

# Multi-Head Automated Segmentation by Incorporating Detection Head into the Contextual Layer Neural Network

Edwin Kys<sup>1</sup> and Febian Febian<sup>2\*</sup>

<sup>1</sup>Researcher Linnear Austin, Texas

<sup>2</sup>Researcher UCL / Linnear London, United Kingdom

\*Corresponding Author

Febian Febian, Researcher UCL / Linnear London, United Kingdom.

Submitted: 2026, Jan 09; Accepted: 2026, Feb 16; Published: 2026, Feb 19

**Citation:** Kys, E., Febian, F. (2026). Multi-Head Automated Segmentation by Incorporating Detection Head into the Contextual Layer Neural Network. *OA J Applied Sci Technol*, 4(1), 01-07.

## Abstract

Deep-learning-based auto-segmentation is increasingly used in radiotherapy, but conventional models often produce anatomically implausible false positives, or “hallucinations,” in slices lacking target structures. We propose a gated multi-head Transformer architecture based on Swin U-Net, augmented with inter-slice context integration and a parallel detection head, which jointly performs slice-level structure detection via a multi-layer perceptron and pixel-level segmentation through a context-enhanced stream. Detection outputs gate the segmentation predictions to suppress false positives in anatomically invalid slices, and training uses slice-wise Tversky loss to address class imbalance. Experiments on the ProstateAnatomical-Edge-Cases dataset from The Cancer Imaging Archive demonstrate that the gated model substantially outperforms a non-gated segmentation-only baseline, achieving a mean Dice loss of  $0.013 \pm 0.036$  versus  $0.732 \pm 0.314$ , with detection probabilities strongly correlated with anatomical presence, effectively eliminating spurious segmentations. In contrast, the non-gated model exhibited higher variability and persistent false positives across all slices. These results indicate that detection-based gating enhances robustness and anatomical plausibility in automated segmentation applications, reducing hallucinated predictions without compromising segmentation quality in valid slices, and offers a promising approach for improving the reliability of clinical radiotherapy auto-contouring workflows.

**Keywords:** Contouring, Auto-Contouring, Auto-Segmentation, CT Scan, Imaging Modality, Contextual Layer, Neural Network, Deep Learning, Machine Learning, Swin U-Net, Transformer

## 1. Introduction

Computed tomography (CT) imaging plays a central role in modern clinical workflows, particularly in radiotherapy treatment planning, where accurate delineation of target volumes and organs-at-risk (OARs) is essential for optimal dose delivery and minimization of normal tissue complications. Manual contouring or segmentation of anatomical structures on CT is widely regarded as the clinical standard; however, it remains a time-consuming, labor-intensive process that is subject to both inter- and intra-observer variability [1]. Deep learning-based auto-contouring seeks to address these limitations by providing rapid, reproducible segmentation of structures directly from CT volumes.

Auto-segmentation methods in medical imaging have evolved substantially over the past decade. Traditional approaches, such

as atlas-based registration and deformable models, laid early groundwork for algorithmic contouring but were often challenged by the considerable anatomical variability, imaging artifacts, and limited soft-tissue contrast inherent in CT imaging [2]. The advent of convolutional neural networks (CNNs) and related deep learning architectures ushered in significant performance gains, with models now routinely achieving expert-level segmentation accuracy in many clinical contexts. A particularly influential architecture in this domain is the U-Net, which couples an encoder-decoder structure with skip connections to capture both global context and fine spatial details for segmentation tasks [3]. A study by Ng et al.(2022) illustrates significant performance gains of deep learning-based model compared to the atlas-based one. In a retrospective comparison involving 45 head and neck radiotherapy patients, the deep learning-based segmentation

---

method consistently outperformed the atlas-based approach by achieving median Dice similarity coefficients above 0.80 for key organs such as the brainstem, mandible, eyes, and spinal cord, while also exhibiting less variability and requiring less time for contouring than the atlas-based method [4].

