



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# BIMCV-R: A Landmark Dataset for 3D CT Text-Image Retrieval

Yinda Chen<sup>1,2</sup>, Che Liu<sup>3</sup>, Xiaoyu Liu<sup>1</sup>, Rossella Arcucci<sup>3</sup>, and Zhiwei Xiong<sup>1,2</sup>(✉)

<sup>1</sup>MoE Key Laboratory of Brain-inspired Intelligent Perception and Cognition,  
University of Science and Technology of China

<sup>2</sup>Anhui Province Key Laboratory of Biomedical Imaging and Intelligent Processing,  
Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

<sup>3</sup>Data Science Institute, Imperial College London  
cyd0806@mail.ustc.edu.cn, zwxiong@ustc.edu.cn

**Abstract.** The burgeoning integration of 3D medical imaging into health-care has led to a substantial increase in the workload of medical professionals. To assist clinicians in their diagnostic processes and alleviate their workload, the development of a robust system for retrieving similar case studies presents a viable solution. While the concept holds great promise, the field of 3D medical text-image retrieval is currently limited by the absence of robust evaluation benchmarks and curated datasets. To remedy this, our study presents a groundbreaking dataset, BIMCV-R, which includes an extensive collection of 8,069 3D CT volumes, encompassing over 2 million slices, paired with their respective radiological reports. Expanding upon the foundational work of our dataset, we craft a retrieval strategy, MedFinder. This approach employs a dual-stream network architecture, harnessing the potential of large language models to advance the field of medical image retrieval beyond existing text-image retrieval solutions. It marks our preliminary step towards developing a system capable of facilitating text-to-image, image-to-text, and keyword-based retrieval tasks. Our project is available at <https://huggingface.co/datasets/cyd0806/BIMCV-R>.

**Keywords:** 3D medical imaging · 3D text-image retrieval · BIMCV-R.

## 1 Introduction

The rapid evolution of medical imaging technologies, especially in the realm of 3D imaging, has brought about a transformative change in radiological diagnostics [1–4]. These technologies provide detailed, three-dimensional visualizations that are crucial for accurate lesion detection and measurement, which are pivotal in disease staging, treatment planning, and prognosis assessment [5, 6]. However, the sheer volume and complexity of this spatial data have substantially increased the workload of clinicians, prompting the need for Artificial Intelligence (AI) to assist in the analysis process. AI applications such as automatic segmentation [7, 8], reconstruction [9–12], and denoising [13] are now integral in providing clearer insights and improving diagnostic accuracy.

AI’s role in diagnostics is further expanded through image recognition tools, including detection [14] and classification [15], which are increasingly being used in conjunction with clinical textual reports to provide a more holistic understanding of medical conditions. Initiatives such as Quilt-1m [16], and pre-training projects like BiomedCLIP [17] and MedClip [18], as well as efforts in 3D image description by GTGM [19], T3D [20], and 3d-MIR [21], reflect the ongoing efforts to connect medical images with their textual counterparts. Despite these advancements, the focus on 2D image-text pairs and the challenges with model-generated descriptions highlight the need for a comprehensive benchmark that integrates 3D medical imaging with textual diagnostics, a significant void in the current landscape [22].

Building on these efforts, we introduce BIMCV-R, a new endeavor that directly addresses the need for a unified benchmark in 3D medical imaging diagnostics. Utilizing the comprehensive medical data repository BIMCV [23] and in collaboration with clinicians, we curate **the first publicly accessible dataset** that features 8,069 3D medical image-report pairs, covering 96 disease types. Conscious of privacy concerns, we anonymize radiological reports and translate the original Spanish dataset into English using GPT-4, with meticulous human proofreading to ensure accuracy and reliability. Furthermore, we develop MedFinder, a dual-stream network architecture that leverages the advanced capabilities of the large language model BiomedCLIP [17] to establish a bridge between medical images and reports through text-image retrieval tasks. This initiative lays the foundation for benchmarks in text-image and keyword retrieval, significantly simplifying the process for physicians and clinicians to search for and reference similar cases, thereby enhancing diagnostic precision and efficiency.

