

SGRNet: Spatially Guided Radiology Network for Structured Radiological Reporting of Head and Neck Cancer

Ayush Gupta^{1,2}, Vinkle Srivastav^{1,2,3}, Prateek Upadhy^{1,2}, Amit Gupta⁴,
Krithika Rangarajan⁴, and Nicolas Padoy^{1,2}

¹ University of Strasbourg, CNRS, INSERM, ICube, Strasbourg, France

² IHU Strasbourg, France

³ Department of Data Science and AI, Wadhvani School of Data Science and AI (WSAI), Indian Institute of Technology (IIT) Madras, Chennai, India

⁴ All India Institute of Medical Sciences, New Delhi, India
ayush.gupta2@etu.unistra.fr

Abstract. Automated radiological report generation can alleviate clinical workloads and eliminate observer variability. However, standard free-text generation models pose hallucination risks in dense regions and fail under data scarcity. We address these challenges in Head and Neck Cancer (HNC) from contrast-enhanced CT (CECT) imaging. To enforce factual safety, we reformulate report generation as an anatomically grounded, multi-label, structured reporting task, predicting localized tumor involvement across a hierarchical clinical schema. To bridge the visual gap from missing metabolic imaging (e.g., PET), we introduce SGRNet (Spatially Guided Radiology Network), incorporating two low-cost spatial priors: automated organ segmentations and weakly supervised tumor localization maps modeled via 3D Gaussian heatmaps. These priors are dynamically integrated via spatial feature modulation to guide the network toward subtle tumor-induced structural alterations. Evaluated on a multi-centric dataset of 184 paired HNC CECT volumes and reports, on five clinically salient, densely packed anatomical subsites, SGRNet achieves a mean Average Precision (mAP) of 0.60, an 8.8 percentage-point absolute improvement over strong volume-only 3D baselines.

Keywords: Structured Radiology Reports · Head and Neck Cancer · Contrast-Enhanced CT · Spatial Priors · 3D Deep Learning.

1 Introduction

Radiological reports underpin the clinical utility of medical imaging, yet manual reporting remains an operational bottleneck. This burden is acute in head and neck cancer (HNC), an anatomically dense disease site accounting for roughly 890,000 new cases annually [4], where assessing tightly packed subsites (mucosal surfaces, deep spaces, neurovascular bundles) averages 24 minutes per scan [5], making it an ideal target for automation.

Hybrid modalities like PET/CT leverage metabolic tracers to highlight tumor boundaries but are cost-prohibitive. Contrast-enhanced computed tomography (CECT) remains the frontline standard of care but lacks direct metabolic signaling (see Appendix A), forcing tumor localization to rely on subtle tissue attenuation changes alone. Reporting models have progressed in 2D chest X-rays using benchmarks like CheXpert [7] or 3D thoracic data like CT-RATE [9], but fail to generalize to the intricate anatomy of the head and neck.

Furthermore, applying unconstrained autoregressive Vision-Language Models (VLMs) introduces severe safety risks where text hallucinations can radically alter staging. To ensure clinical safety, we reformulate report generation into an anatomically grounded structured reporting paradigm based on a hierarchical schema [28], mapping imaging features to binary indicators of tumor involvement. To drive this domain despite extreme data scarcity, we curated a multi-centric dataset of 184 HNC CECT scans paired with radiologist-validated reports. To overcome low soft-tissue contrast, we propose SGRNet (Spatially Guided Radiology Network), which incorporates two low-cost spatial priors: automated organ segmentations and weakly supervised tumor localization maps modeled via 3D Gaussian heatmaps. These streams are fused via a spatial feature modulation mechanism that dynamically guides the network’s attention toward subtle tumor-induced structural alterations.

Our primary contributions are: (1) We introduce a specialized, multi-centric dataset of 184 paired HNC CECT volumes and structured reports. (2) We propose SGRNet, a prior-guided 3D architecture blending anatomical segmentation and weak pathological localization, providing a scalable alternative to voxel-level labeling. (3) On five clinically salient, densely packed anatomical subsites selected from the full structured schema (Appendix C) as an initial, high-value validation of the approach, we demonstrate that our framework achieves an mAP of 0.60, marking an 8.8% relative improvement over baseline 3D architectures, with high gains in small, complex structures.

