Through-Plane Attention for Anisotropic CT Super-Resolution Improves Quantitative Lung Cancer Analysis

Marc Boubnovski Martell[1], Kristofer Linton-Reid[1], Mitchell Chen[1], Sumeet Hindocha[2], Benjamin Hunter[2], Marco A. Calzado[3,4], Richard Lee[2,5], Joram M. Posma[1], and Eric O. Aboagye[1]

[1] Imperial College London Hammersmith Campus, London, UK
[2] Royal Marsden and Institute of Cancer Research, London, UK
[3] Instituto Maimónides de Investigación Biomédica de Córdoba, Córdoba, Spain
[4] Departamento de Biología Celular, Universidad de Córdoba, Córdoba, Spain
[5] National Heart and Lung Institute, Imperial College London, London, UK
mb618@ic.ac.uk

Abstract. Volumetric computed tomography (CT) often exhibits severe through-plane anisotropy due to thick-slice acquisition, disrupting structural continuity and degrading downstream lung cancer analysis. We propose TVSRN-V2, a volumetric super-resolution (SR) framework with an anisotropy-aware Through-Plane Attention Block (TAB) that models inter-slice dependencies by factoring attention into orthogonal sagittal and coronal planes. To broaden the degradation distribution during training, we use hybrid pseudo-low-resolution (pseudo-LR) augmentation combining real paired scans with synthetic degradations. Clinical utility is evaluated across pulmonary lobe and bronchi segmentation, histology classification, and survival prediction in 742 patients (539 from three independent external cohorts). TVSRN-V2 achieves PSNR 39.16 dB and SSIM 0.946, and yields an absolute Dice improvement of 0.16 for bronchi over bicubic interpolation on real low-resolution CT. SR features improve histology classification (F1 +0.03–0.07, $p < 0.05$) and internal survival prediction (C-index +0.06, $p = 0.02$), with smaller non-significant external prognosis gains.

Keywords: Super Resolution · Swin Transformer · Volumetric CT · NSCLC

1 Introduction

Thick-slice CT acquisitions ($\geq$5 mm) remain common in routine clinical practice, but severe through-plane anisotropy degrades fine airway and lobar anatomy and can reduce the performance of artificial intelligence (AI) models trained on thin-slice CT [1, 10]. Volumetric SR can address this degradation, but thick-to-thin CT SR is ill-posed because slice averaging irreversibly removes anatomical information [12, 14]. Clinically useful SR must therefore recover structural continuity

across slices, not only improve image-fidelity metrics [16, 9, 4]. Recent CT SR methods have addressed anisotropic resolution using paired volumetric learning and TVSRN [17], convolutional-transformer cross-view modeling in CTHNet/DLS [18], self-supervised through-plane enhancement [11], and anisotropic cross-view texture transfer [13]. However, robust SR for heterogeneous clinical acquisitions remains challenging because multi-scanner paired low-/high-resolution (LR/HR) data are scarce, and improved reconstruction metrics do not necessarily imply improved downstream quantitative analysis [15, 6].

To address these limitations, we propose TVSRN-V2, a transformer-based volumetric SR framework for anisotropic CT. TVSRN-V2 builds on TVSRN-style volumetric CT SR [17] and Swin Transformer V2 [8] within an asymmetric encoder–decoder architecture. It incorporates an anisotropy-aware TAB that applies sagittal and coronal self-attention with learned channel-wise fusion to capture inter-slice dependencies along the depth axis. Because related cross-view methods such as CTHNet/DLS also process coronal and sagittal views [18], we position TAB as a targeted refinement of existing through-plane modelling, replacing fixed cross-view fusion with adaptive gating and coupling it to a hybrid real/pseudo-LR training curriculum. This curriculum combines real paired LR/HR scans with synthetically degraded pseudo-LR volumes to improve robustness to scanner- and protocol-dependent degradation.