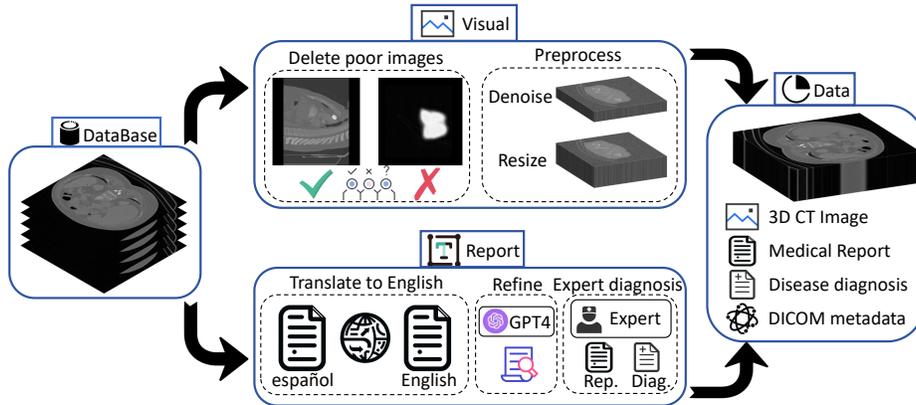
Our contributions are manifold and significant:

1. We curate the first publicly accessible English 3D text-image CT dataset BIMCV-R, inclusive of authentic radiological reports and detailed disease-type diagnoses.
2. We introduce MedFinder, an exhaustive suite of medical retrieval schemes, including innovative approaches for text-image, image-text, and keyword-image retrieval—a pioneering effort on a real-world dataset.
3. By harnessing the power of pre-trained large language models, we showcase their untapped potential in enhancing 3D medical image retrieval, thereby filling a critical void in the field and setting a new standard for future research and application in medical image analysis and retrieval.

## 2 Dataset

This paper presents the BIMCV-R dataset, a substantial resource meticulously crafted for 3D medical multimodal retrieval. This dataset is an extension of the BIMCV dataset [23], encompassing pristine CT scan images, detailed radiological reports, and comprehensive DICOM metadata.

**Data Acquisition and Processing.** The acquisition process of BIMCV-R, as illustrated in Figure 1, commenced with the initial phase of our dataset processing



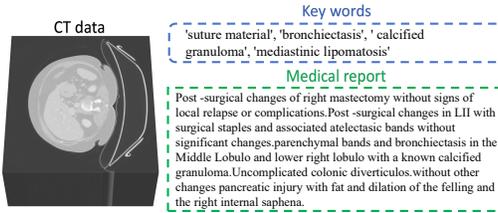
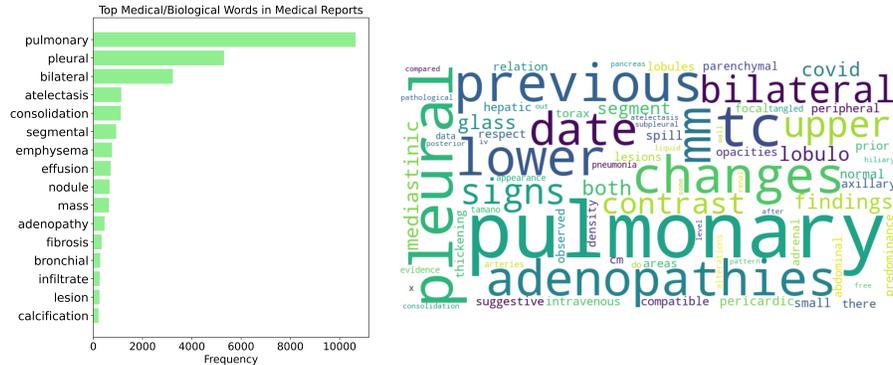
**Fig. 1:** Construction of the BIMCV-R dataset. Utilizing the BIMCV dataset, we enhanced image quality through selective filtering, advanced denoising, and size standardization. For textual data, we translated radiological reports into English and refined them with GPT-4, ensuring consistency. Expert reviews and diagnoses further ensured data reliability and accuracy.

