

RadiomicsRetrieval: A Customizable Framework for Medical Image Retrieval Using Radiomics Features

Inye Na^{1*}, Nejung Rue^{2*}, Jiwon Chung¹, and Hyunjin Park^{1,2}(✉)

¹ Department of Electrical and Computer Engineering,
Sungkyunkwan University, Suwon 16419, South Korea

² Department of Artificial Intelligence,
Sungkyunkwan University, Suwon 16419, South Korea
hyunjinp@skku.edu

Abstract. Medical image retrieval is a valuable field for supporting clinical decision-making, yet current methods primarily support 2D images and require fully annotated queries, limiting clinical flexibility. To address this, we propose ***RadiomicsRetrieval***, a 3D content-based retrieval framework bridging handcrafted radiomics descriptors with deep learning-based embeddings at the *tumor level*. Unlike existing 2D approaches, *RadiomicsRetrieval* fully exploits volumetric data to leverage richer spatial context in medical images. We employ a promptable segmentation model (e.g., SAM) to derive tumor-specific image embeddings, which are aligned with radiomics features extracted from the same tumor via contrastive learning. These representations are further enriched by anatomical positional embedding (APE). As a result, *RadiomicsRetrieval* enables **flexible querying** based on shape, location, or partial feature sets. Extensive experiments on both lung CT and brain MRI public datasets demonstrate that radiomics features significantly enhance retrieval specificity, while APE provides global anatomical context essential for location-based searches. Notably, our framework requires only minimal user prompts (e.g., a single point), minimizing segmentation overhead and supporting diverse clinical scenarios. The capability to query using either image embeddings or selected radiomics attributes highlights its adaptability, potentially benefiting diagnosis, treatment planning, and research on large-scale medical imaging repositories. Our code is available at <https://github.com/nainye/RadiomicsRetrieval>.

Keywords: Medical Image Retrieval · Radiomics Features · Anatomical Position · Flexible Query.

1 Introduction

The growing volume of medical imaging has made content-based image retrieval (CBIR) a promising tool for radiologists to locate clinically relevant cases [20].

* Equal contribution.




Fig. 1. Illustration of the RadiomicsRetrieval. Users can retrieve similar tumors by providing either an image with minimal point prompts or customized conditions (e.g., anatomical location or partial radiomics features). This diverse setup supports both image-based queries and more general feature-driven searches.

While CBIR methods relying on deep learning-based encoders have surpassed those based on handcrafted features, most remain limited to 2D models that lose vital 3D anatomical context essential for accurate lesion characterization [23,16,15,1,6]. Current CBIR methods emphasize global image-level similarity rather than lesion-specific relevance [22]. Although some methods accept lesion masks as queries, this is often impractical for 3D data due to the labor-intensive nature of volumetric delineation [15]. Moreover, most methods still limit queries to images, masks, or textual labels, which reduces flexibility in real-world clinical settings [8].

We propose *RadiomicsRetrieval*, a CBIR framework that unifies traditional radiomics with deep learning via tumor-level contrastive learning. Radiomics features, established handcrafted descriptors, capture detailed tumor shape and intensity information, but demand extensive manual segmentation [21]. In contrast, deep learning methods can bypass segmentation, yet often fail to retain robust tumor specificity [16,9]. To bridge this gap, we leverage promptable segmentation models such as SAM [14] to derive tumor-specific image embeddings from minimal user input. Radiomics features are separately extracted from the same tumor and contrastive learning between these *image embeddings* and *radiomics embeddings* integrates tumor-specific attributes—shape, intensity, and spatial location—into a unified representation.

A key advantage of *RadiomicsRetrieval* is **flexible querying**: users may query with an image for detailed local and global context or with radiomics features for feature-specific searches (e.g., "find highly spherical tumors in the frontal lobe"). By encoding feature names and values as text-value pairs in a language-model style, we handle varying feature sets and enable partial feature usage (Fig. 1). This design broadens clinical utility and improves search effi-