The main contributions are: (1) an anisotropy-aware TVSRN extension combining Swin V2 components with sagittal/coronal through-plane attention and learned fusion; (2) a hybrid real/pseudo-LR training curriculum designed to broaden the degradation distribution seen during SR training; and (3) a multi-task clinical validation, including pulmonary lobe and bronchi segmentation plus histology classification and prognosis modelling in 742 non-small cell lung cancer (NSCLC) patients, showing that SR-enhanced CT improves downstream quantitative lung cancer analysis.

2 Methods

TVSRN-V2 was evaluated in two stages. First, the SR model was trained and tested on paired LR/HR CT volumes. Second, the trained SR model was applied as a fixed preprocessing step before downstream lung cancer analysis, without retraining downstream models on SR-enhanced inputs.

2.1 Datasets

SR dataset. SR training and testing used RPLHR-CT [17], which contains 250 anonymised clinical CT scans acquired on Philips scanners and reconstructed as paired thin-slice HR (1 mm) and thick-slice LR (5 mm) volumes. The standard split of 100 training, 50 validation, and 100 test volumes was used. To increase degradation diversity, HR volumes were additionally downsampled along the through-plane axis to generate pseudo-LR inputs, with downsampling continued until the effective slice thickness exceeded 3 mm or the number of slices fell

below 130. The SR training curriculum combined real paired LR/HR scans with pseudo-LR/HR pairs.

Segmentation data. Pulmonary lobe, bronchi, and trachea segmentation used 50 thoracic CT scans with manual annotations validated by a board-certified radiologist: 38 scans from the SPIE-AAPM Lung CT Challenge and 12 local anonymised scans from an external clinical centre. Evaluation used 50 LUNA16 scans with slice thickness ≤ 3 mm and an additional external pathological cohort of 24 patients with chronic obstructive pulmonary disease (COPD), lung cancer, COVID-19 pneumonitis, or lobar collapse.

Classification and prognosis data. Histology classification and survival models were trained using a public NSCLC cohort from The Cancer Imaging Archive (TCIA; $n = 203$). External validation used three independent clinical centres from a multi-institutional NSCLC study: Royal Marsden Hospital (RMH), Guy's and St Thomas' Hospital (GSTT), and Imperial College Healthcare NHS Trust (ICHT) ($n = 539$ total) [5].

2.2 TVSRN-V2 Architecture

TVSRN-V2 is an asymmetric encoder–decoder architecture for thick-to-thin volumetric CT SR. The encoder extracts LR features using linear embedding and Swin Transformer V2 layers [8]. The decoder reconstructs HR volumes using upsampling layers, Feature Interaction Modules (FIMs), and TABs. FIMs follow the original TVSRN design and aggregate cross-scale information by combining encoder skip features with upsampled decoder features [17]. CTHNet/DLS is the closest related cross-view through-plane architecture [18], but was not retrained as a matched baseline because the complete paired development data are not publicly available under the RPLHR-CT protocol and the published evaluation targets a different clinical setting.

Through-Plane Attention Block. In anisotropic CT, the through-plane axis is the main dimension of resolution loss. TAB therefore models inter-slice structure using two attention pathways that operate on planes containing the depth axis. Given an intermediate feature volume $F \in R^{B\times C\times D\times H\times W}$, the sagittal pathway reshapes F to $(B \cdot W) \times C \times D \times H$ and applies shifted-window self-attention over the (D, H) plane. The coronal pathway reshapes F to $(B \cdot H) \times C \times D \times W$ and applies shifted-window self-attention over the (D, W) plane. The resulting features, F_{sag} and F_{cor}, are fused using learned channel-wise gating:

$$F_{\text{TAB}} = G \odot F_{\text{sag}} + (1 - G) \odot F_{\text{cor}}, \quad G = \sigma(W_g([F_{\text{sag}}; F_{\text{cor}}])), \tag{1}$$

where W_g is a learned $1 \times 1 \times 1$ convolution, σ is the sigmoid activation, $[\cdot;\cdot]$ denotes channel-wise concatenation, and $\odot$ denotes element-wise multiplication. This gate adaptively weights sagittal and coronal evidence according to local anatomy. Compared with fixed cross-view fusion in prior through-plane designs [18], TAB uses learned channel-wise fusion within a Swin V2 decoder and is trained with a hybrid real/pseudo-LR curriculum.