where we eliminated image instances with pixel missing values exceeding 30%, and discarded CT scan samples with any dimension (width, height, or depth) less than 96. Subsequently, CT images of superior imaging quality were manually selected, and images lacking corresponding medical descriptions were removed. Regarding the textual content, descriptions shorter than five words were omitted, and personal information within the text descriptions, such as names and addresses, was anonymized. Following this, all textual descriptions were translated into English using GPT-4 and underwent a manual verification process, culminating in a dataset comprising 8,069 paired samples with more than 2M slices. More importantly, we engaged over 20 medical professionals to diagnose 1,475 of these samples, identifying 96 different diseases, including tumors, infectious diseases, cardiovascular diseases, and respiratory conditions. These diagnoses facilitated the development of an extensive keyword library. Ultimately, we constructed a dataset exceeding 700GB, encompassing original CT scan images, radiological reports, and DICOM metadata, offering a comprehensive resource for medical research and application development in the field of medical imaging.

**Data Statistics Analysis.** The BIMCV-R dataset is distinguished by its inclusion of high-resolution 3D medical images paired with corresponding radiological reports, providing a rich resource for training deep learning models capable of understanding and processing 3D medical imagery. Through this approach, models can learn the correlation between visual features extracted from images and the linguistic descriptions found in radiological reports. We have compiled basic statistics of the dataset as shown in Table 1, and we present sample data as illustrated in Figure 2. Furthermore, based on diagnoses from medical experts,

**Table 1:** Summary of Image and Report Statistics.

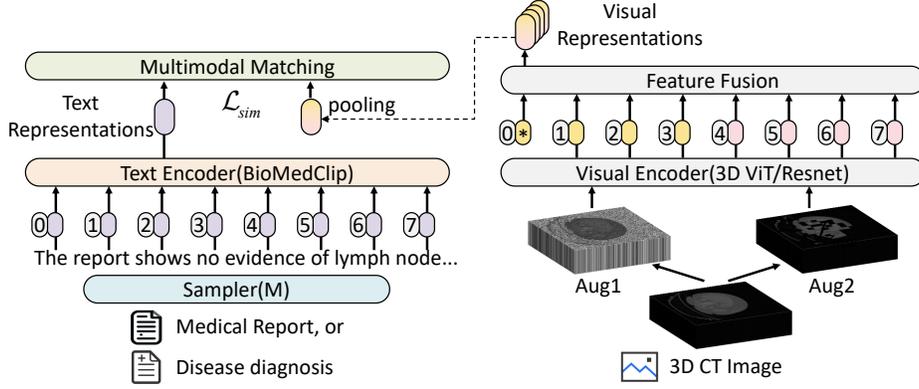
Statistic	Average	Median	Range
Image Width	529	512	514 ~ 710
Image Height	528	512	520 ~ 672
Number of Slices	279	224	101 ~ 670
Length of Report	104	97	7 ~ 260

**Fig. 2:** Sample data of BIMCV-R.**Fig. 3:** Left: Word Frequency Analysis. Right: World Cloud Analysis.

we have conducted a keyword frequency analysis of the radiological reports, with the results depicted in Figure 3.