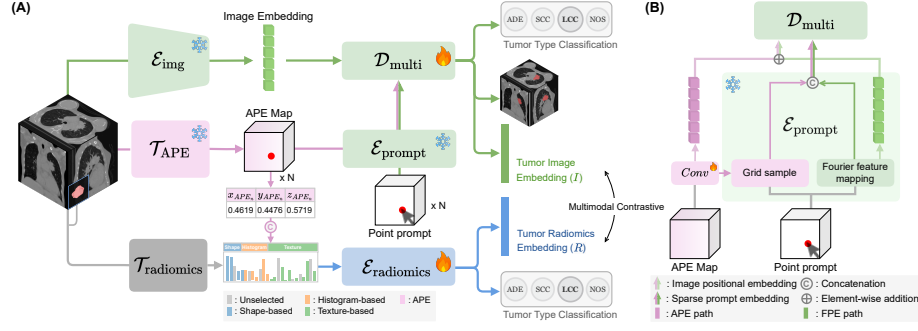


Fig. 2. (A) Overall architecture of *RadiomicsRetrieval*, showing parallel paths for tumor image (green), radiomics (blue), and anatomical (purple) embeddings. (B) Detailed view of how the anatomical positional embedding (APE) is integrated with the Fourier-based positional encoding (FPE).

ciency, making *RadiomicsRetrieval* a highly adaptable CBIR solution. Our main contributions are:

1. **Hybrid Radiomics–Deep Learning Approach:** Combines traditional radiomics with deep learning via tumor-level contrastive learning for enhanced clinical relevance.
2. **Flexible Query Modes:** Supports both images for detailed context and radiomics embeddings for feature-specific or partial-feature queries.
3. **Minimizing Labeling Efforts:** Uses promptable segmentation models to derive tumor-specific image embeddings from minimal prompts, reducing the burden of manual segmentation.

2 Method

RadiomicsRetrieval operates on each tumor through two parallel paths (Fig. 2 (A)): an **image-query path** and a **radiomics-query path**. The image path generates a *tumor image embedding* (I) via a promptable segmentation model, while the radiomics path produces a *tumor radiomics embedding* (R) from hand-crafted features. These embeddings are aligned via *tumor-level contrastive learning*, yielding a shared latent space where clinically similar lesions cluster and dissimilar lesions are pushed apart.

2.1 Image-Query Path

We extend SAM-Med3D [26], which typically outputs only a segmentation mask for the prompted region of interest (ROI), by augmenting its lightweight mask decoder with additional *classification* and *contrastive* tokens—denoted $\mathcal{D}_{multi}$. Given a volumetric input ($128 \times 128 \times 128$) and a minimal prompt (e.g., a single

point), the model produces a *tumor image embedding* (I), along with auxiliary mask and classification outputs. Only during training, full segmentation masks guide the network in learning tumor-specific (ROI-level, not image-level) embeddings. However, when constructing the reference database (i.e., for retrieval), the model requires only a few point prompts per tumor, thereby reducing the need for manual ROI specification. These embeddings then populate a reference database, against which future queries are compared (Fig. 1, top-right).

2.2 Radiomics-Query Path

A *tumor radiomics embedding* (R) is encoded with TransTab [27] by tokenizing each radiomics feature as a feature name-value pair. Users can specify any subset of features without retraining, supporting flexible search conditions. We extract 72 standard radiomics features per tumor, including 14 shape-based, 18 histogram, and 40 texture features via PyRadiomics [25]. During training, we randomly sample radiomics subsets ranging from 72 to 1 feature, enhancing robustness for partial feature queries (Fig. 3(B)).

2.3 Anatomical Positional Embedding (APE)

We incorporate global anatomical context using a 3-channel voxel-based APE [13], integrating it into both image and radiomics paths (Fig. 2(A)). In the **image path**, the APE map is downsampled through convolutional layers to align with SAM’s Fourier-based positional encoding (FPE) [24] dimensions and then integrated element-wise (Fig. 2(B)). Additionally, sampled APE values at prompt coordinates are concatenated with sparse prompt embeddings to reinforce local anatomical context. In the **radiomics path**, voxel-level APE explicitly enriches tumor-intrinsic radiomics features, improving location-aware retrieval. Specifically, APE values at point-prompt locations are concatenated with radiomics subsets as input to $\mathcal{E}_{\text{radiomics}}$. During training, we randomly include APE and also train APE-only cases, enabling flexible querying (radiomics-only, APE-only, or combined) as illustrated in Fig. 3(B). For lung CT (NSCLC), we use an APE model pretrained on abdominal and thoracic CT [13]; for brain tumor MRI, we pre-trained the APE model with BraTS MRI data.