Deep learning-based auto-segmentation has been applied across a range of CT-based radiotherapy scenarios, demonstrating substantial reductions in contouring time and variability when compared with manual delineation. For example, reviews of such methods have shown that deep learning auto-contours can achieve accuracy comparable to manual expert contours across multiple sites and structures while drastically reducing user time [5]. Another study by Zhu et al. (2018) suggests a 3D squeeze-and-excitation U-Net architecture trained on 261 head and neck CT volumes, and demonstrated that it improved average Dice similarity coefficients by approximately 3.3% compared with the previous MICCAI 2015 head-and-neck segmentation benchmark while processing an entire organ-at-risk set in roughly 0.12 seconds in a single forward pass [6]. These methods also show promise in reducing inter-observer variability, one of the key bottlenecks in clinical contouring workflows, and in maintaining consistency when applied to independent datasets [7,8]. Despite these advances, several challenges remain. CT images are inherently limited in soft-tissue contrast, which can confound boundary definition for certain organs or targets [9]. Additionally, generalization across imaging protocols, institutions, and patient populations remains a barrier to broad clinical adoption [10]. Many deep learning approaches also require large volumes of high-quality annotated data, which are costly to curate and may not be uniformly available across all anatomical sites or treatment settings [11]. Robust uncertainty estimation and clinically acceptable performance thresholds are further areas of active investigation.

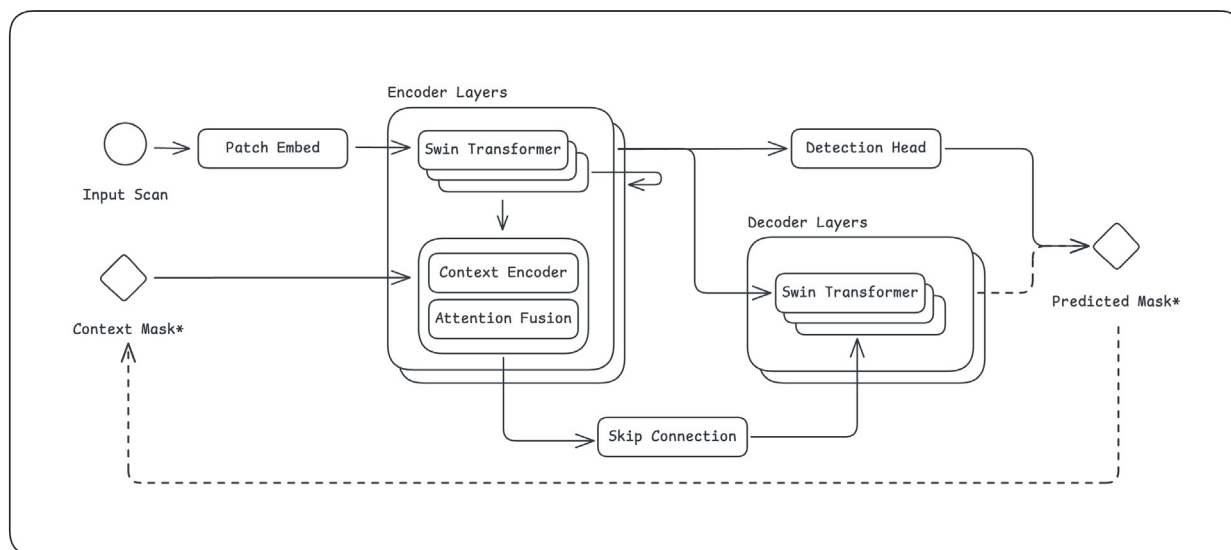
In this work, we propose a model that incorporates contextual layer neural network integrated with detection head for auto-contouring on CT images. The proposed approach is designed to reduce hallucination, hence, improving the volumetric dice performance thereby improving contour quality and reducing the amount of manual correction required by clinicians. The work is trained and evaluated on a public dataset provided by TCIA namely Prostate-Anatomical-Edge-Cases collection and compare against established baselines using standard quantitative metrics and qualitative expert review.

