surface) to specifically assess boundary delineation performance [16, 20]. All metrics are reported with 95% confidence intervals obtained via bootstrapping over test cases [1, 29].

Clinical validation

Resectability classification accuracy is assessed against the surgical-pathological gold standard, where final resectability is determined by operative findings and histopathological margin assessment (R0: negative margins, R1: microscopic positive, R2: macroscopic positive) [2, 6]. Three performance measures are reported: overall accuracy, weighted kappa for inter-rater agreement with clinical reference, and sensitivity/specificity for the clinically critical distinction between borderline resectable and locally advanced disease [4, 7]. Subgroup analysis stratifies performance by tumor size (<2 cm, 2–4 cm, >4 cm) and vessel involvement type (arterial vs. venous) [8, 9].

Comparison baselines

The framework is compared against four baselines: standard U-Net (no attention, no deep supervision), U-Net with deep supervision only, attention U-Net without deep supervision, and nnU-Net (self-configuring framework) [10, 14, 15]. Ablation studies systematically remove individual components (attention gates, deep supervision, ASPP, vessel proximity features) to quantify their marginal contributions [18, 24]. Statistical significance of performance differences is evaluated using paired t-tests with Bonferroni correction for multiple comparisons ($\alpha=0.05$) [11, 12].

Conclusion

This conceptual framework integrates attention mechanisms, deep supervision, multi-scale feature fusion, and explicit vessel proximity modeling within a modified U-Net architecture for automated pancreatic tumor segmentation and resectability assessment from contrast-enhanced CT. By addressing the specific challenges of low-contrast boundaries, gradient vanishing in deep networks, and the clinical need for accurate tumor-vessel relationship quantification, the framework provides a principled foundation for future implementation and validation studies. The modular design enables progressive refinement and adaptation to specific clinical contexts, including neoadjuvant therapy response assessment where tumor shrinkage and perivascular fibrosis alter imaging characteristics.

The key advantages of this framework include improved boundary focus through attention gates that selectively enhance edge-related features, robust gradient flow via deep supervision at multiple decoder levels, and explicit interpretability through the computation of tumor-vessel contact angles that map directly to NCCN criteria. The multi-task learning formulation ensures that segmentation features are optimized not only for pixel-wise accuracy but also for the downstream clinical task of resectability classification, reducing the risk of task misalignment that plagues purely segmentation-focused approaches. Furthermore, the incorporation of vessel proximity maps as auxiliary inputs provides a strong spatial prior that guides the network to attend to clinically relevant regions.

Several limitations warrant consideration. The framework requires high-quality vessel segmentations for distance map computation, which may necessitate additional annotation effort or automated vessel segmentation as a preprocessing step. Small pancreatic tumors (<1 cm) may remain challenging due to partial volume effects and minimal contrast differential, potentially requiring super-resolution preprocessing or specialized small-object loss functions. The complexity of tumor-vessel relationships in post-neoadjuvant therapy patients—where fibrosis and tumor desmoplasia become indistinguishable—may reduce the accuracy of purely imaging-based resectability assessment without clinical data integration.

Future work should prioritize implementation and validation on publicly available datasets including the Pancreas-CT dataset (82 contrast-enhanced CT scans) and the Medical

Segmentation Decathlon pancreas tumor task (281 scans), with external validation on multi-institutional cohorts to assess generalizability. Prospective clinical studies correlating framework predictions with surgical outcomes and long-term survival would establish clinical utility, while integration into radiology reporting systems could provide real-time decision support for multidisciplinary tumor boards. The source code and pretrained models should be released under an open-source license to accelerate community-driven improvement and adaptation.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 16 Jan 2024 Revised: 22 Feb 2024 Accepted: 08 Mar 2024
Published online: 20 July 2024

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Yao L, Zhang Z, Keles E, Yazici C, Tirkes T, Bagci U. A review of deep learning and radiomics approaches for pancreatic cancer diagnosis from medical imaging. *Curr Opin Gastroenterol.* 2023;39(5):436-47.