Training Details

Loss Functions. We freeze large pretrained modules (SAM and APE) and fine-tune only the lightweight decoder ($\mathcal{D}_{\text{multi}}$) and the radiomics encoder. We jointly optimize: (1) $\mathcal{L}_{\text{seg}}$ (Dice + Cross-Entropy) for segmenting the prompted tumor, (2) $\mathcal{L}_{\text{contrast}}$ (InfoNCE-style [7]) for clustering embeddings of the same tumor and separating others, and (3) $\mathcal{L}_{\text{cls}}$ (Binary Cross-Entropy or Cross-Entropy) for classifying tumor type in both the image and radiomics pathways.




Fig. 3. Multi-positive contrastive learning. (A) Same-tumor embeddings are *positives* (yellow); different tumors are *negatives* (gray). (B) Radiomics feature subsets. (C) Paired random cropping (twice per tumor).

Multi-positive Contrastive Setup. As shown in Fig. 3, we use a batch size of 80, containing 40 distinct tumors, each repeated with randomly chosen point prompts (1–10) and randomized crops. This produced two image embeddings (e.g., I' and I'') per tumor and multiple radiomics embeddings (e.g., $R^{72'} \dots R^{0'}$, $R^{72''} \dots R^{0''}$) incorporating matching APE values. Embeddings from the same tumor formed *positives*, while embeddings from other tumors served as *negatives*, enabling multi-positive contrastive learning.

Retrieval Inference. At inference time, a user may provide (1) an image with minimal point prompts, yielding a *tumor image embedding* via the image path, or (2) a set (or subset) of radiomics features, APE, or both, yielding a *tumor radiomics embedding*. The resulting query embedding is then compared against the reference database using cosine similarity in the shared latent space, allowing both precise image-based searches and flexible feature-driven queries (e.g., all lesions with a diameter exceeding 30 mm in the left frontal lobe).

3 Dataset

Brain Tumor MRI. We used the Adult Glioma (GLI) [3,19,4] and Meningioma (MEN) [17] subsets from the BraTS 2023 Challenge. Only the tumor core, comprising enhancing tumor and necrotic areas, was considered. Patients were split (8:1:1, stratified by tumor type) into 1440 (GLI 724, MEN 716) tumors for training, 176 (GLI 89, MEN 87) for validation, and 179 (GLI 90, MEN 89) for testing. For retrieval evaluation, ANTsPy’s non-linear registration aligned anatomical masks from the ICBM T1 atlas [12,11,10], yielding precise left/right brain parcellations.

Lung Tumor CT. We combined non-small cell lung cancer (NSCLC) data from multiple sources (NSCLC-Radiomics [2], NSCLC-Radiomics-Interobserver1 [29],

Table 1. Location-Based Evaluation. $P@k$ indicates the mean precision@k, i.e., the fraction of top- k retrieved samples that lie in the same region as the query. **Bold** denotes the best and underline is the second best in a given category.

Query type	Training modality	Contrastive	BraTS		NSCLC	
			P@5	P@10	P@5	P@10
Image & 1 tumor point	<Image,FPE>	unimodal	0.8358	0.8028	0.9531	0.9167
	<Image,APE>	unimodal	0.7899	0.7598	<u>0.9271</u>	0.8766
	<Image,Radiomics,FPE>	multimodal	0.6458	0.6514	0.6927	0.6807
	<Image,Radiomics,APE>	multimodal	0.9117	0.9145	0.9115	<u>0.9000</u>
	<Image,Radiomics,FPE+APE>	multimodal	<u>0.9084</u>	<u>0.8933</u>	0.9062	0.8984
1 point	<Image,Radiomics,APE>	multimodal	0.8905	0.9000	<u>0.9000</u>	0.8891
APE	<Image,Radiomics,FPE+APE>	multimodal	<u>0.8603</u>	<u>0.8615</u>	0.9010	<u>0.8880</u>

RIDER-LungCT-Seg [30], NSCLC Radiogenomics [5], and LUNG-PET-CT-Dx [18]), focusing only on NSCLC cases. The dataset was split (7:1:2) at the patient level, stratified by subtype (adenocarcinoma [ADC], squamous cell carcinoma [SCC], large cell carcinoma [LCC], not otherwise specified [NOS]). This yielded 672 (SCC 168, LCC 85, ADC 302, NOS 48, Unlabeled 69) tumors for training, 100 (SCC 24, LCC 14, ADC 45, NOS 7, Unlabeled 10) for validation, and 192 (SCC 65, LCC 35, ADC 74, NOS 18) for testing. Unlabeled tumors were ignored in the classification loss. We used TotalSegmentator [28] to obtain anatomical masks, including five lung lobes (upper/lower on each side plus the right middle lobe).

