


WristMIR: Coarse-to-Fine Region-Aware Retrieval of Pediatric Wrist Radiographs with Radiology Report–Driven Learning

Mert Sonmezer^{†1} 

MERT.SONMEZER@METU.EDU.TR

¹ *Department of Computer Engineering, Middle East Technical University, Ankara, Turkiye 06800*

Serge Vasylechko^{2,3} 

SERGE.VASYLECHKO@CHILDRENS.HARVARD.EDU

Duygu Atasoy^{2,3}

DUYGU.ATASOY@CHILDRENS.HARVARD.EDU


² *Quantitative Intelligent Imaging Laboratory, Department of Radiology, Boston Children’s Hospital, Longwood Avenue, Boston, MA 02115*

³ *Harvard Medical School, Longwood Avenue, Boston, MA 02115*

Seyda Ertekin^{*1,4} 

SERTEKIN@METU.EDU.TR

⁴ *AI & Big Data Analytics Laboratory, METU-DTX Research & Application Center, Middle East Technical University, Ankara, Turkiye 06800*

Sila Kurugol^{*2,3} 

SILA.KURUGOL@CHILDRENS.HARVARD.EDU

Editors: Under Review for MIDL 2026

Abstract

Retrieving wrist radiographs with analogous fracture patterns is challenging because clinically important cues are subtle, highly localized and often obscured by overlapping anatomy or variable imaging views. Progress is further limited by the scarcity of large, well-annotated datasets for case-based medical image retrieval. We introduce **WristMIR**, a region-aware pediatric wrist radiograph retrieval framework that leverages dense radiology reports and bone-specific localization to learn fine-grained, clinically meaningful image representations without any manual image-level annotations. Using MEDGEMMA-based structured report mining to generate both global and region-level captions, together with pre-processed wrist images and bone-specific crops of the distal radius, distal ulna, and ulnar styloid, WristMIR jointly trains global and local contrastive encoders and performs a two-stage retrieval process: (1) coarse global matching to identify candidate exams, followed by (2) region-conditioned reranking aligned to a predefined anatomical bone region. WristMIR improves retrieval performance over strong vision-language baselines, raising image-to-text Recall@5 from 0.82 % to 9.35 %. Its embeddings also yield stronger fracture classification (AUROC 0.949, AUPRC 0.953). In region-aware evaluation, the two-stage design markedly improves retrieval-based fracture diagnosis, increasing mean F_1 from 0.568 to 0.753, and radiologists rate its retrieved cases as more clinically relevant, with mean scores rising from 3.36 to 4.35. These findings highlight the potential of anatomically guided retrieval to enhance diagnostic reasoning and support clinical decision-making in pediatric musculoskeletal imaging.

Keywords: Region-aware retrieval, Pediatric Wrist radiography, Contrastive pretraining.

[†]This work was conducted while the author was a research intern at Boston Children’s Hospital.

*Co-shared last authorship.

1. Introduction

Wrist fractures are among the most common pediatric injuries, and detecting and classifying them on radiographs is essential for appropriate management. Interpretation in pediatric patients, however, is challenging. Developmental anatomy, including open growth plates, variable ossification centers, and age-dependent changes in bone morphology, introduces substantial variability that can obscure subtle cortical disruptions and complicate distinguishing normal variants from true fractures. Because fracture appearance evolves with age and varies across patients, access to prior radiographs with similar injury patterns can provide valuable diagnostic context and support more consistent decisions.

Despite this need, large and richly labeled pediatric wrist datasets are scarce, as detailed annotations for fracture type, location, and severity requires expert pediatric radiologists and is too time-consuming to scale. This has motivated weakly and self-supervised approaches that leverage naturally occurring signals such as paired radiographs and reports. Contrastive language-image models (CLIP and its medical variants) use anatomical and pathological details in radiology reports to learn joint visual-textual representations without manual labels (Lin et al., 2023; Wang et al., 2022; Zhang et al., 2025a; Johnson et al., 2019). These models have advanced medical image retrieval, classification, and domain-specific vision-language modeling, offering scalability well suited to pediatric imaging.

Retrieving radiographs with analogous fracture patterns is especially valuable in pediatrics, where subtle differences often determine treatment decisions. Prior work in content-based image retrieval shows that similar-case retrieval can support diagnosis, reduce uncertainty, and enhance education (Choe et al., 2021; Qayyum et al., 2017; Müller et al., 2004; Dubey, 2022; Hu et al., 2022). However, retrieval remains challenging when clinically meaningful differences are highly localized and subtle (Şaban Öztürk, 2021; Yan et al., 2018; Zhong et al., 2021; Lee et al., 2023; Zhang et al., 2025b). Two radiographs may appear globally similar yet differ markedly in fracture type, severity, or region.