### 3 Methodology

**Overview.** The comprehensive workflow of MedFinder is depicted in Figure 4. Initially, we sample the lengthy textual description  $T$  for manageable processing. For the text  $T$ , we employ a sampler  $S$  to randomly select  $M$  words, where  $M$  is typically set to 64, resulting in a sampled textual representation  $T' = S(T, M)$ . Subsequently, for a 3D medical image  $I$ , we apply a 3D image encoder  $F_{3D}$  to extract its feature representation  $Z = F_{3D}(I)$ . To enhance the model’s discriminative capability for medical image features, we introduce the concept of view consistency. Specifically, we apply two different data augmentation techniques  $A_1$  and  $A_2$  to the original 3D medical image  $I$ , generating two augmented views  $I_1$  and  $I_2$ . These views are processed through the 3D image encoder to obtain feature representations  $Z_1 = F_{3D}(I_1)$  and  $Z_2 = F_{3D}(I_2)$ . A view consistency loss  $L_{\text{cons}}$  ensures the consistency between these two feature representations, aiding the model in learning more robust image feature representations. Following feature extraction, we employ a feature discrimination loss  $L_{\text{dis}}$  to further refine the model, encouraging it to learn to distinguish between different medical image features. Finally, we integrate textual and image features, using a similarity metric  $S$  to compute the similarity score  $s(T', Z)$  between the text and image.



**Fig. 4:** An overview of our method, divided into textual feature extraction, visual feature extraction, and similarity matching.

**Textual Feature Extracting.** Consider a medical text  $T \in \mathbb{R}^{N \times M}$ , where  $N$  represents the maximum length of the text, and  $M$  denotes the size of the vocabulary. Our objective is to extract features from this text to facilitate subsequent image retrieval tasks.

Initially, we define a sampling function  $S : \mathbb{R}^{N \times M} \times \mathbb{N} \rightarrow \mathbb{R}^{L \times M}$ , which selects a continuous text segment of length  $L$  (here,  $L = 100$ ) from the input text  $T$ . This sampling process can be represented as  $T' = S(T, L)$ , where  $T'$  is the sampled text segment, a matrix of size  $L \times M$ , representing 100 selected words and their corresponding vocabulary indices from the original text.

Subsequently, we employ a pre-trained text encoder  $F_{\text{biomedClip}} : \mathbb{R}^{L \times M} \rightarrow \mathbb{R}^D$  to extract features from the text segment. The encoder  $F_{\text{biomedClip}}$  is a deep neural network that maps the text segment  $T'$  to a  $D$ -dimensional feature space. The feature extraction process can be represented as  $W = F_{\text{biomedClip}}(T')$ , where  $W$  is the feature vector of the text segment  $T'$ , a  $D$ -dimensional vector containing semantic information of the text segment. We freeze the text encoder during training.

**Visual Feature Extracting.** Consider a 3D medical image block  $I \in \mathbb{R}^{H \times W \times D}$ , where  $H$ ,  $W$ , and  $D$  represent the dimensions, our goal is to extract features for retrieval tasks. We use a resizing function  $R : \mathbb{R}^{H \times W \times D} \rightarrow \mathbb{R}^{H' \times W' \times D'}$ , to resize  $I$  to a standard size  $H' \times W' \times D'$ .

Subsequently, we apply two data augmentation techniques  $A_1$  and  $A_2$ , such as noise addition, rotation, and cutmix, to improve generalization. The augmented images are

$$I_{\text{aug1}} = A_1(I'), \quad I_{\text{aug2}} = A_2(I'), \quad (1)$$

where  $I_{\text{aug1}}$  and  $I_{\text{aug2}}$  are the enhanced images.

A visual encoder  $F_{3D} : \mathbb{R}^{H' \times W' \times D'} \rightarrow \mathbb{R}^{D''}$  extracts features from these images, producing

$$Z_{\text{aug1}} = F_{3D}(I_{\text{aug1}}), \quad Z_{\text{aug2}} = F_{3D}(I_{\text{aug2}}), \quad (2)$$

with  $Z_{\text{aug1}}$  and  $Z_{\text{aug2}}$  as the feature vectors.

To capture class-discriminative features, we obtain the CLS feature vector  $Z_{\text{cls}}$  from the encoder’s output using

$$Z_{\text{cls}} = \text{CLS}(F_{3D}(I')), \quad (3)$$

where CLS extracts class-discriminative features.

During training, we ensure semantic consistency of augmented features by computing the Mean Squared Error (MSE) loss between them. The MSE loss is articulated as

$$L_{\text{mse}} = \frac{1}{2} \|Z_{\text{aug1}} - Z_{\text{aug2}}\|_2^2, \quad (4)$$

where  $\|\cdot\|_2$  signifies the L2 norm.