2 Related Work

Radiological Report Generation in 2D and 3D Domains: Large-scale paired image-text benchmarks like MIMIC-CXR [6] have fueled advanced 2D multi-modal models for free-text and controlled structured prediction [11], but do not extend to 3D volumetric relationships. Volumetric 3D report generation models like CT2Rep [12] and 3D-CT-GPT [13] rely on thoracic benchmarks like CT-RATE [9]. Their unconstrained text decoders introduce unacceptable hallucination risks if applied to highly intricate head and neck imaging workflows where sub-centimeter errors alter treatment trajectories.

Head and Neck Imaging Benchmarks: Publicly accessible HNC datasets are exceptionally rare and primarily center on metabolic PET/CT imaging [15] and lack corresponding text reports. Standalone anatomical segmentation benchmarks exist (e.g., Walter et al. segmenting 71 distinct structures [16]), but are bounded to pixel labeling. Our framework bridges accessible, standard-of-care

contrast-enhanced CT (CECT) volumes with clinical reports, explicitly utilizing automated segmentation outputs like those from [16] as architectural constraints rather than end-goals.

Weakly Supervised Spatial Priors: To bypass the steep costs of generating voxel-level 3D annotations, weak supervision has emerged as a key strategy. Methods leveraging bounding boxes or extreme points have successfully guided models with minimal manual overhead [18]. Prior works have adapted radiologists’ eye-tracking gaze points to establish spatial focus [19]. Our framework uniquely repurposes 3D Gaussian tumor heatmaps as a low-cost, weak localization prior explicitly integrated to regularize and guide a multi-label structured reporting pipeline.

3 Dataset and Curation Pipeline

Cohort Selection and Multi-Centric Heterogeneity: We curated a multi-centric dataset comprising 184 contrast-enhanced CT (CECT) scans across two distinct clinical streams: (1) *Public Domain Stream*: 49 scans extracted from the ACRIN-HNSCC study (ACRIN 6685) [22] hosted on TCIA, isolated after a quality audit to remove scans with severe motion or artifact profiles. (2) *Institutional Clinical Stream*: 135 contrast-enhanced CT scans consecutively collected from the All India Institute of Medical Sciences (AIIMS), New Delhi, publicly available through the Indian Biological Image Archive (IBIA) under accession CTS_1000000034⁵. All institutional samples were de-identified and approved by local ethical review boards before downstream modeling. Pathologically, primary tumor involvement spans key sub-sites including the supraglottis, transglottis, nasopharynx, and oropharynx.

Standardized Curation and Ground-Truth Structuring: To establish reliable, clinically standardized labels from unstructured text, a dual-stage pipeline was enacted based on the clinical schema in [28]: (1) For the AIIMS scans ($N = 135$), native narrative free-text reports (average length: 275 words) were mapped to a structured template using an LLM to identify binary tumor involvement across subsites, followed by manual expert validation. (2) For the public ACRIN-HNSCC scans ($N = 49$), which lacked matching text records, a senior radiologist manually evaluated each CECT volume and directly populated the identical structured template.

4 Method

We reformulate report generation into a structured multi-label classification task over an itemized anatomical schema to ensure factual safety. SGRNet processes raw CECT volumes alongside low-cost structural and pathological priors to regularize latent features.

⁵ https://ibdc.dbt.gov.in/ibia/study_details_browse_1/CTS_1000000034/




Fig. 1. Architecture of the SGRNet. The tumour heatmap and the organ segmentation stream modulate the CECT stream via spatial feature modulation, allowing the network to emphasise tumour-salient regions while preserving global anatomical context.

4.1 Clinically Controlled Schema Mapping

Unstructured text reports are converted using the schema extraction prompt in [28] with GPT-4. Unmentioned structural fields are designated as “Missing”. Every output vector is audited by expert radiologists, serving as a reliable mapping intermediate between volumetric image patterns and structured categorical descriptors without generation artifacts.

4.2 SGRNet Architectural Framework

SGRNet employs a multi-branch architecture consisting of a primary imaging stream and two auxiliary prior streams that synthesize anatomical and pathological geometry.