Global CLIP-style embeddings often fail to capture these fine-grained cues. Subtle findings such as cortical step-off, buckle deformation, physeal widening, or mild tilt/angulation may occupy small regions and are easily diluted by global pooling. Radiographic projections also introduce bone superimposition, causing small but clinically significant features to be lost in coarse representations. Effective retrieval, therefore, requires integrating global wrist context with fine-grained, anatomy-specific detail. Yet, manual annotation of such details is subjective, time-consuming, and difficult to scale (Abacha et al., 2023; Johnson et al., 2019; Nagy et al., 2022).

To address these challenges, we introduce **WristMIR**, a region-aware retrieval framework for pediatric wrist radiographs. WristMIR leverages dense radiology reports to extract anatomy-specific findings and treats sentence-level similarity as a proxy for image-level similarity. These textual signals are paired with bone-specific crops (distal radius, distal ulna, and ulnar styloid) to train a contrastive language-image model that learns both global wrist representations and localized bone embeddings. At inference, WristMIR performs two-stage retrieval: a global search to identify clinically plausible candidates, followed by region-conditioned reranking for the specified bone. Our main contributions are as follows:

1. Annotation-free supervision enabled by a scalable preprocessing pipeline that structures radiology reports into anatomy-specific findings and pairs them with detector gener-

ated, bone-level image crops, thereby eliminating the need for manual image annotation, a critical bottleneck in pediatric datasets.

2. A region-aware representation learning through a contrastive framework that aligns global wrist images with localized bone representations, enabling fine-grained discrimination of subtle fracture patterns that global embeddings fail to capture.

3. The WristMIR, a two-stage, region-conditioned retrieval that improves clinical relevance over global-only retrieval, by first ensuring global compatibility (laterality, morphology) and then refining retrieval based on the local anatomical region.

2. Data Preprocessing

Our preprocessing pipeline (Fig. 1) converts wrist radiographs and metadata into training-ready datasets for global and region-aware retrieval by (i) standardizing inputs, (ii) extracting structured anatomy-specific findings from radiology reports, and (iii) producing bone-specific crops and captions for CLIP-based training.

2.1. Data Sources

We retrospectively collected 7927 Posterior Anterior (PA) view pediatric wrist radiography examinations from our institution’s database under an approved IRB protocol. Each exam is paired with a free-text radiology report authored by board-certified pediatric radiologists. These radiograph-report pairs provide the sole supervision for WristMIR. No manual image-level annotations were used; all labels, regional descriptors, and fracture characteristics are automatically derived through our report-mining pipeline (see § 2.3). We focus on PA views because lateral and oblique projections stack the radius and ulna along the imaging axis, preventing reliable bone-level localization and obscuring fracture cues, making them unsuitable for region-aware retrieval. A full breakdown of dataset composition and fracture distribution is provided in Appendix B.

2.2. Wrist-ROI Extraction & Bone Detection

Detection and cropping. To isolate relevant anatomical regions, we use YOLOv11N-based detectors for wrist-ROI and bone-level localization (Fig. 1a) (Jocher et al., 2023). The ROI detector extracts the primary diagnostic area and removes non-informative regions, achieving a precision of 0.991 and recall of 0.975. The bone-level detector identifies the distal radius, distal ulna, and ulnar styloid with precision 0.996 and recall 0.981. These detected regions are cropped and paired with anatomy-specific captions to construct the image-text pairs for region-aware contrastive learning.

Enhancement and padding. To standardize appearance and improve visibility, all wrist-ROI crops are processed with Contrast-Limited Adaptive Histogram Equalization (CLAHE) (Pizer et al., 1990) and unsharp masking. These steps normalize contrast, enhance cortical boundaries, and reduce variation due to exposure or sharpness differences. Aspect ratios are preserved, and zero-padding is applied to the shorter dimension to produce square CLIP inputs while retaining the original geometry.