## 2. Materials and Methods

Medical image segmentation has seen substantial advances with the adoption of deep learning, particularly CNN and, more recently, Transformer-based architectures. Models such as U-Net and its numerous variants have become the de facto standard due to their strong performance in pixel-wise prediction tasks. However, a persistent challenge in volumetric and slice-based medical imaging is the phenomenon commonly referred to as hallucination, where models predict anatomically implausible structures in image slices where those structures do not exist. This issue is especially pronounced in organs with limited spatial extent along a given axis, such as lungs near the apex or base, small tumours, or vessels that appear intermittently across slices [12]. Hallucination arises primarily because traditional segmentation models are optimized solely for pixel-level accuracy, without an explicit mechanism to reason about structure presence or absence at the slice-level. As a result, models may extrapolate learned spatial patterns beyond their valid anatomical context, leading to false positive segmentations that are clinically misleading [13].

The emerging transformer-based models, including Swin U-Net variants, have been introduced to address limitations of CNNs in capturing long-range dependencies. Hierarchical Transformers leverage self-attention mechanisms to model global context while maintaining computational efficiency through windowed attention [14]. In medical imaging, these architectures have demonstrated improved boundary delineation and robustness to noise. Despite these improvements, Transformers alone do not inherently solve the hallucination problem. While temporal or inter-slice context can reduce inconsistency as presented in a study by An (2025) excessive reliance on contextual cues may actually exacerbate hallucinations by encouraging the model to propagate structures across slices where they are anatomically absent [15]. This limitation motivates architectures that explicitly separate structure existence reasoning from pixel-wise segmentation.

Multi-task learning has been widely explored in medical imaging as a means to improve generalization by sharing representations across related tasks. Recent work has explored architectures that combine detection and segmentation sequentially. Detection-segmentation pipelines often use detection outputs to define regions of interest for segmentation, thereby reducing false positives [16]. However, such pipelines are typically staged and do not benefit from fully shared representations.



**Figure 1:** Overview of N2 Model Workflow Consisting of Patch Embedding, Encoder Layers with Context Integration, Skip Connections to the Decoder, and Temporal Context Fusion with the Addition of Detection Head in its Pipeline

Our proposed model, N2 architecture as shown in Figure 1, builds on this paradigm by introducing dual processing streams in parallel: a context-free detection stream and a context-enhanced segmentation stream. The segmentation stream is based on Swin U-Net and enhanced with temporal context integration. Through a cross-attention mechanism, N2 incorporates previous segmentation masks into the encoder, which strengthens its ability to handle inter-frame variations and promotes consistent segmentations over time. This design choice highlights segmentation benefits from temporal or contextual fusion whereas detection should rely primarily on the current image content to avoid biased presence decisions driven by neighbouring slices. In this model, detection outputs can be used to gate or modulate segmentation predictions. Soft gating strategies scale segmentation probabilities by detection confidence, while hard gating applies binary masks based on detection thresholds. These mechanisms are especially effective for suppressing hallucinated segmentations without degrading performance on slices where the structure truly exists which will be presented in more detail in later section.

The N2 workflow starts with an initial patch-embedding layer that converts the raw input tensor of shape  $(B, CS, H, W)$ , where  $B$  is batch size,  $CS$  is the number of channels, and  $H, W$  are spatial dimensions, into a structured sequence of tokens by partitioning the image into non-overlapping patches of size  $(PH, PW)$ , flattening them, and projecting them into a  $D$ -dimensional embedding space using a linear or convolutional projection. This produces a feature

tensor of shape  $(B, N, D)$  that enables efficient capture of local spatial patterns while reducing the computational cost compared to processing the whole image. The encoder then builds hierarchical feature representations by passing these patch tokens through multiple stages of Swin Transformer blocks that apply window-based and shifted window self-attention, interleaved with patch-merging layers that downsample spatial resolution and increase channel depth, all while supporting multi-scale feature extraction. Skip connections link encoder and decoder stages to preserve fine spatial details.

Simultaneously, the output features from the encoder are routed into a multi-head processing module: one stream feeds the segmentation decoder, while a parallel detection stream leads into a detection head implemented as a multi-layer perceptron (MLP) classifier to predict structure presence or absence from the encoded representations. In the decoder, patch-expanding layers upsample the features progressively, merging them with corresponding encoder outputs via skip connections to restore spatial resolution and recover both contextual and fine structural information effectively. This dual-head design enables the model to both delineate anatomical structures and determine their existence within a given input, improving overall reliability. Finally, normalized features are linearly projected to per-class segmentation logits, reshaped into spatial grids and upsampled back to the original resolution to produce dense, pixel-level segmentation outputs.