Automated Structural Priors (OrgSeg): Anatomical context is derived using TotalSegmentator [16] via a pre-trained nn-UNet engine [17] configured to extract 21 high-risk anatomical subsites relevant to HNC. The discrete categorical label map $S \in \{0, \dots, 21\}^{D \times H \times W}$ is passed through a learned low-dimensional semantic embedding layer to save memory: $E_{\text{emb}}(S) = \text{Embed}(S) \in \mathbb{R}^{d_e \times D \times H \times W}$ (where $d_e = 8$). This dense map is passed into a 3D convolutional encoder (E_{seg}) to produce downsampled structural prior feature maps:

$$f_{\text{seg}} = E_{\text{seg}}(E_{\text{emb}}(S)) \in \mathbb{R}^{C_{\text{seg}} \times d \times h \times w} \quad (1)$$

Weak Pathological Priors via 3D Gaussian Heatmaps: To model tumor-salient regions without requiring costly pixel-level segmentation masks, a 3D nn-UNet pre-trained on external CT channels from standard PET/CT data [15] is fine-tuned on our 39 box-annotated training samples to output coarse binary cuboids. Let $\Omega \subset \mathbb{R}^3$ represent the connected foreground voxels of the thresholded bounding box. We extract its center of mass (μ_x, μ_y, μ_z) and determine its maximum spatial spans along the primary axes to yield bounding dimensions $(w_{\text{cuboid}}, h_{\text{cuboid}}, d_{\text{cuboid}})$. A continuous 3D Gaussian heatmap G is mathematically modeled across the volume space as:

$$G(x, y, z) = \exp \left(- \left(\frac{(x - \mu_x)^2}{2\sigma_x^2} + \frac{(y - \mu_y)^2}{2\sigma_y^2} + \frac{(z - \mu_z)^2}{2\sigma_z^2} \right) \right) \quad (2)$$

where directional standard deviations $\sigma_x, \sigma_y, \sigma_z$ are set directly proportional to the bounding dimensions $w_{\text{cuboid}}, h_{\text{cuboid}}, d_{\text{cuboid}}$, respectively. This continuous volume is mapped through a dedicated 3D Gaussian encoder E_g to produce the pathological prior map: $f_g = E_g(G) \in \mathbb{R}^{C_G \times d \times h \times w}$.

Spatial Feature Modulation: The prior feature maps are passed through a lightweight spatial gating block $A(\cdot)$ comprising $1 \times 1 \times 1$ convolutional layers and a sigmoid activation function to yield a unified spatial focus map $m = A(f_{\text{seg}}) + A(f_g) \in \mathbb{R}^{1 \times d \times h \times w}$. Concurrently, the primary raw CECT volume X_{CECT} is processed by the primary visual trunk E_{CECT} to yield $f_{\text{CECT}} \in \mathbb{R}^{C_{\text{CECT}} \times d \times h \times w}$. Spatial feature modulation is executed via an element-wise soft-gating mechanism:

$$\tilde{f}_{\text{CECT}} = f_{\text{CECT}} \odot (1 + m) \quad (3)$$

where $\odot$ denotes element-wise multiplication, amplifying subtle lesion signatures while preserving global anatomical alignment.

Unified Feature Fusion and Structured Prediction: The modulated feature volume $\tilde{f}_{\text{CECT}}$ is directly concatenated with the raw prior feature maps f_{seg} and f_g , passed through a 3D convolutional fusion network E_{fus} , globally average pooled, and evaluated by an MLP to compute final categorical task probabilities:

$$\hat{y} = \sigma \left(\text{MLP} \left(\text{GAP} \left(E_{\text{fus}} \left(\text{Concat}(\tilde{f}_{\text{CECT}}, f_{\text{seg}}, f_g) \right) \right) \right) \right) \quad (4)$$

where $\hat{y} \in [0, 1]^N$ represents the structured task vector mapped onto the N core targets of the report template.