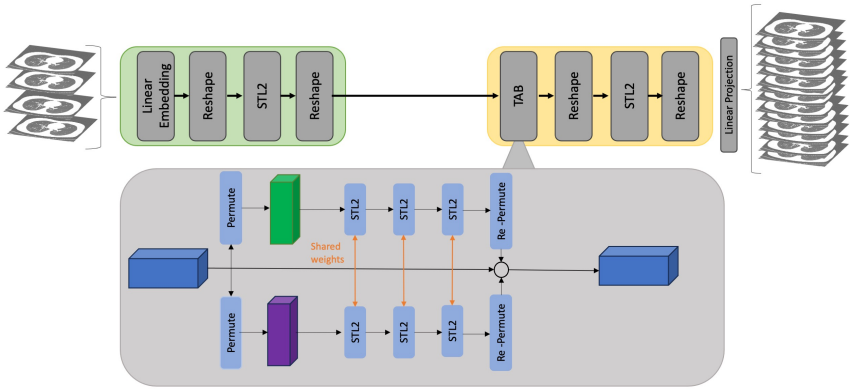

Fig. 1. TVSRN-V2 architecture. The Swin V2 encoder extracts LR features. The decoder reconstructs HR volumes via FIMs and TABs, which fuse sagittal and coronal through-plane dependencies.

2.3 Training and Inference

TVSRN-V2 was implemented in PyTorch and trained on one NVIDIA A6000 GPU. Training used Adam with learning rate 10^{-4} and batch size 1. LR patches of size $4 \times 256 \times 256$ were randomly sampled with corresponding HR targets of size $16 \times 256 \times 256$. The model was trained for 2000 epochs on the 100 real paired RPLHR-CT training scans and fine-tuned using mixed real paired and pseudo-LR/HR volumes. The learning rate, patch size, number of epochs, and window configuration were selected using the RPLHR-CT validation split and fixed for all test and downstream evaluations; no hyperparameters were tuned on the external clinical cohorts.

Mean inference time, measured over the RPLHR-CT test volumes, was approximately 12 s per CT volume, corresponding to approximately 4 s per 100 reconstructed HR slices on a single NVIDIA A6000 GPU.

2.4 Evaluation Protocols

SR reconstruction quality was evaluated on the RPLHR-CT test split using PSNR and SSIM, comparing TVSRN-V2 with bicubic interpolation and the original TVSRN under the same public benchmark protocol [17]. CTHNet/DLS was treated as the closest related cross-view through-plane architecture [18], but was not retrained as a matched baseline because the complete paired development data are not publicly available and the published evaluation targets a different clinical setting. We therefore focus quantitative SR comparisons on the public RPLHR-CT protocol and include qualitative LR/SR/HR comparisons in coronal and sagittal views to assess through-plane structural consistency.

To isolate the clinical effect of SR preprocessing, all downstream models were fixed and applied consistently to native LR and SR-enhanced inputs. Lobar and airway segmentation were evaluated using a V-Net multi-task learning model [2]. For clinical endpoints, histology classification and survival prognosis models were applied to CT scans containing fewer than 200 slices, where through-plane anisotropy is most pronounced. Tumour metabolome representation CT (TMR-CT) features were extracted using a pre-trained deep representation framework [3]. Histology classification was performed using a Random Forest and evaluated via F1-score. Survival prediction utilized a Random Survival Forest and was evaluated via the concordance index (C-index).

Segmentation and histology comparisons used one-sided Wilcoxon signed-rank tests against the relevant baseline. Prognosis comparisons used one-sided Z-tests with 95% bootstrap confidence intervals from 1000 resamples.