Stage	Description	Input Shape	Output Shape	Notes
1	Patch Embedding	$(8,1,256,256)$	$(8,4096,96)$	Patch Size $(4,4)$ , $D = 96$
2.1	Encoder Stage 1	$(8,4096,96)$	$(8,1024,192)$	Channels=192, Tokens=1024
	Temporal Context Fusion	$(8,1,256,256)$	$(8,1024,192)$	Context Encoder

		+ (8,1024,192)		+ Attention Fusion
2.2	Encoder Stage 2	(8,1024,192)	(8,256,384)	Channels=384, Tokens=256
	Temporal Context Fusion	(8,1,256,256)	(8,256,384)	Context Encoder
		+ (8,256,384)		+ Attention Fusion
2.3	Encoder Stage 3	(8,256,384)	(8,64,768)	Channels=768, Tokens=64
	Temporal Context Fusion	(8,1,256,256)	(8,64,768)	Context Encoder
		+ (8,64,768)		+ Attention Fusion
2.4	Encoder Stage 4	(8,64,768)	(8,64,768)	Channels=768, Tokens=64
	Temporal Context Fusion	(8,1,256,256)	(8,64,768)	Context Encoder
		+ (8,64,768)		+ Attention Fusion
3.1	Detection Head	(8,64,768)	(8,3)	ROI Channel=3
3.1	Decoder Stage 4	(8,64,768)	(8,256,384)	Channels=384, Tokens=256
3.2	Decoder Stage 3	(8,256,384)	(8,1024,192)	Channels=192, Tokens=1024
3.3	Decoder Stage 2	(8,1024,192)	(8,4096,96)	Channels=96, Tokens=4096
3.4	Decoder Stage 1	(8,4096,96)	(8,4096,96)	Channels=96, Tokens=4096
4.1	Normalization	(8,4096,96)	(8,4096,96)	LayerNorm
4.2	Linear Projection	(8,4096,96)	(8,4096,1)	Segmentation Classes=1
4.3	Reshape	(8,4096,1)	(8,1,64,64)	Tokens → Grid
4.4	Upsample	(8,1,64,64)	(8,1,256,256)	

**Table 1: End-to-End Example Flow of the N2 Architecture Input and Output at Each Stage**

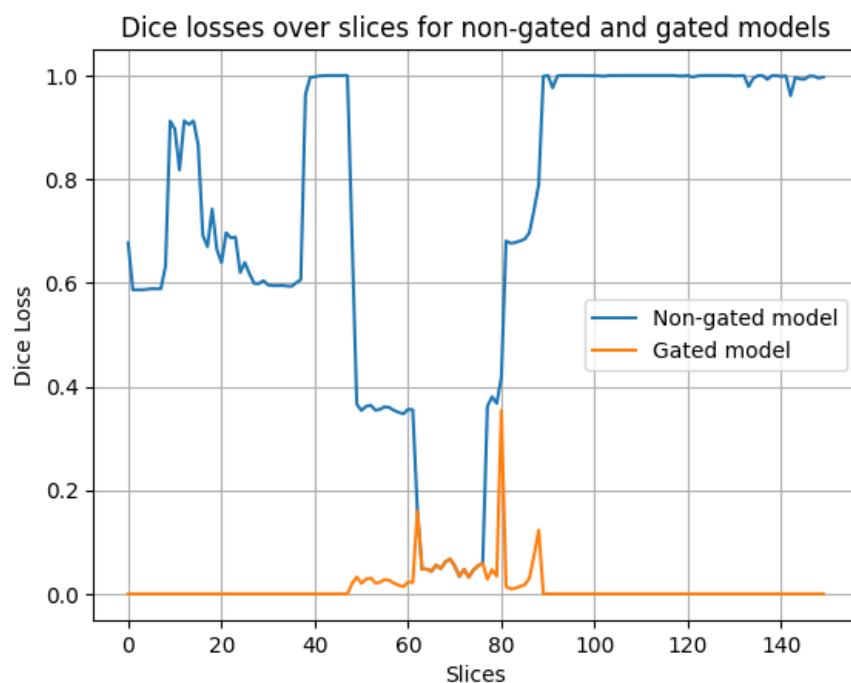
In this work, we utilized the publicly accessible Prostate-Anatomical-Edge-Cases dataset from The Cancer Imaging Archive (TCIA), which includes CT scans with accompanying RTSTRUCT segmentations from 131 individuals diagnosed with prostate cancer, totaling roughly 17 GB of data [17]. The collection originates from a single-institution, retrospective review of 950 consecutive prostate adenocarcinoma patients treated with definitive radiotherapy between 2011 and 2019, from which 112 cases exhibiting notable anatomical variations—such as hip arthroplasty, enlarged median prostate lobes, “droopy” seminal vesicles, and the presence of urinary catheters—were identified as edge cases, alongside 19 randomly selected normal cases for comparison. For every subject, expert manual segmentations were provided for the prostate, rectum, bladder, and bilateral femoral heads, establishing a robust ground truth for treatment planning and algorithm evaluation.