4 Results

We evaluated retrieval performance using the test set as queries and the training set as the reference database. Specifically, we constructed the reference database using image embeddings derived from a single tumor-center prompt per training image. To examine how radiomics features and APE influenced retrieval, we compared five training settings. We started with a baseline that used only an image path (*unimodal-contrastive*) with FPE (<Image,FPE>), and then either kept or replaced FPE with APE or added a radiomics path (*multimodal-contrastive*) as follows:

- <Image,APE>: Replaced FPE with APE.
- <Image,Radiomics,FPE>: Kept FPE and added a radiomics path.
- <Image,Radiomics,APE>: Replaced FPE with APE and added a radiomics path.
- <Image,Radiomics,FPE+APE>: Combined both FPE and APE alongside the radiomics path.

4.1 Evaluation of Global Context

Table 1 reports location-based retrieval performance by measuring the fraction of retrieved tumors in the same region as the query. The *Image & 1 tumor*

Table 2. Radiomics-Based Evaluation. *Top-k* is the average correlation between the query’s radiomics features and those of the top-*k* retrieved samples. *Upper Bound* selects the top-*k* by directly comparing radiomics features in the reference set, while *Random Baseline* randomly selects samples from the reference set. **Bold** denotes the best and underline is the second best in a given category.

Query type	Training modality	Contrastive	BraTS		NSCLC	
			Top-5	Top-10	Top-5	Top-10
Image & 1 tumor point	<Image,FPE>	unimodal	0.8582	0.8476	0.8284	0.8161
	<Image,APE>	unimodal	0.8495	0.8418	0.7649	0.7543
	<Image,Radiomics,FPE>	multimodal	0.9197	0.9177	0.8855	0.8818
	<Image,Radiomics,APE>	multimodal	0.8959	0.8866	0.8316	0.8159
	<Image,Radiomics,FPE+APE>	multimodal	<u>0.8987</u>	<u>0.8904</u>	<u>0.8403</u>	<u>0.8185</u>
Radiomics features	<Image,Radiomics,FPE>	multimodal	0.9663	0.9621	0.9448	0.9324
	<Image,Radiomics,APE>	multimodal	0.9484	0.9426	0.8839	0.8847
	<Image,Radiomics,FPE+APE>	multimodal	<u>0.9499</u>	<u>0.9457</u>	<u>0.9051</u>	<u>0.8956</u>
Upper Bound			0.9794	0.9756	0.9711	0.9648
Random Baseline			0.7745	0.7806	0.6815	0.6790

point query type used an image with a single tumor center prompt and the resulting image embedding for retrieval. Among *unimodal-contrastive* models, <Image,FPE> surpassed <Image,APE>, indicating that FPE captured finer spatial details than the coarser, “average-patient”-style APE. Although APE alone lacked high-frequency precision, it was crucial in the *radiomics-query path*, as radiomics features solely describe tumor-internal properties. The model learned global location features by integrating APE into the radiomics path during multimodal contrastive training with image embeddings.

Turning to *multimodal-contrastive* approaches, those that used APE achieved the highest performance, showing that global anatomical cues improved location-specific matching. In contrast, <Image,Radiomics,FPE> focused on tumor-internal features and was less effective for anatomical retrieval. Finally, even a 1 *point* APE query localized tumors accurately, underscoring the flexibility of our method for purely anatomical searches.

4.2 Evaluation of Local Context

Table 2 shows correlations of tumor-internal radiomics features between queries and retrievals. *Multimodal-contrastive* methods substantially outperformed *unimodal-contrastive* approaches, demonstrating the benefit of combining image and radiomics paths. Notably, when radiomics features (without APE) were used as the query, the method focused more on tumor-internal attributes, yielding even higher correlation scores than image-based queries. In particular, <Image,Radiomics,FPE> achieved the highest scores, suggesting this setting prioritized local tumor details over global anatomical cues. Meanwhile, <Image,Radiomics,FPE+APE> remained competitive for local context while offering stronger location-based retrieval (see Table 1), suggesting a robust trade-off between capturing tumor-specific details and anatomical positioning.