3 Results

3.1 Ablation

Table 1 summarises the contribution of each architectural component on the internal test set. The full TVSRN-V2 achieves a PSNR of 39.162 ± 1.834 dB and an SSIM of 0.946 ± 0.028, statistically outperforming all ablated variants ($p < 0.05$, one-sided Wilcoxon signed-rank test). Removing the TAB decreases PSNR by 1.34 dB and SSIM by 0.031, indicating that through-plane attention is a major contributor to reconstruction fidelity. When the Swin V2 encoder is replaced with a standard Vision Transformer (ViT), PSNR decreases by an additional 3.13 dB relative to the no-TAB variant. This result indicates that scaled cosine attention, residual post-normalisation, and continuous relative positional bias contribute to reconstruction fidelity beyond the attention mechanism alone.

Table 1. TVSRN-V2 ablation study on the internal test set ± std. * indicates $p < 0.05$ (one-sided Wilcoxon signed-rank test) versus worse-performing methods. Best results in bold.

Designs	PSNR(↑)	SSIM(↑)
w/o TAB	37.820 ± 1.834*	0.915 ± 0.023
Encoder only	36.456 ± 1.675*	0.900 ± 0.026
ViT Encoder	34.688 ± 1.278	0.871 ± 0.025
TVSRN-V2	**39.162 ± 1.834***	**0.946 ± 0.028***

3.2 SR Image Quality

Table 2 compares TVSRN-V2 with bicubic interpolation and the original TVSRN on the internal test set. TVSRN-V2 achieves a PSNR of 39.162 ± 1.834 dB and an SSIM of 0.946 ± 0.028, significantly outperforming both baselines ($p < 0.05$).

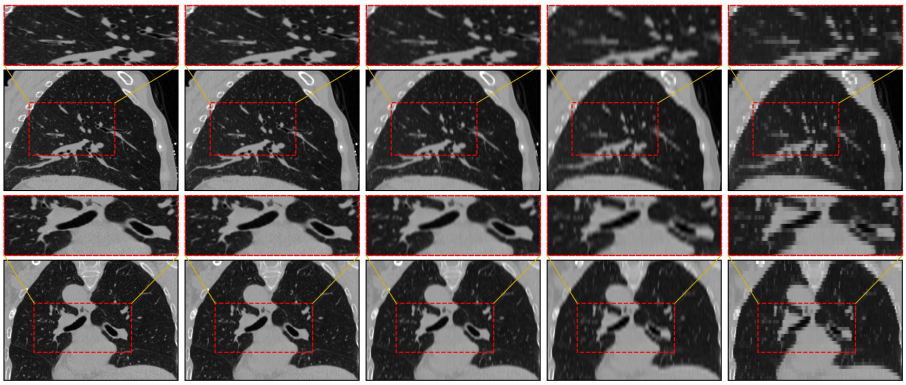

Fig. 2. Qualitative comparison of through-plane CT SR. Matched coronal and sagittal views (ordered by fidelity: 1 mm HR, TVSRN-V2, TVSRN, bicubic, 5 mm LR). Red crops highlight lobar/proximal bronchi and adjacent bronchovascular structures. TVSRN-V2 preserves airway continuity and anatomical boundaries better than baselines, without generating spurious structures.

Compared with the original TVSRN, TVSRN-V2 improves PSNR by 0.584 dB. This improvement demonstrates that the combined introduction of TAB and Swin V2 components yields measurable reconstruction gains beyond the original architecture.

Figure 2 shows matched coronal and sagittal views centred on lobar and proximal segmental bronchi. Bicubic interpolation reduces coarse blur but does not fully restore bronchial continuity, whereas TVSRN-V2 better preserves airway lumen and wall structure relative to TVSRN and approaches the paired HR target.

Table 2. Performance on the internal test set (± std). 95% confidence intervals are shown in brackets. Asterisk (*) indicates $p < 0.05$ vs. TVSRN-V2 (Wilcoxon signed-rank test). Best results in bold.