Medical image segmentation often suffers from severe class imbalance, particularly when structures occupy a small fraction of the image. To train the proposed model, we used Tversky loss, a generalization of overlap-based metrics that introduces tunable weights to explicitly balance penalties for false positives and false negatives, thereby enhancing boundary precision making it well-suited for segmentation tasks where missing a structure is more costly than slight over-segmentation [18]. The loss was computed on a slice-by-slice basis so that each 2D slice contributed equally

to training, preventing larger structures from disproportionately influencing the objective.

To improve model generalization and robustness, we applied a suite of preprocessing and data augmentation steps during training. Image intensities were first standardized using Min–Max scaling, and spatial variability was addressed through random 90° rotations and horizontal and vertical flips. Elastic deformations were incorporated to mimic realistic anatomical variability, while photometric variations such as brightness and contrast perturbations, along with controlled Gaussian noise injection, were used to simulate diverse acquisition conditions. Finally, all image volumes were normalized to zero mean and unit variance. This combination of geometric, photometric, and noise augmentations created a varied training distribution while preserving the anatomical fidelity necessary for accurate medical image segmentation.

The model was trained for 100 epochs with a batch size of 8 using the AdamW optimizer, a variant of Adam that separates weight decay from the gradient update to improve regularization, with an initial learning rate of  $1 \times 10^{-5}$ . The network has about 110 million trainable parameters and was implemented in PyTorch. Training was performed on a single NVIDIA RTX 4000 Ada GPU with 20 GB of VRAM, supported by a host machine equipped with 50 GB of RAM and 9 VCPUs.



**Figure 2:** Dice Loss Comparison Between the Gated Multi-Head Model and the Non-Gated Segmentation-only Model

### 3. Results and Discussion


Quantitative evaluation was conducted using slice-wise Dice loss and detection probability outputs to compare the proposed gated multi-head architecture with the corresponding non-gated segmentation-only baseline. A total of 150 axial slices were analyzed for each model.

The gated multi-head model demonstrated a substantial improvement in segmentation performance, exhibiting a markedly lower mean volumetric Dice loss of  $0.013 \pm 0.036$  compared with  $0.732 \pm 0.314$  for the non-gated configuration. In addition to the lower average loss, the gated model showed significantly reduced variability across slices, indicating more stable and consistent performance. By contrast, the non-gated model displayed pronounced fluctuations in Dice loss across all slices, particularly in slices near anatomical boundaries, reflecting an increased incidence of false positive segmentations. These trends are quantitatively summarized in Figure 2.

Analysis of the detection probability outputs further supports these findings. In the gated architecture, the detection head consistently assigned near-zero confidence scores to slices where the target anatomy was absent. This behaviour effectively suppressed segmentation predictions via the gating mechanism, resulting

in near-zero Dice loss in non-anatomical slices. In contrast, the non-gated model produced segmentation outputs regardless of anatomical presence, leading to persistently elevated Dice loss values in slices where no structure should be present. This behaviour highlights the inherent limitation of segmentation-only architectures in handling anatomically absent regions.