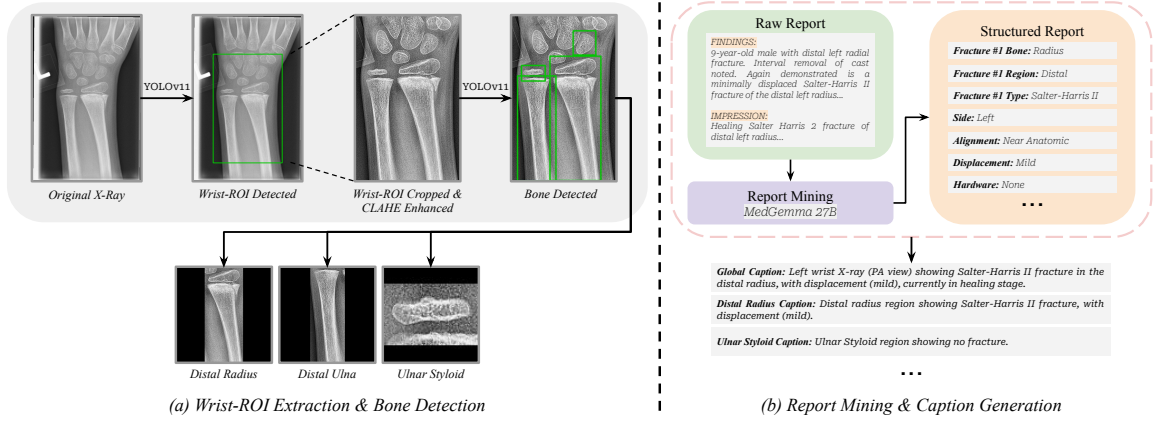


Figure 1: **Data preprocessing pipeline.** (a) YOLOv11 detector first identifies the wrist region of interest (ROI), applies CLAHE enhancement, and then localizes and crops three anatomical regions (distal radius, distal ulna, and ulnar styloid). (b) MEDGEMMA-27B converts each radiology report into a structured representation capturing anatomy-specific findings, which are then used to generate global exam-level captions and region-specific captions aligned with each bone crop.

2.3. Report Mining

VLM-assisted structuring. We transform free-text wrist radiology reports into structured representations using the medical VLM MEDGEMMA-27B (Sellersgren et al., 2025) (Fig. 1b). Reports are normalized to RADLEX terminology (Langlotz et al., 2006) and parsed into a JSON-like schema capturing anatomical entities, localized fracture descriptors, and global findings (Vasylechko et al., 2025). To reduce hallucinations and enforce schema adherence, we use chain-of-thought prompting with curated examples and validate outputs using PYDANTIC (Colvin, 2017). Post-processing canonicalizes terminology (e.g., "ulna styloid" → "ulnar styloid") and ensures consistency. Manual review of 250 cases shows a hallucination rate below 1%, indicating robust performance at scale.

Caption generation. Each structured report is converted into (i) global captions summarizing the entire wrist examination, including projection, alignment, fracture characteristics, and (ii) region-specific captions describing findings for each anatomical structures (distal radius, distal ulna, or ulnar styloid). Captions are generated via deterministic templates for consistency and serve as text inputs for contrastive training.

3. Methodology: WristMIR

WristMIR is a region-aware two-stage retrieval framework (Fig. 2) that learns multi-granular visual-textual representations of pediatric wrist radiographs. The method consists of two components: (i) contrastive learning of global and region-specific embeddings and (ii) region-aware retrieval guided by anatomical queries.