Model	PSNR (↑)	SSIM (↑)
Bicubic	33.214 ± 1.083* [31.628, 35.225]	0.904 ± 0.028* [0.821, 0.932]
TVSRN	38.578 ± 1.621* [36.012, 41.202]	0.937 ± 0.021* [0.896, 0.969]
TVSRN-V2	**39.162 ± 1.834** **[37.438, 42.122]**	**0.946 ± 0.028** **[0.902, 0.977]**

3.3 Post-SR Lobar and Bronchi Segmentation

Table 3 reports Dice scores on clinically relevant real LR CT scans. TVSRN-V2 consistently outperforms both bicubic interpolation and TVSRN across all

Table 3. Dice scores on real low-resolution CT. Bold indicates best performance. * denotes $p < 0.05$ vs. bicubic interpolation. Lobe abbreviations: RUL = right upper lobe, RML = right middle lobe, RLL = right lower lobe, LUL = left upper lobe, LLL = left lower lobe.

Class	Bicubic	TVSRN	TVSRN-V2
RLL	0.897 ± 0.031	$0.945 \pm 0.028^*$	$\mathbf{0.961 \pm 0.022}^*$
RML	0.891 ± 0.038	$0.919 \pm 0.035^*$	$\mathbf{0.935 \pm 0.030}^*$
RUL	0.902 ± 0.028	$0.948 \pm 0.031^*$	$\mathbf{0.966 \pm 0.018}^*$
LLL	0.904 ± 0.032	$0.944 \pm 0.028^*$	$\mathbf{0.958 \pm 0.025}^*$
LUL	0.895 ± 0.035	$0.954 \pm 0.031^*$	$\mathbf{0.965 \pm 0.021}^*$
Trachea	0.916 ± 0.029	$0.954 \pm 0.025^*$	$\mathbf{0.967 \pm 0.019}^*$
Bronchi	0.508 ± 0.076	$0.608 \pm 0.076^*$	$\mathbf{0.670 \pm 0.070}^*$

Table 4. Impact of SR on Histology Classification of NSCLC. F1-score (± standard error) for adenocarcinoma versus squamous cell carcinoma classification. P-values from one-sided Wilcoxon signed-rank tests compare performance with and without SR.

		TCIA (test)		RMH		GSTT		ICHT	
		F1-score	P-value	F1-score	P-value	F1-score	P-value	F1-score	P-value
TMR-CT	w/o SR	0.84±0.03		0.78±0.02		0.77±0.03		0.79±0.03	
	with SR	**0.87±0.02**	**0.03**	**0.83±0.02**	**0.03**	**0.82±0.03**	**0.04**	**0.86±0.02**	**0.02**

lobar structures, trachea, and bronchi. The largest improvement is observed in bronchi segmentation, where Dice increases from 0.508 with bicubic interpolation to 0.670 with TVSRN-V2, corresponding to an absolute gain of 0.16 ($p < 0.05$).

3.4 Impact on Clinical Endpoints: Histology and Prognosis

This section evaluates TVSRN-V2 as a preprocessing step for two key clinical tasks in NSCLC: histology classification and survival prognosis.

Histology Classification Table 4 reports the impact of SR on the F1-score for histology classification using TMR-CT features. SR preprocessing yielded consistent, statistically significant improvements across all four cohorts ($p < 0.05$). The most notable gains were observed in datasets with the coarsest native slice thicknesses. For instance, in the ICHT dataset, the F1-score improved from 0.79 ± 0.03 to 0.86 ± 0.02 after applying TVSRN-V2.

Prognosis Results Table 5 reports the impact of SR on the concordance index (C-index) for NSCLC prognosis prediction using a Random Survival Forest. Across all cohorts, SR-enhanced features provided consistent prognostic gains. The C-index in the TCIA validation set increased significantly from 0.74 to 0.80 with SR ($p = 0.02$). While not all external improvements reached statistical significance individually, the directional consistency across heterogeneous real-world acquisition protocols supports the robustness of through-plane spatial recovery.

Table 5. Impact of SR on Prognosis of NSCLC. C-index (± standard error) for Random Survival Forest using TMR-CT features. P-values from one-sided Z-tests compare with and without SR.