Qualitative results across continuous slices are illustrated in Figure 3. In the first row, the non-gated segmentation-only model exhibits clear hallucinated segmentations, whereas the gated multi-head model correctly suppresses segmentation, in agreement with the ground truth. Across subsequent rows, hallucinations persist in the non-gated model, including spurious bladder segmentations that propagate across sequential slices. These false positives emphasize the tendency of segmentation-only models to produce anatomically implausible predictions when contextual awareness is lacking. When the bladder structure is finally present in the ground truth, the non-gated model achieves segmentation quality comparable to the gated approach. Refer back to the Figure 2, the dice loss of non-gated model overlaps with the gated model showcasing the effectiveness of the baseline model architecture relying on contextual layer integrated Swin U-Net model. However, this occurs only after a prolonged sequence of erroneous predictions.



**Figure 3:** Comparison of Multi-Channel Segmentation Results for Prostate, Bladder, and Rectum Across Sequential Slices (45, 60, 61, and 62) from a Representative Case. Ground Truth Annotations are Shown in the Left Column, Predictions from the Gated Multi-Head Model in the Middle Column, and Predictions from the Segmentation-only Model in the Right Column

Representative multi-channel segmentation results for prostate, bladder, and rectum are shown in Figure 3. The gated multi-head model consistently aligns with the ground truth across both anatomical and non-anatomical slices, while the segmentation-only model demonstrates false positives in regions where structures are absent. Together, these qualitative and quantitative findings indicate that detection-based gating provides a clear advantage, yielding lower Dice loss, reduced variance, and improved anatomical plausibility. These improvements are particularly relevant for clinical radiotherapy applications involving structures with limited spatial extent, where conventional non-gated segmentation models are especially prone to hallucinated predictions.

#### 4. Conclusion

This work demonstrates that hallucination in medical image segmentation is not merely a consequence of data limitations or optimization challenges, but a fundamental structural limitation

of models trained exclusively for pixel-level accuracy without explicit modeling of anatomical presence. By introducing a dual-head architecture that decouples slice-level structure detection from pixel-wise segmentation and incorporates detection-based gating, the proposed approach effectively suppresses false positive predictions in anatomically invalid slices.

The non-gated segmentation-only model produced persistent false positive predictions in anatomically absent regions, resulting in a high mean volumetric Dice loss of  $0.732 \pm 0.314$ . In contrast, incorporating a multi-head architecture that combines a slice-level detection component based on an MLP with a contextual Swin U-Net segmentation backbone effectively mitigates this hallucination behaviour. The proposed multi-head model achieves a substantially lower mean volumetric Dice loss of  $0.013 \pm 0.036$ , along with markedly reduced variance across slices. This numerical improvement is particularly pronounced in boundary

and edge-case slices, where hallucinated predictions dominate the error profile of the non-gated baseline. The detection head reliably assigns near-zero confidence to slices lacking target anatomy, preventing the propagation of erroneous segmentations while preserving accurate predictions when structures are present.

Overall, these results provide strong quantitative evidence that detection-aware, multi-head architectures significantly enhance anatomical plausibility and robustness in deep learning-based auto-segmentation. By addressing hallucination at a structural level, rather than relying solely on pixel-wise optimization, the proposed framework offers a promising and clinically meaningful advancement toward safer and more reliable radiotherapy segmentation workflows. Future work will focus on extending this approach to additional anatomical sites, validating generalizability across institutions, and exploring uncertainty-aware gating mechanisms to further improve clinical safety.