After obtaining feature vectors  $Z_{\text{aug1}}$  and  $Z_{\text{aug2}}$ . We employ a cross-attention mechanism to merge these features, allowing the model to link the vectors for enhanced feature representation.

We define a cross-attention function  $\text{CrossAttn} : \mathbb{R}^{D''} \times \mathbb{R}^{D''} \rightarrow \mathbb{R}^{D''}$ , which inputs two feature vectors and outputs a fused feature vector as  $Z_{\text{fusion}} = \text{CrossAttn}(Z_{\text{aug1}}, Z_{\text{aug2}})$ .

The calculation of cross-attention is detailed as follows:

$$\text{Attention}(Z_{\text{aug1}}, Z_{\text{aug2}}) = \text{softmax} \left( \frac{Z_{\text{aug1}}^T Z_{\text{aug2}}}{\sqrt{D''} \|Z_{\text{aug1}}\| \|Z_{\text{aug2}}\|} \right), \quad (5)$$

$$Z_{\text{fusion}} = Z_{\text{aug1}} \odot \text{Attention}(Z_{\text{aug1}}, Z_{\text{aug2}}), \quad (6)$$

where  $\odot$  represents the element-wise multiplication (Hadamard product), and softmax is a normalization function ensuring the sum of attention weights equals 1. Thus, a feature vector  $Z_{\text{fusion}}$  that merges information from two views is derived.

**Similarity Matching.** Upon obtaining the visual features  $Z_{\text{fusion}}$  and text features  $W$ , we proceed with the following steps for retrieval.

Initially, we perform a pooling operation on the visual features  $Z_{\text{fusion}}$  to obtain a fixed-length feature vector, expressed as  $Z_{\text{pooled}} = \frac{1}{D''} \sum_{i=1}^{D''} Z_{\text{fusion},i}$ , where  $Z_{\text{pooled}}$  is the pooled visual feature vector, and  $D''$  represents the dimensionality of the visual features.

Next, the pooled visual features  $Z_{\text{pooled}}$  are matched with the text features  $W$  to compute their similarity score. The similarity score is calculated using cosine similarity, denoted as  $\text{Score}(W, Z_{\text{pooled}}) = \frac{W^T Z_{\text{pooled}}}{\|W\| \|Z_{\text{pooled}}\|}$ , where Score is the similarity score between the text features  $W$  and the pooled visual features  $Z_{\text{pooled}}$ .

During training, we have a set of positive pairs  $(W_i, Z_{\text{pooled},i})$  and negative pairs  $(W_j, Z_{\text{pooled},j})$ , the loss function is represented as:

$$L_{sim} = - \sum_{i,j} \log \frac{\exp(\text{Score}(W_i, Z_{\text{pooled},i}))}{\sum_k \exp(\text{Score}(W_k, Z_{\text{pooled},i}))}, \quad (7)$$

where  $L_{sim}$  is the loss function, encouraging the model to increase the similarity score of positive pairs while decreasing that of negative pairs [24].

The total loss can be written as

$$L_{\text{total}} = L_{\text{mse}} + \alpha L_{sim}, \quad (8)$$

where  $L_{\text{total}}$  is the total loss, combining MSE loss ( $L_{\text{mse}}$ ) and matching loss ( $L_{sim}$ ), weighted by  $\alpha$ .

## 4 Experiments and Results

**Data Splitting and Metrics.** In this study, we configured the BIMCV-R dataset into training, validation, and test sets, accounting for 70%, 10%, and 20% of the total data, respectively. To thoroughly evaluate the performance of our multimodal retrieval system, we selected Recall@K (R@K), Median Rank (MdR), and Mean Rank (MnR) as our primary evaluation metrics. These metrics provide a multidimensional reflection of the model’s effectiveness in retrieval tasks. Furthermore, to assess the performance of keyword-based image retrieval, we incorporated the Precision@K (P@K) metric, which precisely evaluates the proportion of correct items in the returned results. In our performance comparison experiments, we chose the CLIP4Clip [25] model from the video retrieval domain and the 3D-MIR [21] model from the medical multimodal retrieval domain as benchmarks to validate the effectiveness of our approach.