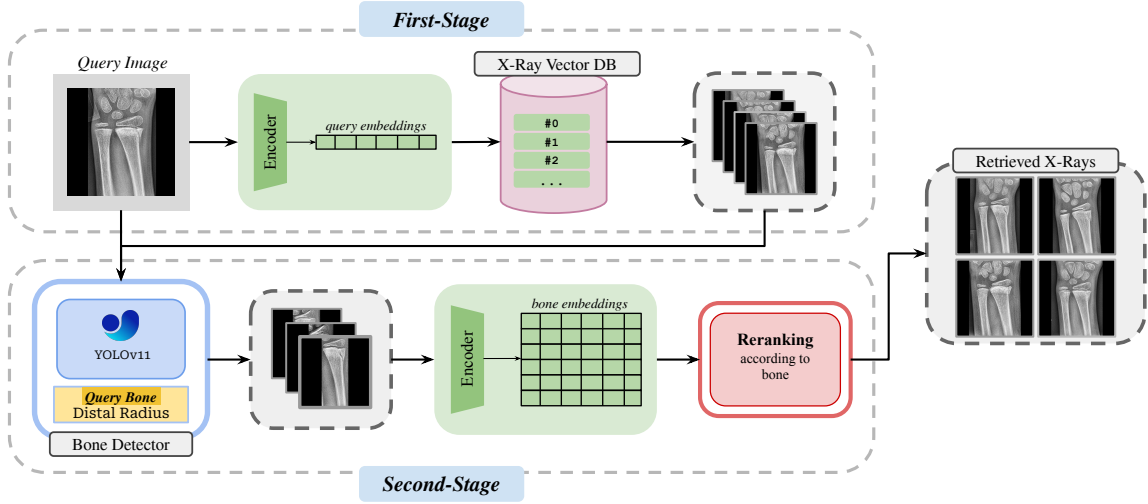


Figure 2: **WristMIR architecture.** A query wrist radiograph is encoded to generate both global and bone-level embeddings. A YOLOv11 detector identifies the relevant bone regions (e.g., distal radius, ulna, ulnar styloid). The global embedding is used to retrieve the top- k most similar exams from a precomputed database, after which these candidates are reranked using the region-specific embeddings to enable fine-grained, anatomy-aware retrieval.

3.1. Global and Region-Specific Contrastive Learning

Architecture. WristMIR adopts a dual-encoder CLIP framework built on BIOMEDCLIP (Zhang et al., 2025a). The image encoder $\Phi_{\text{img}}(\cdot)$ is a ViT-B/16 and the text encoder $\Phi_{\text{text}}(\cdot)$ is a transformer, both producing 512-dimensional embeddings. For each wrist radiograph I and caption R , derived from structured MEDGEMMA-27B reports, the encoders map inputs into a shared embedding space:

$$v = \Phi_{\text{img}}(I), \quad t = \Phi_{\text{text}}(R), \quad (1)$$

aligning paired image-text representations while pushing apart unpaired ones. Training includes both global wrist crops and localized regions (distal radius, distal ulna, ulnar styloid), enabling multi-granular representation learning.

Training objective. WristMIR is fine-tuned from BIOMEDCLIP by unfreezing the last eight image encoder blocks. Each image or crop I_i and caption R_i are encoded as $v_i = \Phi_{\text{img}}(I_i)$ and $t_i = \Phi_{\text{text}}(R_i)$, projected into a shared embedding space.

Because reports often describe normal ("no fracture") examinations, identical or highly similar captions occur frequently, producing ambiguous one-to-one supervision. Wrist radiographs also contain many near-duplicate cases (similar fracture types, healing stages, or regions), making strict single-positive contrastive learning unstable and poorly aligned with the data distribution.

To address this, WristMIR adopts a multi-positive contrastive loss in which all samples sharing the same caption are treated as valid positives. Formally, the symmetric CLIP loss is extended with a positive mask P_{ij} that distributes equal weight over all captions identical to R_i :

$$\mathcal{L} = -\frac{1}{B} \sum_{i=1}^B \left[\sum_j P_{ij} \log \frac{\exp(\langle v_i, t_j \rangle / \tau)}{\sum_k \exp(\langle v_i, t_k \rangle / \tau)} + \sum_j P_{ji} \log \frac{\exp(\langle t_i, v_j \rangle / \tau)}{\sum_k \exp(\langle t_i, v_k \rangle / \tau)} \right] \quad (2)$$

where τ is a learnable temperature. This formulation respects the clinical reality that many examinations convey equivalent semantic information, stabilizes contrastive alignment under limited caption diversity, and enables the model to focus on distinguishing genuinely different fracture patterns rather than arbitrarily separating semantically identical samples. Implementation details and training parameters are provided in Appendix C.

3.2. Region-Aware Retrieval

Two-stage retrieval. WristMIR employs a two-stage retrieval pipeline designed not only for efficiency but, more importantly, to enforce anatomical and view-level consistency before applying fine-grained region analysis. In the *global retrieval* stage, cosine similarity is computed between the query’s global embedding and all stored global embeddings:

$$S_g = \langle v_q, v_i \rangle. \quad (3)$$

This produces a candidate pool aligned with the query in coarse clinical properties such as laterality, projection, and wrist morphology. Restricting retrieval to these anatomically consistent cases prevents mismatches (e.g., opposite sides, different projections) and provides a stable basis for downstream region-level analysis.

The second stage performs *region-conditioned reranking*. Given a clinician-specified anatomical region (e.g., distal radius), similarity is computed between the corresponding region-level embeddings:

$$S_r = \langle v_{q,c}, v_{i,c} \rangle, \quad (4)$$

enabling the model to focus on subtle, localized morphological cues for the clinician-specified region. This ensures that fine-grained comparisons occur only among globally compatible candidates, improving clinical relevance.