## References

1. Bibault, J. E., & Giraud, P. (2023). Deep learning for automated segmentation in radiotherapy: a narrative review. *British Journal of Radiology*, 97(1153), 13-20.
2. Wachinger, C., Fritscher, K., Sharp, G., & Golland, P. (2015). Contour-driven atlas-based segmentation. *IEEE transactions on medical imaging*, 34(12), 2492-2505.
3. Ronneberger, O., Fischer, P., & Brox, T. (2015, October). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (pp. 234-241). Cham: Springer international publishing.
4. Ng, C. K., Leung, V. W., & Hung, R. H. (2022). Clinical evaluation of deep learning and atlas-based auto-contouring for head and neck radiation therapy. *Applied sciences*, 12(22), 11681.
5. Moran, K., Poole, C., & Barrett, S. (2025). Evaluating deep learning auto-contouring for lung radiation therapy: A review of accuracy, variability, efficiency and dose, in target volumes and organs at risk. *Physics and Imaging in Radiation Oncology*, 100736.
6. Zhu, W., Huang, Y., Zeng, L., Chen, X., Liu, Y., Qian, Z., ... & Xie, X. (2018). AnatomyNet: deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy. *Medical physics*, 46(2), 576-589.
7. Wong, J., Fong, A., McVicar, N., Smith, S., Giambattista, J., Wells, D., ... & Alexander, A. (2020). Comparing deep learning-based auto-segmentation of organs at risk and clinical target volumes to expert inter-observer variability in radiotherapy planning. *Radiotherapy and Oncology*, 144, 152-158.
8. Arjmandi, N., Mosleh-Shirazi, M. A., Mohebbi, S., Nasser, S., Mehdizadeh, A., Pishevar, Z., ... & Momennezhad, M. (2025). Evaluating the dosimetric impact of deep-learning-based auto-segmentation in prostate cancer radiotherapy: Insights into real-world clinical implementation and inter-observer variability. *Journal of Applied Clinical Medical Physics*, 26(3), e14569.
9. Weiss, E., Wu, J., Sleeman, W., Bryant, J., Mitra, P., Myers, M., ... & Williamson, J. (2010). Clinical evaluation of soft tissue organ boundary visualization on cone-beam computed tomographic imaging. *International Journal of Radiation Oncology\* Biology\* Physics*, 78(3), 929-936.
10. Suleman, M. U., Mursaleen, M., Khalil, U., Saboor, A., Bilal, M., Khan, S. A., ... & Tahir, M. (2025). Assessing the generalizability of artificial intelligence in radiology: a systematic review of performance across different clinical settings. *Annals of Medicine and Surgery*, 87(12), 8803-8811.
11. Wang, S., Li, C., Wang, R., Liu, Z., Wang, M., Tan, H., ... & Zheng, H. (2021). Annotation-efficient deep learning for automatic medical image segmentation. *Nature communications*, 12(1), 5915.
12. Zhou, J., Fejza, K., Salvatori, M., Della Latta, D., Hermann, G. M., & Di Fulvio, A. (2024). Deep-learning segmentation of small volumes in ct images for radiotherapy treatment planning. *arXiv preprint arXiv:2404.04202*.
13. Rickmann, A. M., Xu, M., Wolf, T. N., Kovalenko, O., & Wachinger, C. (2023, June). Halos: hallucination-free organ segmentation after organ resection surgery. In *International Conference on Information Processing in Medical Imaging* (pp. 667-678). Cham: Springer Nature Switzerland.
14. Zhang, L., Yin, X., Liu, X., & Liu, Z. (2025). Medical image segmentation by combining feature enhancement Swin Transformer and UperNet. *Scientific Reports*, 15(1), 14565.
15. An, D., Gu, P., Sonka, M., Wang, C., & Chen, D. Z. (2025, February). Sli2Vol+: Segmenting 3D medical images based on an object estimation guided correspondence flow network. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (pp. 3624-3634). IEEE.
16. Araújo, T., Aresta, G., Galdran, A., Costa, P., Mendonça, A. M., & Campilho, A. (2018, September). Uolo-automatic object detection and segmentation in biomedical images. In *International Workshop on Deep Learning in Medical Image Analysis* (pp. 165-173). Cham: Springer International Publishing.
17. Kanwar, A., Merz, B., Claunch, C., Rana, S., Hung, A., & Thompson, R. F. (2023). Stress-testing pelvic autosegmentation algorithms using anatomical edge cases. *Physics and Imaging in Radiation Oncology*, 25, 100413.
18. Terven, J., Cordova-Esparza, D. M., Romero-González, J. A., Ramírez-Pedraza, A., & Chávez-Urbiola, E. A. (2025). A comprehensive survey of loss functions and metrics in deep learning. *Artificial Intelligence Review*, 58(7), 195.

**Copyright:** ©2026 Febian Febian, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.