5 Experiments & Results

5.1 Implementation Details

All network variants were evaluated using a stratified 5-fold cross-validation configuration on the core 145-scan modeling cohort across the five most consistently annotated anatomical subsites: *Tongue*, *Hypopharynx*, *Larynx_air*, *Strap*

muscles + Thyroid gland, and *Carotid arteries + Internal jugular vein (IJV)*. The low-dimensional segmentation embedding block was configured with $d_e = 8$, yielding downsampled feature representations of size $42 \times 32 \times 32$ with 64 channels across all auxiliary trunks. To prevent overfitting, standard 3D spatial transformations (random scaling, rotation, and elastic deformation) were enforced during optimization. Models were trained end-to-end utilizing the Adam optimizer with a learning rate of 1×10^{-5} and a batch size of 2. Severe clinical class imbalances were mitigated using the Asymmetric Loss (ASL) formulation:

$$\mathcal{L}_{\text{ASL}} = - \sum_{c=1}^K \left[y_c (1 - p_c)^{\gamma_+} \log(p_c) + (1 - y_c) p_c^{\gamma_-} \log(1 - p_c) \right] \quad (5)$$

where $\gamma_+ = 1$ and $\gamma_- = 2$ were fixed empirically. Task execution is reported via class-specific and mean Average Precision (mAP). For the weak pathological stream, the localized 3D nn-UNet cuboid engine was pre-trained on the external CT channels of a public PET/CT cohort [15] and fine-tuned for 250 epochs over the 39 box-annotated volumes using a cosine annealing learning rate scheduler, achieving a mean Dice score of 0.51 compared to 0.37 when optimized from scratch.

5.2 Quantitative Results and Ablation Analysis

The upper portion of Table 1 presents a performance evaluation comparing volume-only 3D backbones against our prior-guided variants. Baseline networks relying purely on unconstrained CECT inputs demonstrated moderate predictive capabilities, reaching mAPs of 0.512 (DenseNet121) and 0.494 (EfficientNet-B0). The 3D foundation model (CT-FM) achieved competitive metrics on macro-anatomical pathways but deteriorated sharply on intricate soft-tissue subsets like the strap muscles (0.256 AP).

We additionally report a non-learned *Simple Overlap Baseline*, which flags a structured field as involved whenever the fractional spatial overlap between the corresponding TotalSegmentator organ mask and the fine-tuned tumor mask exceeds a fixed threshold, with no learned classification component; this isolates how much of the observed gains stem from simply possessing the two spatial priors versus their learned, end-to-end fusion in SGRNet.

By contrast, SGRNet variants achieved substantial, uniform improvements across nearly all highly packed targets. SGRNet instantiated on a DenseNet121 trunk achieved an optimal mAP of **0.600**, registering a notable 8.8% relative improvement over its corresponding raw image baseline. These margins are particularly pronounced in ultra-dense structures like the hypopharynx (+0.151 AP) and the laryngeal airway (+0.184 AP).

The lower portion of Table 1 presents a detailed ablation matrix parsing the contributions of individual spatial priors and fusion architectures. The empirical results show that spatial modulation mechanisms consistently outperform simple element-wise addition across both backbones. Integrating anatomical organ maps alone provides a steady baseline boost, but the subsequent integration of the

weak Gaussian tumor heatmap acts as a critical regularizer, driving the DenseNet mAP to its peak performance tier of 0.600.

Table 1. Average Precision (AP) for baseline CECT models, the non-learned overlap baseline, prior-blending ablations, and the final SGRNet variants. (IJV: Internal Jugular Vein. Notable values highlighted in bold; final SGRNet rows shown in bold).

Model	Input	Tongue	Hypo-pharynx	Larynx Air	Strap + Thyroid	Carotid + IJV	Mean AP
DenseNet121	CECT Only	0.549	0.559	0.574	0.357	0.519	0.512
EfficientNet-B0	CECT Only	0.543	0.534	0.563	0.339	0.489	0.494
CT-FM [25]	CECT Only	0.419	0.576	0.612	0.256	0.579	0.488
Simple Overlap Baseline	Org+Tum-Overlap	0.531	0.690	0.724	0.401	0.523	0.574
DenseNet121	Org-Add	0.504	0.725	0.699	0.314	0.627	0.574
DenseNet121	Org-Modul	0.494	0.674	0.756	0.447	0.601	0.594
EfficientNet-B0	Org-Add	0.475	0.708	0.658	0.283	0.593	0.543
EfficientNet-B0	Org-Modul	0.478	0.636	0.701	0.413	0.600	0.567
DenseNet121	Tum-Add	0.528	0.688	0.701	0.342	0.588	0.569
DenseNet121	Tum-Modul	0.536	0.702	0.714	0.355	0.596	0.581
EfficientNet-B0	Tum-Add	0.512	0.671	0.682	0.328	0.571	0.553
EfficientNet-B0	Tum-Modul	0.521	0.684	0.695	0.336	0.579	0.563
DenseNet121	Org+Tum-Add	0.485	0.731	0.715	0.367	0.614	0.582
EfficientNet-B0	Org+Tum-Add	0.431	0.713	0.686	0.345	0.543	0.544
SGRNet (DenseNet)	Org+Tum-Modul	0.478	0.710	0.758	0.501	0.554	0.600
SGRNet (Efficient)	Org+Tum-Modul	0.498	0.674	0.724	0.467	0.558	0.584