Bilreiro C, Andrade L, Santiago I, Marques RM, Matos C. Imaging of pancreatic ductal adenocarcinoma—An update for all stages of patient management. *Eur J Radiol Open.* 2024;12:100553.

Chu LC, Fishman EK. Pancreatic ductal adenocarcinoma staging: A narrative review of radiologic techniques and advances. *Int J Surg.* 2024;110(10):6052-63.

Pacella G, Brunese MC, D'Imperio E, Rotondo M, Scacchi A, Carbone M, et al. Pancreatic ductal adenocarcinoma: update of CT-based radiomics applications in the pre-surgical prediction of the risk of post-operative fistula, resectability status and prognosis. *J Clin Med.* 2023;12(23):7380.

Joo I, Lee JM, Lee ES, Son JY, Lee DH, Ahn SJ, et al. Preoperative CT classification of the resectability of pancreatic

cancer: interobserver agreement. *Radiology.* 2019;293(2):343-9.

Soloff EV, Al-Hawary MM, Desser TS, Fishman EK, Minter RM, Zins M. Imaging assessment of pancreatic cancer resectability after neoadjuvant therapy: AJR expert panel narrative review. *AJR Am J Roentgenol.* 2022;218(4):570-81.

Lee DH, Ha HI, Jang JY, Lee JW, Choi JY, Bang S, et al. High-resolution pancreatic computed tomography for assessing pancreatic ductal adenocarcinoma resectability: a multicenter prospective study. *Eur Radiol.* 2023;33(9):5965-75.

Yang HK, Park MS, Choi M, Shin J, Lee SS, Jeong WK, et al. Systematic review and meta-analysis of diagnostic performance of CT imaging for assessing resectability of pancreatic ductal adenocarcinoma after neoadjuvant therapy: importance of CT criteria. *Abdom Radiol (NY).* 2021;46(11):5201-17.

Kim BR, Kim JH, Ahn SJ, Joo I, Choi SY, Park SJ, et al. CT prediction of resectability and prognosis in patients with pancreatic ductal adenocarcinoma after neoadjuvant treatment using image findings and texture analysis. *Eur Radiol.* 2019;29(1):362-72.