		TCIA (test)		RMH		GSTT		ICHT	
		C-index	P-value	C-index	P-value	C-index	P-value	C-index	P-value
TMR-CT	w/o SR	0.74±0.03		0.72±0.04		0.71±0.05		0.71±0.04	
	with SR	**0.80±0.04**	**0.02**	**0.76±0.04**	**0.11**	**0.75±0.05**	**0.08**	**0.75±0.05**	**0.07**

4 Discussion

TVSRN-V2 demonstrates that anisotropy-aware SR can improve quantitative lung cancer analysis in thick-slice CT, where radiomics and CT-derived biomarkers are increasingly used for diagnosis, management, and prognosis but remain sensitive to acquisition variability and standardisation challenges [9, 7]. Segmentation provides the clearest structural evidence: on real LR CT, bronchi Dice increased by an absolute margin of 0.16 over bicubic interpolation (0.508 to 0.670). Because small airway structures occupy few voxels along the 5 mm through-plane axis, they are especially vulnerable to slice averaging. The ablation study supports this interpretation: removing TAB caused a substantial drop in reconstruction quality, indicating that attention over planes containing the depth axis contributes to inter-slice structural recovery.

These anatomical gains were associated with modest but consistent downstream improvements. TMR-CT histology classification improved across cohorts, while survival prediction showed a significant internal C-index gain and smaller, non-significant directional improvements externally. The non-significant external prognosis gains suggest that downstream benefit depends on endpoint variability, cohort size, and native acquisition thickness, rather than SR alone.

The main limitation is the absence of a matched comparison with stronger recent 3D and cross-view CT SR baselines, particularly CTHNet/DLS, which also uses coronal/sagittal through-plane processing [18]. TAB should therefore be interpreted as an anisotropy-aware refinement of TVSRN/CTHNet-style through-plane modelling, combining adaptive sagittal/coronal fusion, Swin V2 components, and hybrid real/pseudo-LR training, rather than as a wholly independent architectural class. In addition, although pseudo-LR augmentation broadens the degradation distribution during training, its independent contribution was not isolated by a dedicated ablation. Future work should benchmark against recent self-supervised and cross-view methods [11, 13], assess hallucinated structures qualitatively and quantitatively, and reduce transformer inference cost using efficient attention or distillation. Multi-scanner paired training or domain-adaptive fine-tuning may further improve robustness beyond the single-manufacturer SR training set.

5 Conclusion

We presented TVSRN-V2, an anisotropy-aware extension of TVSRN-style volumetric CT SR that combines Swin V2 components, learned sagittal/coronal through-plane fusion, and hybrid real/pseudo-LR training. The strongest benefit was structural: TVSRN-V2 improved bronchi Dice by an absolute 0.16 over bicubic interpolation on real LR CT and qualitatively improved lobar/proximal airway continuity. Downstream analyses showed consistent histology classification gains, significant survival improvement in TCIA, and smaller non-significant directional prognosis gains across external cohorts. These findings support super-resolution as a practical preprocessing step for improving structural continuity in routine anisotropic CT, while motivating future head-to-head comparison with stronger cross-view and 3D CT SR baselines.

Acknowledgments. This work was supported by Medical Research Council Grant (MR/M015858/1), the Imperial College NIHR Biomedical Research Centre award (WSCC_P62585), the Imperial College Experimental Cancer Medicines award (C1312/A25149), the National Cancer Imaging Translational Accelerator (C2536/A28680) and Stratified Medicine Graduate Training Programme (STRATiGRAD) award WROC_I36048.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abbott, R.E., Nishimwe, A., Wiputra, H., Breighner, R.E., Ellingson, A.M.: A super-resolution algorithm to fuse orthogonal ct volumes using orthofusion. Scientific reports **15**(1), 1382 (2025)
2. Boubnovski, M., Chen, M., Linton-Reid, K., Posma, J., Copley, S., Aboagye, E.: Development of a multi-task learning v-net for pulmonary lobar segmentation on ct and application to diseased lungs. Clinical Radiology **77**(8), e620–e627 (2022)
3. Boubnovski Martell, M., Linton-Reid, K., Hindocha, S., Chen, M., Moreno, P., Álvarez-Benito, M., Salvatierra, Á., Lee, R., Posma, J.M., Calzado, M.A., et al.: Deep representation learning of tissue metabolome and computed tomography annotates nsclc classification and prognosis. NPJ Precision Oncology **8**(1), 28 (2024)
4. de Farias, E.C., Di Noia, C., Han, C., Sala, E., Castelli, M., Rundo, L.: Impact of gan-based lesion-focused medical image super-resolution on the robustness of radiomic features. Scientific reports **11**(1), 21361 (2021)
5. Hindocha, S., Charlton, T.G., Linton-Reid, K., Hunter, B., Chan, C., Ahmed, M., Greenlay, E.J., Orton, M., Bunce, C., Lunn, J., et al.: Gross tumour volume radiomics for prognostication of recurrence & death following radical radiotherapy for nsclc. NPJ Precision Oncology **6**(1), 77 (2022)
6. Jeong, J., Park, D., Kang, J.H., Kim, M., Kim, H.Y., Choi, W., Ham, S.Y.: Deep learning-based slice thickness reduction for computer-aided detection of lung nodules in thick-slice ct. Diagnostics **14**(22), 2558 (2024)