6 Discussion & Limitations

The experimental profiles confirm that injecting explicit spatial boundaries and weak pathological context optimizes multi-label classification accuracy within low-contrast 3D medical datasets. Raw, unconstrained volumes fail to provide standard 3D convolutional trunks or large foundation models (CT-FM) with the necessary fine-grained visual attention to parse small, packed structures like the strap muscles or larynx. While CT-FM utilizes a wide pre-training corpus, its representations remain generalized; our localized spatial feature modulation layer acts as a target-specific corrective gate, steering feature maps back to tumor-adjacent boundaries.

On the *Tongue* subsite, the baseline DenseNet121 (0.549 AP) slightly outperformed our prior-guided framework (0.478 AP), likely because oral tongue tissue is highly mobile and deformable during CECT acquisition – unlike rigid structures such as the carotid sheath or laryngeal cartilage, its cross-sectional boundary shifts with deglutition and jaw position, introducing misalignment noise into the TotalSegmentator-derived gating stream. For stable structures, however, the joint prior strategy delivers substantial margins, and the underlying pipeline is broadly scalable: external PET/CT datasets can be stripped to their CT channels to pre-train weak tumor localizers, reducing reliance on costly hybrid scanners.

Limitations: Despite its strengths, several limitations constrain our current pipeline. First, the core modeling size (145 cases) leaves the network sensitive to rare, highly atypical tumor morphologies. Second, modeling weak tumor regions as rigid axis-aligned 3D cuboids occasionally includes healthy adjacent boundaries within the computed Gaussian distributions. Third, errors generated by automated segmentation tools like TotalSegmentator can cascade into downstream feature modulation steps. Fourth, our evaluation targets five high-yield, densely packed subsites rather than the complete structured template (Appendix C), chosen for annotation consistency and diagnostic complexity; we treat this as an initial validation of the reformulation, with full-schema scaling left to future work. Finally, we omit a head-to-head comparison against free-text/VLM report generators (e.g., [12, 13]), as they operate in an unconstrained text space with no shared metric or hallucination benchmark for HNC CECT; building such a benchmark is left to future work.

7 Conclusion

We have presented SGRNet, a prior-guided 3D architecture engineered for the safe, structured radiological reporting of head and neck cancer on contrast-enhanced CT scans. By converting text reports into deterministic, structured subsite targets, we eliminate the severe text hallucination risks tied to unconstrained vision-language models. By dynamically modulating visual features with low-cost anatomical segmentations and weak 3D Gaussian heatmaps, SGRNet successfully circumvents the lack of metabolic signals in standard imaging, yielding an 8.8% relative mAP improvement over raw 3D baselines. Future expansions will focus on multi-institutional validation, semi-supervised representation learning to utilize unlabeled volumes, and investigating deformable registration steps to handle highly mobile oral structures like the tongue.

8 Acknowledgements

This work was partially supported by the Interdisciplinary Thematic Institute HealthTech, as part of the ITI 2021-2028 program of the University of Strasbourg, CNRS, and Inserm, funded by IdEx Unistra (ANR-10-IDEX-0002) and SFRI (STRAT’US project, ANR-20-SFRI-0012) under the framework of the French Investments for the Future Program. This work was also partially supported by French state funds managed by the ANR under Grant ANR-10-IAHU-02. This work was granted access to the HPC resources of IDRIS under the allocations AD011011631R4 made by GENCI. The authors would like to acknowledge the High-Performance Computing Center of the University of Strasbourg for supporting this work by providing scientific support and access to computing resources. Part of the computing resources was funded by the Equipex Equip@Meso project (Programme Investissements d’Avenir) and the CPER Alsacalcul/Big Data.