**Results.** We conducted experiments on both multimodal retrieval and keyword-based retrieval, where the CLIP4Clip and 3D-MIR models processed 3D medical volumes using frame-by-frame and average input methods, respectively, as described in the original studies. Our approach provided experimental results using ResNet-50 and ViT-base as backbones. The results of the multimodal retrieval experiments are shown in Table 2. These results indicate that our method outperforms the baselines. Specifically, CLIP4Clip’s performance was compromised due to the use of CLIP weights in the text processing component, which has a significant gap with medical descriptions, leading to inferior results. Despite significant improvements in our method, the lengthy nature of medical diagnostics and the high similarity among input images mean there is still a noticeable gap compared to image-text retrieval and text-video retrieval tasks.

To further assess the performance of our model, we employed a keyword-based retrieval task, which aligns more closely with practical scenarios than the rigid CT image-text pairings. Physicians often rely on specific keywords to search

**Table 2:** Results of multimodal retrieval, the best results are highlighted in **bold**.

Methods	Text to Image					Image to Text				
	R@1	↑R@5	↑R@10	MdR ↓	MnR ↓	R@1	↑R@5	↑R@10	MdR ↓	MnR ↓
CLIP4clip [25]	0.3	1.5	2.2	717.0	735.9	0.3	0.8	1.5	722.0	738.7
3D-MIR [21]	1.1	4.7	10.3	121.1	152.3	1.2	4.0	8.8	134.9	162.4
MedFinder (Resnet-50)	<b>2.8</b>	<b>8.7</b>	<b>20.3</b>	<b>68.9</b>	<b>81.3</b>	<b>2.9</b>	<b>8.8</b>	<b>19.7</b>	<b>71.2</b>	<b>80.7</b>
MedFinder (ViT-base)	2.7	<b>8.9</b>	<b>21.4</b>	<b>75.4</b>	80.1	2.7	<b>9.0</b>	<b>20.3</b>	<b>72.3</b>	<b>81.9</b>

**Table 3:** Results of keyword retrieval, the best results are highlighted in **bold**.

Methods	atelectasis			consolidation			adenopathy		
	P@20	P@50	P@100	P@20	P@50	P@100	P@20	P@50	P@100
CLIP4Clip	0.20	0.24	0.21	0.15	0.14	0.18	0.20	0.18	0.19
3D-MIR	0.35	0.34	0.31	0.40	0.36	0.39	0.35	0.34	0.37
MedFinder (Resnet-50)	<b>0.75</b>	<b>0.72</b>	<b>0.69</b>	<b>0.70</b>	<b>0.66</b>	<b>0.62</b>	<b>0.75</b>	<b>0.72</b>	<b>0.69</b>
MedFinder (ViT-base)	0.70	0.68	0.63	0.65	0.64	0.61	0.70	0.68	0.63

**Table 4:** Ablation study in text to image retrieval tasks using Resnet-50 backbone.

Text Sampler	$L_{mse}$	Text Encoder	R@1	R@5	R@10	MdR	MnR
	✓	Clip	1.2	5.1	11.3	120.3	148.4
✓		Clip	1.3	5.4	11.8	117.3	142.5
✓	✓	Clip	1.8	6.3	13.9	107.9	129.6
✓	✓	BiomedCLIP	2.8	8.7	20.3	68.9	81.3

for similar cases. Accordingly, we selected three typical diagnoses—atelectasis, consolidation, and adenopathy—as keywords to retrieve related 3D CT slices. As shown in Table 3, our approach achieves an accuracy of approximately 70% in retrieving relevant cases, surpassing other baseline methods. This capability can significantly reduce the workload of physicians in real diagnostic processes. Further experimental details and visualization results will be included in the supplementary materials, with the code to be made available open source upon the publication of the paper.