Fig. 4. Retrieved Samples Under Different Query Types. Using the same query tumor, varying the retrieval condition (image-based, location-only, or size-only, etc.) produces distinct yet valid results, highlighting *RadiomicsRetrieval*’s flexibility. The histogram under each tumor image depicts its radiomics feature values. All retrievals shown here were generated with the `<Image,Radiomics,FPE+APE>` configuration.

4.3 Comparisons of Retrieved Results under Different Conditions

Fig. 4 illustrates how the retrieved samples vary when using different query types on the *same* tumor. An **image query** with a single point prompt considers the tumor’s overall shape, size, and intensity; a **1 point APE** query focuses solely on anatomical position, potentially returning diverse tumor appearances in that region; and a **VoxelVolume feature** query emphasizes size alone, disregarding location and shape. These examples demonstrate how *RadiomicsRetrieval* offers flexible, customizable searches that highlight different tumor attributes depending on the user’s clinical or research priorities.

5 Conclusion

We present *RadiomicsRetrieval*, a flexible CBIR framework aligning image embeddings and radiomics embeddings at the tumor level via minimal prompts. By integrating APE for global context and radiomics features for tumor-specific characteristics, our method achieves robust performance across diverse query modes—including image-based, location-based, and size-based retrieval. *RadiomicsRetrieval* provides a versatile clinical tool capable of addressing various querying needs. Future directions include uncertainty quantification, systematic radiomics analyses, and query expansions.

Acknowledgments. This study was supported by National Research Foundation of Korea (RS-2024-00408040), AI Graduate School Support Program (Sungkyunkwan University) (RS-2019-II190421), ICT Creative Consilience program (IITP-2025-RS-2020-II201821), and the Artificial Intelligence Innovation Hub program (RS-2021-II212068).

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

References

1. Abacha, A.B., Santamaria-Pang, A., Lee, H.H., Merkow, J., Cai, Q., Devarakonda, S.T., Islam, A., Gong, J., Lungren, M.P., Lin, T., et al.: 3d-mir: A benchmark and empirical study on 3d medical image retrieval in radiology. arXiv preprint arXiv:2311.13752 (2023)
2. Aerts, H.J., Wee, L., Rios Velazquez, E., Leijenaar, R.T., Parmar, C., Grossmann, P., Carvalho, S., Bussink, J., Monshouwer, R., Haibe-Kains, B., et al.: Data from nscic-radiomics. The Cancer Imaging Archive (2014)
3. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F.C., Pati, S., et al.: The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv preprint arXiv:2107.02314 (2021)
4. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J.S., Freymann, J.B., Farahani, K., Davatzikos, C.: Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. Scientific data **4**(1), 1–13 (2017)
5. Bakr, S., Gevaert, O., Echegaray, S., Ayers, K., Zhou, M., Shafiq, M., Zheng, H., Zhang, W., Leung, A., Kadoch, M., et al.: Data for nscic radiogenomics. The Cancer Imaging Archive (2017)
6. Böttcher, B., van Assen, M., Fari, R., von Knebel Doeberitz, P.L., Kim, E.Y., Berkowitz, E.A., Meinel, F.G., De Cecco, C.N.: Evaluation of a content-based image retrieval system for radiologists in high-resolution ct of interstitial lung diseases. European Radiology Experimental **9**(1), 4 (2025)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PmLR (2020)