References

1. Clinical Pathology Market Size, Share, Growth, and Industry Analysis, By Type (Hematology, Microbiology, Biochemistry, Immunology, Molecular Diagnostics), By Application (Healthcare, Hospitals, Laboratories, Research, Clinical Testing, Diagnostics), Regional Insights and Forecast to 2033. Fortune Business Insights (2023). Available at: <https://www.fortunebusinessinsights.com/industry-reports/clinical-pathology-market-102345>
2. American College of Radiology: Reporting and Data Systems (RADS). Available at: <https://www.acr.org/Clinical-Resources/Clinical-Tools-and-Reference/Reporting-and-Data-Systems>
3. Histopathology Reporting Time Study. Archives of Pathology & Laboratory Medicine (2018). Available at: <https://www.archivesofpathology.org/doi/10.5858/arpa.2017-0251-RA>
4. International Agency for Research on Cancer (IARC): GLOBOCAN 2022: Global Cancer Statistics (2022). Available at: <https://gco.iarc.fr/>
5. MacDonald, S.L.S., Cowan, I.A., Floyd, R.A., Graham, R.: Measuring and managing radiologist workload: measuring radiologist reporting times using data from a Radiology Information System. Journal of Medical Imaging and Radiation Oncology 57(5), 558–566 (2013). <https://doi.org/10.1111/1754-9485.12092>
6. Johnson, A.E.W., Pollard, T.J., Berkowitz, S.J., Greenbaum, N.R., Lungren, M.P., Deng, C., Mark, R.G., Horng, S.: MIMIC-CXR: A large publicly available database of labeled chest radiographs. arXiv preprint arXiv:1901.07042 (2019)
7. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Illcus, S., Chute, C., et al.: CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison. Proceedings of the AAAI Conference on Artificial Intelligence 33(1), 590–597 (2019)
8. Bustos, A., Pertusa, A., Salinas, J.M., de la Iglesia-Vaya, M.: PadChest: A large chest X-ray image dataset with multi-label annotated reports. Medical Image Analysis 66, 101797 (2020)
9. Hamamci, I.E., Er, S., Wang, C., Almas, F., Simsek, A.G., Esirgun, S.N., et al.: Developing Generalist Foundation Models from a Multimodal Dataset for 3D Computed Tomography. arXiv preprint arXiv:2403.17834 (2024). Available at: <https://arxiv.org/abs/2403.17834>
10. Yuan, J., Liao, H., Luo, J.: Automatic Radiology Report Generation Based on Multi-view Image Fusion and Medical Concept Enrichment. arXiv preprint arXiv:1907.09085 (2019). Available at: <https://arxiv.org/abs/1907.09085>
11. Wu, X., et al.: MAIRA-2: Grounded Chest X-ray Report Generation with RadFact Evaluation. arXiv preprint arXiv:2406.04449 (2024). Available at: <https://arxiv.org/abs/2406.04449>
12. Hamamci, I.E., Er, S., Menze, B.: CT2Rep: Automated Radiology Report Generation for 3D Medical Imaging. In: Medical Image Computing and Computer Assisted Intervention (MICCAI 2024), LNCS, vol. 15012, pp. 476–486. Springer Nature Switzerland (2024). https://doi.org/10.1007/978-3-031-72390-2_45
13. Chen, H., Zhao, W., Li, Y., Zhong, T., Wang, Y., Shang, Y., Guo, L., Han, J., Liu, T., Liu, J., Zhang, T.: 3D-CT-GPT: Generating 3D Radiology Reports through Integration of Large Vision-Language Models. arXiv preprint arXiv:2409.19330 (2024). Available at: <https://arxiv.org/abs/2409.19330>
14. Clark, K., Vendt, B.: The Cancer Imaging Archive (TCIA) (2013). Available at: <https://www.cancerimagingarchive.net/>