**Ablation Study.** Our ablation studies, detailed in Table 4, show that the Text Sampler and using pretrained text encoder weights significantly affect performance. This is likely due to the detailed nature of retrieval texts, where text sampling enables the network to learn more comprehensive descriptions. Additionally, a medical-specific text encoder extracts more relevant information, essential for our task.

## 5 Conclusion

In this study, we introduce the BIMCV-R dataset, aimed at establishing a benchmark for 3D medical image-text retrieval. With a carefully curated collection

of 8,069 3D CT volumes and their corresponding radiological reports, we offer a valuable resource to researchers. Our MedFinder demonstrates that effective information retrieval can be achieved in multimodal and keyword retrieval tasks by integrating advanced language models with image processing technologies. While our exploration is in its initial stages, it opens up a novel direction for the field, underscoring the potential for fruitful advancements in 3D medical image analysis technologies. We hope the BIMCV-R dataset will inspire further research to advance the development of 3D medical image analysis technologies.

## Disclosure of Interests

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

## References

1. Hao Yu, Haoyu Liu, Zhe Liu, Zeyu Wang, and Jiabin Jia. High-resolution conductivity reconstruction by electrical impedance tomography using structure-aware hybrid-fusion learning. *Computer methods and programs in biomedicine*, 243:107861, 2024.
2. Alym Abdulla, Wendy M Davis, Namali Ratnaweera, Elena Szefer, Brooke Ballantyne Scott, and Agnes YY Lee. A meta-analysis of case fatality rates of recurrent venous thromboembolism and major bleeding in patients with cancer. *Thrombosis and Haemostasis*, 120(04):702–713, 2020.
3. Zhixin Lai, Xuesheng Zhang, and Suiyao Chen. Adaptive ensembles of fine-tuned transformers for llm-generated text detection. *arXiv preprint arXiv:2403.13335*, 2024.
4. Chenxin Li, Xinyu Liu, Cheng Wang, Yifan Liu, Weihao Yu, Jing Shao, and Yixuan Yuan. Gtp-4-o: Modality-prompted heterogeneous graph learning for omni-modal biomedical representation. *arXiv preprint*, 2024.
5. Yi-Ling Cheng, Ruijie Zhang, Louis Tisinger, Salvatore Cali, Zhou Yu, Hua Yun Chen, and An Li. Characterization of microplastics in sediment using stereomicroscopy and laser direct infrared (ldir) spectroscopy. *Gondwana Research*, 108:22–30, 2022.
6. Neil R Smalheiser, Elena E Graetz, Zhou Yu, and Jing Wang. Effect size, sample size and power of forced swim test assays in mice: Guidelines for investigators to optimize reproducibility. *PLoS one*, 16(2):e0243668, 2021.
7. Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger R Roth, and Daguang Xu. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In *MICCAI Brainlesion Workshop*, 2021.
8. Yinda Chen, Haoyuan Shi, Xiaoyu Liu, Te Shi, Ruobing Zhang, Dong Liu, Zhiwei Xiong, and Feng Wu. Tokenunify: Scalable autoregressive visual pre-training with mixture token prediction. *arXiv preprint arXiv:2405.16847*, 2024.
9. Yipeng Sun, Linda-Sophie Schneider, Fuxin Fan, Mareike Thies, Mingxuan Gu, Siyuan Mei, Yuzhong Zhou, Siming Bayer, and Andreas Maier. Data-driven filter design in fbp: Transforming ct reconstruction with trainable fourier series. *arXiv preprint arXiv:2401.16039*, 2024.