7. Linton-Reid, K., Chen, M., Martell, M.B., Posma, J.M., Aboagye, E.O.: Radiomics in clinical radiology: advances, challenges, and future directions. Clinical Radiology p. 107165 (2025)
8. Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., Ning, J., Cao, Y., Zhang, Z., Dong, L., et al.: Swin transformer v2: Scaling up capacity and resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12009–12019 (2022)
9. Martell, M.B., Linton-Reid, K., Chen, M., Aboagye, E.: Radiomics for lung cancer diagnosis, management, and future prospects. Clinical Radiology **86**, 106926 (2025)
10. Park, S., Lee, S.M., Kim, S., Choi, S., Kim, W., Do, K.H., Seo, J.B.: Performance of radiomics models for survival prediction in non-small-cell lung cancer: influence of ct slice thickness. European Radiology **31**, 2856–2865 (2021)
11. Shi, J., Pelt, D.M., Batenburg, K.J.: Sr4zct: Self-supervised through-plane resolution enhancement for ct images with arbitrary resolution and overlap. In: Machine Learning in Medical Imaging. pp. 52–61. Springer (2023)
12. Song, H., Mao, X., Yu, J., Li, Q., Wang, Y.: I^3net: Inter-intra-slice interpolation network for medical slice synthesis. IEEE Transactions on Medical Imaging **43**(9), 3306–3318 (2024)
13. Uhm, K.H., Cho, H., Hong, S.H., Jung, S.W.: An anisotropic cross-view texture transfer with multi-reference non-local attention for ct slice interpolation. IEEE Transactions on Medical Imaging **45**(1), 336–349 (2026). https://doi.org/10.1109/TMI.2025.3596957
14. Umirzakova, S., Ahmad, S., Khan, L.U., Whangbo, T.: Medical image super-resolution for smart healthcare applications: A comprehensive survey. Information Fusion **103**, 102075 (2024)
15. Xing, X., Li, L., Sun, M., Yang, J., Zhu, X., Peng, F., Du, J., Feng, Y.: Deep-learning-based 3d super-resolution ct radiomics model: predict the possibility of the micropapillary/solid component of lung adenocarcinoma. Heliyon **10**(13) (2024)
16. Yang, H., Zhang, Z., Tian, L., Dang, P., Wan, L., Zhou, Z., Li, J., Cao, L., Yang, X., Chen, J.: Deep learning–based ct slice synthesis improves radiomic feature reproducibility and discriminative performance in lung nodule assessment. Insights into Imaging **17**(1), 163 (2026)
17. Yu, P., Zhang, H., Kang, H., Tang, W., Arnold, C.W., Zhang, R.: Rplhr-ct dataset and transformer baseline for volumetric super-resolution from ct scans. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 344–353. Springer (2022)
18. Yu, P., Zhang, H., Wang, D., Zhang, R., Deng, M., Yang, H., Wu, L., Liu, X., Oh, A.S., Abtin, F.G., et al.: Spatial resolution enhancement using deep learning improves chest disease diagnosis based on thick slice ct. npj Digital Medicine **7**(1), 335 (2024)