15. Saeed, N., Hassan: A Multimodal and Multi-centric Head and Neck Cancer Dataset for Segmentation, Diagnosis and Outcome Prediction. arXiv preprint arXiv:2509.00367v3 (2025)
16. Walter, A., Hoegen-Saßmannshausen, P., Stanic, G., Rodrigues, J.P., Aderberg, S., Jäkel, O., Frank, M., Giske, K.: Segmentation of 71 Anatomical Structures Necessary for the Evaluation of Guideline-Conforming Clinical Target Volumes in Head and Neck Cancers. *Cancers* 16(2), 415 (2024). <https://doi.org/10.3390/cancers16020415>
17. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical segmentation. *Nature Methods* 18, 203–211 (2021). <https://doi.org/10.1038/s41592-020-01008-z>
18. Roth, H.R., Yang, D., Xu, Z., Wang, X., Xu, D.: Going to Extremes: Weakly Supervised Medical Image Segmentation. *Machine Learning and Knowledge Extraction* 3(2), 507–524 (2021). <https://doi.org/10.3390/make3020026>
19. Zhong, Y., Tang, C., Yang, Y., Qi, R., Zhou, K., Gong, Y., Heng, P.A., Hsiao, J.H., Dou, Q.: Weakly-supervised Medical Image Segmentation with Gaze Annotations. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2024)*, LNCS, vol. 15003, pp. 530–540. Springer (2024). https://doi.org/10.1007/978-3-031-72384-1_50
20. Qu, H., Wu, P., Huang, Q., Yi, J., Yan, Z., Li, K., Riedlinger, G.M., De, S., Zhang, S., Metaxas, D.N.: Weakly Supervised Deep Nuclei Segmentation Using Partial Points Annotation in Histopathology Images. *IEEE Transactions on Medical Imaging* 39(11), 3655–3666 (2020). <https://doi.org/10.1109/TMI.2020.3002244>
21. Chou, Y.C., Lin, C., Ho, C.W., Wang, C.M., Cheng, C.P.: A Weakly-supervised Multi-lesion Segmentation Framework Based on Target-level Incomplete Annotations. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2024)*, pp. 1747–1757. Springer (2024)
22. Kinahan, P., Muzi, M., Bialecki, B., Coombs, L.: Data from the ACRIN 6685 Trial HNSCC-FDG-PET/CT. The Cancer Imaging Archive (TCIA) (2019). <https://doi.org/10.7937/K9/TCIA.2016.JQEJZZNG>
23. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., et al.: TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images. *Radiology: Artificial Intelligence* (2023). <https://doi.org/10.1148/ryai.230024>
24. Adams, L.C., Truhn, D., Busch, F., Kader, A., Niehues, S.M., Makowski, M.R., Bressen, K.K.: Leveraging GPT-4 for Post Hoc Transformation of Free-text Radiology Reports into Structured Reporting: A Multilingual Feasibility Study. *Radiology* 307(4), e230725 (2023). <https://doi.org/10.1148/radiol.230725>
25. Pai, S., Hadzic, I., Bontempi, D., Bressen, K., Kann, B.H., Fedorov, A., Mak, R.H., Aerts, H.J.W.L.: Vision Foundation Models for Computed Tomography. arXiv preprint arXiv:2501.09001 (2025)
26. Huang, G., Liu, Z., Weinberger, K.Q., Van Der Maaten, L.: Densely Connected Convolutional Networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4700–4708 (2017)
27. Tan, M., Le, Q.V.: EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In: *International Conference on Machine Learning (ICML)*, pp. 6105–6114 (2019)
28. Gupta, A., Malhotra, H., Garg, A.K., Rangarajan, K.: Enhancing Radiological Reporting in Head and Neck Cancer: Converting Free-Text CT Scan Reports to Structured Reports Using Large Language Models. *Indian Journal of Radiology and Imaging* 35(1), 43–49 (2025). <https://doi.org/10.1055/s-0044-1788589>

A Comparison of PET/CT, PET/MRI and CECT




Fig. A.1. Visual comparison of imaging modalities. (A) PET/CT highlights metabolically active tumour regions via radiotracer uptake. (B) PET/MRI provides high soft-tissue contrast with PET-guided localisation. (C) CECT lacks metabolic tracers, resulting in more subtle tumour boundaries and motivating the need for automated interpretability methods.

B Example of a Free-Text Radiology Report

Patient 1

Procedure: CECT of Face & Neck.

Clinical background: Ca oropharynx (cT4aN3bM0) post CT/RT.