10. Yinda Chen, Wei Huang, Shenglong Zhou, Qi Chen, and Zhiwei Xiong. Self-supervised neuron segmentation with multi-agent reinforcement learning. In *IJCAI*, 2023.
11. Runzhao Yang, Yinda Chen, Zhihong Zhang, Xiaoyu Liu, Zongren Li, Kunlun He, Zhiwei Xiong, Jinli Suo, and Qionghai Dai. Unicompress: Enhancing multi-data medical image compression with knowledge distillation. *arXiv preprint arXiv:2405.16850*, 2024.
12. Chenxin Li, Brandon Y Feng, Zhiwen Fan, Panwang Pan, and Zhangyang Wang. Steganerf: Embedding invisible information within neural radiance fields. In *ICCV*, 2023.
13. Shiyu Deng, Wei Huang, Chang Chen, Xueyang Fu, and Zhiwei Xiong. A unified deep learning framework for sstem image restoration. *IEEE Transactions on Medical Imaging*, 41(12):3734–3746, 2022.
14. Zudi Lin, Donglai Wei, Won-Dong Jang, Siyan Zhou, Xupeng Chen, Xueying Wang, Richard Schalek, Daniel Berger, Brian Matejek, Lee Kametsky, et al. Two stream active query suggestion for active learning in connectomics. In *ECCV*, 2020.
15. Weinan Dai, Jinglei Tao, Xu Yan, Zhenyuan Feng, and Jinkun Chen. Addressing unintended bias in toxicity detection: An lstm and attention-based approach. In *ICAICA*, 2023.
16. Wisdom Ikezogwo, Saygin Seyfioglu, Fatemeh Ghezloo, Dylan Geva, Fatwir Sheikh Mohammed, Pavan Kumar Anand, Ranjay Krishna, and Linda Shapiro. Quilt-1m: One million image-text pairs for histopathology. In *NeurIPS*, 2024.
17. Sheng Zhang, Yanbo Xu, Naoto Usuyama, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, Cliff Wong, et al. Large-scale domain-specific pretraining for biomedical vision-language processing. *arXiv preprint arXiv:2303.00915*, 2023.
18. Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. Medclip: Contrastive learning from unpaired medical images and text. In *EMNLP*, 2022.
19. Yinda Chen, Che Liu, Wei Huang, Sibao Cheng, Rossella Arcucci, and Zhiwei Xiong. Generative text-guided 3d vision-language pretraining for unified medical image segmentation. *arXiv preprint arXiv:2306.04811*, 2023.
20. Che Liu, Cheng Ouyang, Yinda Chen, Cesar César Quilodrán-Casas, Lei Ma, Jie Fu, Yike Guo, Anand Shah, Wenjia Bai, and Rossella Arcucci. T3d: Towards 3d medical image understanding through vision-language pre-training. *arXiv preprint arXiv:2312.01529*, 2023.
21. Asma Ben Abacha, Alberto Santamaria-Pang, Ho Hin Lee, Jameson Merkow, Qin Cai, Surya Teja Devarakonda, Abdullah Islam, Julia Gong, Matthew P Lungren, Thomas Lin, et al. 3d-mir: A benchmark and empirical study on 3d medical image retrieval in radiology. *arXiv preprint arXiv:2311.13752*, 2023.
22. Kathryn Wantlin, Chenwei Wu, Shih-Cheng Huang, Oishi Banerjee, Farah Dadabhoy, Veeral Vipin Mehta, Ryan Wonhee Han, Fang Cao, Raja R Narayan, Errol Colak, et al. Benchmd: A benchmark for modality-agnostic learning on medical images and sensors. *arXiv preprint arXiv:2304.08486*, 2023.
23. Maria De La Iglesia Vayá, Jose Manuel Saborit, Joaquim Angel Montell, Antonio Pertusa, Aurelia Bustos, Miguel Cazorla, Joaquin Galant, Xavier Barber, Domingo Orozco-Beltrán, Francisco García-García, et al. Bimcv covid-19+: a large annotated dataset of rx and ct images from covid-19 patients. *arXiv preprint arXiv:2006.01174*, 2020.
24. Zihan Li, Yunxiang Li, Qingde Li, Puyang Wang, Dazhou Guo, Le Lu, Dakai Jin, You Zhang, and Qingqi Hong. Lvit: language meets vision transformer in medical image segmentation. *IEEE transactions on medical imaging*, 2023.

25. Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022.