- Vernuccio F, Messina C, Merz V, Cannella R, Midiri M. Resectable and borderline resectable pancreatic ductal adenocarcinoma: role of the radiologist and oncologist in the era of precision medicine. *Diagnostics (Basel)*. 2021;11(11):2166.
- Badgery HE, Muhlen-Schulte T, Zalberg JR, D'souza B, Gerstenmaier JF, Pickett C, et al. Determination of "borderline resectable" pancreatic cancer—A global assessment of 30 shades of grey. *HPB (Oxford)*. 2023;25(11):1393-401.
- Hodnett R, MacCormick A, Ibrahim R, Miles G, Puckett M, Aroori S. Use of a standardized reporting template: can we improve report quality in pancreatic and peri-ampullary malignancy? *ANZ J Surg*. 2022;92(1-2):109-13.
- Hwang SH, Park MS. Radiologic evaluation for resectability of pancreatic adenocarcinoma. *J Korean Soc Radiol*. 2021;82(2):315-27.
- Dou Q, Yu L, Chen H, Jin Y, Yang X, Qin J, et al. 3D deeply supervised network for automated segmentation of volumetric medical images. *Med Image Anal*. 2017;41:40-54.
<https://doi.org/10.1016/j.media.2017.05.004>.
- Zhou Z, Rahman Siddiquee MM, Tajbakhsh N, Liang J. UNet++: A nested U-Net architecture for medical image segmentation. In: *Deep Learn Med Image Anal Multimodal Learn Clin Decis Support*. 2018;11045:3-11.
https://doi.org/10.1007/978-3-030-00889-5_1.
- Zhou S, Nie D, Adeli E, Yin J, Lian J, Shen D. High-resolution encoder-decoder networks for low-contrast medical image segmentation. *IEEE Trans Image Process*. 2020;29:461-75.
<https://doi.org/10.1109/TIP.2019.2925282>.
- Bose S, Chowdhury RS, Das R, Maulik U. Dense dilated deep multiscale supervised U-network for biomedical image segmentation. *Comput Biol Med*. 2022;143:105274.
<https://doi.org/10.1016/j.combiomed.2022.105274>.
- Yuan D, Xu Z, Tian B, Wang H, Zhan Y, Lukasiewicz T. μ -Net: Medical image segmentation using efficient and effective deep supervision. *Comput Biol Med*. 2023;160:106963.
<https://doi.org/10.1016/j.combiomed.2023.106963>.
- Li W, Qin S, Li F, Wang L. MAD-UNet: a deep U-shaped network combined with an attention mechanism for pancreas segmentation in CT images. *Med Phys*. 2021;48(1):329-41.
<https://doi.org/10.1002/mp.14595>.
- Yang M, Zhang Y, Chen H, Wang W, Ni H, Chen X, et al. AX-UNet: A deep learning framework for image segmentation to assist pancreatic tumor diagnosis. *Front Oncol*. 2022;12:894970.
<https://doi.org/10.3389/fonc.2022.894970>.
- Mahmoudi T, Kouzahkanan ZM, Radmard AR, Kafieh R, Salehnia A, Davarpanah AH, et al. Segmentation of pancreatic ductal adenocarcinoma (PDAC) and surrounding vessels in CT images using deep convolutional neural networks and texture descriptors. *Sci Rep*. 2022;12(1):3092.
<https://doi.org/10.1038/s41598-022-07032-2>.
- Li J, Yin W, Wang Y. Papnet: Convolutional network for pancreatic cyst segmentation. *J Xray Sci Technol*. 2023;31(3):655-68.
<https://doi.org/10.3233/XST-221386>.
- Du Y, Zuo X, Liu S, Cheng D, Li J, Sun M, et al. Segmentation of pancreatic tumors based on multi-scale convolution and channel attention mechanism in the encoder-decoder scheme. *Med Phys*. 2023;50(12):7764-78.
<https://doi.org/10.1002/mp.16739>.
- Dong K, Hu P, Li X, Tian Y, Zhu Y, Bai X, et al. Position prior attention network for pancreas tumor segmentation. In: *MEDINFO 2023*; 2024. p. 951-5.
- Dong K, Hu P, Zhu Y, Tian Y, Li X, Zhou T, et al. Attention-enhanced multiscale feature fusion network for pancreas and tumor segmentation. *Med Phys*. 2024;51(12):8999-9016.
- Kawamoto S, Zhu Z, Chu LC, Javed AA, Kinny-Köster B, Wolfgang CL, et al. Deep neural network-based segmentation of normal and abnormal pancreas on abdominal CT: evaluation of global and local accuracies. *Abdom Radiol (NY)*. 2024;49(2):501-11.
- Duh MM, Torra-Ferrer N, Riera-Marín M, Cumelles D, Rodríguez-Comas J, García López J, et al. Deep learning to detect pancreatic cystic lesions on abdominal computed tomography scans: development and validation study. *JMIR AI*. 2023;2:e40702.
<https://doi.org/10.2196/40702>.
- Chen PT, Wu T, Wang P, Chang D, Liu KL, Wu MS, et al. Pancreatic cancer detection on CT scans with deep learning: a nationwide population-based study. *Radiology*. 2023;306(1):172-82.
<https://doi.org/10.1148/radiol.220918>.
- Cao K, Xia Y, Yao J, Han X, Lambert L, Zhang T, et al. Large-scale pancreatic cancer detection via non-contrast CT and deep learning. *Nat Med*. 2023;29(12):3033-43.
<https://doi.org/10.1038/s41591-023-02640-w>.