There is ill-defined enhancing soft tissue thickening in the oropharynx measuring approximately 2.7 cm in longest dimension involving the right lateral and posterior oropharyngeal wall, right vallecula, median glossoepiglottic fold, the tip of the epiglottis, and right base of tongue. Non-enhancing oedematous thickening of the bilateral aryepiglottic folds and vocal cords is noted, likely representing post-radiotherapy changes. The hyoid bone is normal. The prevertebral fat space is maintained.

Enlarged multiple conglomerated necrotic lymph nodes are present in bilateral level II and III regions, the largest measuring approximately 2.0×2.4 cm on the right side. There is encasement of vessels on the right side with thrombosis of the internal jugular vein extending from the jugular foramen to the sigmoid sinus.

Paranasal sinuses/nasal cavity: Mucosal thickening is seen in the right maxillary sinus.

Maxilla/mandible: Normal.

Infra-temporal neck spaces: Normal.

Additional information: None.

Comparison: Compared with the scan dated 17 June 2022, there is a reduction in the size of the primary mass and lymph nodes.

Impression: Irregular enhancing thickening involving the oropharynx with extensions as described above and bilateral cervical conglomerated necrotic lymphadenopathy.

C Structured Reporting Template for Head and Neck Cancers

Table A.1. Structured reporting template used for head and neck cancer CT examinations.

General Information	Additional Structures
Patient ID	Airway compromise
Primary tumour location	Endotracheal tube in-situ
Tumour size	Thyroid cartilage
	Cricoid cartilage
Oral Cavity	
Lip	Esophagus
Buccal mucosa and gingiva	Prevertebral space
Retromolar trigone	Pyriform sinus
Floor of mouth	Post-cricoid region
Oral tongue	Tracheal cartilages
	Strap muscles / thyroid gland
Oropharynx	
Base of tongue	Carotid arteries / IJV
Tonsil (R/L/Both)	Parapharyngeal space
Valleculae	Pterygoid muscles
Soft palate	Skull base
	Orbit
	Parotid gland
Larynx	
Epiglottis	Submandibular gland
False cords	Cranial nerves
Aryepiglottic folds	Intracranial extension
Arytenoid cartilages	
Preepiglottic space	Lymph Node Metastases
Paraglottic space	Number of nodes
Vestibule	Laterality
True vocal cords	Maximum size
Anterior commissure	Extranodal extension
Posterior commissure	Loss of fat planes
Subglottis	

D Summary of Dataset Characteristics

Table A.2. Summary of dataset-level imaging characteristics.

Parameter	Value	Unit
Total scans	184	–
Average voxel spacing	$1 \times 1 \times 1$	mm ³
Average resolution	512×510	pixels
Average number of slices	288	slices
Average mean HU	−763 (Std. 576)	HU
HU range	[−1732, 3444]	HU

E Organ Labels Used for Anatomical Segmentation

Table A.3. Segmented organs and corresponding label identifiers.

ID	Structure
0	Background
1	Larynx air
2	Thyroid cartilage
3	Cricoid cartilage
4	Hyoid bone
5	Tongue
6	Digastric (left)
7	Digastric (right)
8	Sternothyroid (left)
9	Sternothyroid (right)
10	Thyrohyoid (left)
11	Thyrohyoid (right)
12	Submandibular gland (left)
13	Submandibular gland (right)
14	Thyroid gland
15	Internal carotid artery (left)
16	Internal carotid artery (right)
17	Internal jugular vein (left)
18	Internal jugular vein (right)
19	Trachea
20	Oropharynx
21	Hypopharynx

F Segmentation Results




Fig. A.2. Qualitative results from the anatomical organ segmentation model. Panels (A)–(C) illustrate representative outputs from three different patients showing the predicted masks for the 21 anatomical structures.




Fig. A.3. Tumour segmentation results from an nnU-Net model trained exclusively on CT images from the public PET/CT dataset. Panels (A)–(C) show representative examples.




Fig. A.4. Tumour segmentation results from the nnU-Net model pre-trained on the public PET/CT dataset and subsequently fine-tuned using the manually curated cuboid regions. Panels (A)–(C) show representative examples.

G Evaluation Metric

Dice Similarity Coefficient. The Dice similarity coefficient between two binary masks A and B is defined as

$$\text{Dice}(A, B) = \frac{2|A \cap B|}{|A| + |B|},$$

where A and B denote the predicted and ground-truth masks, respectively.