8. Chen, Y., Liu, C., Liu, X., Arcucci, R., Xiong, Z.: Bimcv-r: A landmark dataset for 3d ct text-image retrieval. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 124–134. Springer (2024)
9. Choe, J., Hwang, H.J., Seo, J.B., Lee, S.M., Yun, J., Kim, M.J., Jeong, J., Lee, Y., Jin, K., Park, R., et al.: Content-based image retrieval by using deep learning for interstitial lung disease diagnosis with chest ct. *Radiology* **302**(1), 187–197 (2022)
10. Collins, D.L., Zijdenbos, A.P., Baaré, W.F., Evans, A.C.: Animal+ insect: improved cortical structure segmentation. In: Information Processing in Medical Imaging: 16th International Conference, IPMI'99 Visegrád, Hungary, June 28–July 2, 1999 Proceedings 16. pp. 210–223. Springer (1999)
11. Fonov, V., Evans, A.C., Botteron, K., Almli, C.R., McKinstry, R.C., Collins, D.L., Group, B.D.C., et al.: Unbiased average age-appropriate atlases for pediatric studies. *Neuroimage* **54**(1), 313–327 (2011)
12. Fonov, V.S., Evans, A.C., McKinstry, R.C., Almli, C.R., Collins, D.: Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage* **47**, S102 (2009)
13. Goncharov, M., Samokhin, V., Soboleva, E., Sokolov, R., Shirokikh, B., Belyaev, M., Kurmukov, A., Oseledets, I.: Anatomical positional embeddings. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 68–77. Springer (2024)
14. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4015–4026 (2023)
15. Kobayashi, K., Gu, L., Hataya, R., Mizuno, T., Miyake, M., Watanabe, H., Takahashi, M., Takamizawa, Y., Yoshida, Y., Nakamura, S., et al.: Sketch-based semantic retrieval of medical images. *Medical Image Analysis* **92**, 103060 (2024)
16. Kumar, A., Kim, J., Cai, W., Fulham, M., Feng, D.: Content-based medical image retrieval: a survey of applications to multidimensional and multimodality data. *Journal of digital imaging* **26**, 1025–1039 (2013)
17. LaBella, D., Adewole, M., Alonso-Basanta, M., Altes, T., Anwar, S.M., Baid, U., Bergquist, T., Bhalerao, R., Chen, S., Chung, V., et al.: The asnr-miccai brain tumor segmentation (brats) challenge 2023: Intracranial meningioma. *arXiv preprint arXiv:2305.07642* (2023)
18. Li, P., Wang, S., Li, T., Lu, J., HuangFu, Y., Wang, D.: A large-scale ct and pet/ct dataset for lung cancer diagnosis (lung-pet-ct-dx). *The cancer imaging archive* (2020)
19. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging* **34**(10), 1993–2024 (2014)
20. Müller, H., Michoux, N., Bandon, D., Geissbühler, A.: A review of content-based image retrieval systems in medical applications—clinical benefits and future directions. *International journal of medical informatics* **73**(1), 1–23 (2004)
21. Na, I., Noh, J.J., Kim, C.K., Lee, J.W., Park, H.: Combined radiomics-clinical model to predict platinum-sensitivity in advanced high-grade serous ovarian carcinoma using multimodal mri. *Frontiers in Oncology* **14**, 1341228 (2024)
22. Öztürk, Ş., Çelik, E., Çukur, T.: Content-based medical image retrieval with opponent class adaptive margin loss. *Information Sciences* **637**, 118938 (2023)
23. Qayyum, A., Anwar, S.M., Awais, M., Majid, M.: Medical image retrieval using deep convolutional neural network. *Neurocomputing* **266**, 8–20 (2017)

24. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems* **33**, 7537–7547 (2020)
25. Van Griethuysen, J.J., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R.G., Fillion-Robin, J.C., Pieper, S., Aerts, H.J.: Computational radiomics system to decode the radiographic phenotype. *Cancer research* **77**(21), e104–e107 (2017)
26. Wang, H., Guo, S., Ye, J., Deng, Z., Cheng, J., Li, T., Chen, J., Su, Y., Huang, Z., Shen, Y., et al.: Sam-med3d: towards general-purpose segmentation models for volumetric medical images. *arXiv preprint* (2023)
27. Wang, Z., Sun, J.: Transtab: Learning transferable tabular transformers across tables. *Advances in Neural Information Processing Systems* **35**, 2902–2915 (2022)
28. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., et al.: Totalsegmentator: robust segmentation of 104 anatomic structures in ct images. *Radiology: Artificial Intelligence* **5**(5) (2023)
29. Wee, L., Aerts, H., Kalendralis, P., Dekker, A.: Data from nslc-radiomics-interobserver1. The Cancer Imaging Archive (2019)
30. Wee, L., Aerts, H., Kalendralis, P., Dekker, A.: Rider lung ct segmentation labels from: Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. The Cancer Imaging Archive (2020)