Efficient region-level reranking via caching. Although region-level retrieval relies on YOLO-based bone detection, all detections are performed offline. Each exam’s distal radius, distal ulna, and ulnar styloid crops are extracted once and encoded using the WristMIR CLIP encoder. Their embeddings are stored and indexed, allowing reranking to operate entirely on cached region representations with no additional inference at query time.

4. Experiments

4.1. Baselines

We evaluate WristMIR on a pediatric wrist radiograph dataset paired with global and region-specific captions (§ 2). To assess the impact of region-aware learning, we compare against three strong medical vision-language baseline models: BIOMEDCLIP (Zhang et al., 2025a), PMC-CLIP (Lin et al., 2023), and MEDCLIP (Wang et al., 2022). These models represent state-of-the-art CLIP-style approaches pretrained on large biomedical corpora but lack anatomy-specific reasoning. All baselines are evaluated in a zero-shot setting.



Figure 3: **WristMIR attention maps.** The model consistently attends to fracture-relevant regions, focusing on localized morphological cues. Bounding boxes are shown only to guide visual interpretation of the fracture location and were not included in the dataset or were not used during CLIP training.

4.2. Experimental Setup and Metrics

We conducted all experiments on an evaluation dataset of 876 pediatric wrist images paired with clinical captions, ensuring no overlap with the training set. All retrieval experiments use a fixed two-stage setup: global retrieval selects the top-100 candidates, followed by region-conditioned reranking to produce the top-10 results. We assess model performance on both zero-shot classification and retrieval tasks using the following metrics:

Linear Probing: A logistic regression classifier is trained on frozen image embeddings for binary fracture detection. We report AUPRC, AUROC, and F_1 scores to assess the discriminative strength and clinical relevance of the learned visual representations.

Recall@ k : Measures the proportion of queries for which the correct caption appears in the top- k results, evaluating retrieval accuracy for both global and region-level queries.

Binary fracture & fracture classification matching: Compares retrieved cases against the query’s binary fracture labels and fracture categories (Salter–Harris, buckle, and transverse) assessing the benefit of region-aware reranking over single-stage retrieval.

Radiologist assessment: A board-certified pediatric radiologist blindly rates the top- k results for diagnostic relevance on a 5-point scale, five indicating the highest relevance.

Retrieval-based fracture diagnosis: Fracture presence in each region (distal radius, distal ulna, ulnar styloid) is predicted by aggregating labels from the top- k retrieved cases, and per-region F_1 scores are reported.

4.3. Unconditional Retrieval and Classification Performance

To assess WristMIR for pediatric wrist classification, we report unconditional image-to-text retrieval and binary fracture classification results in Table 1. WristMIR consistently outperforms all baseline medical CLIP models.

In image-to-text retrieval, WristMIR achieves higher performance across all k values. It attains a Recall@5 of 9.35 %, compared to 0.82 % for the strongest baseline, BioMEDCLIP. Given the extreme visual homogeneity of wrist radiographs, where distinct cases often

Table 1: **Comparison of WristMIR with medical CLIP baselines.** Retrieval and classification performance for WristMIR and medical CLIP baselines. WristMIR achieves the highest Recall@ k across all k values and outperforms all baselines in linear probing metrics, demonstrating stronger wrist-specific representations.

Method	Recall@ k (%) \uparrow				Linear Probing		
	$k = 5$	$k = 10$	$k = 50$	$k = 100$	AUROC \uparrow	AUPRC \uparrow	F_1 \uparrow
MEDCLIP	0.13	0.32	1.11	2.31	0.827	0.800	0.758
PMC-CLIP	0.44	0.92	3.71	7.73	0.891	0.890	0.822
BIOMEDCLIP	0.82	1.14	5.01	10.17	0.862	0.853	0.791
Our method	9.35	15.28	38.13	52.84	0.949	0.953	0.867

appear globally identical, this 10-times gain reflects meaningful extraction of fine-grained clinical signal. The advantage persists with larger pools: WristMIR reaches a Recall@100 of 52.84 % versus 10.17 % for BIOMEDCLIP. These large gaps highlight that broad biomedical pre-training alone does not capture the subtle, highly localized morphological cues required for pediatric wrist fractures, leading to uniformly low recall for standard medical CLIP models. The attention heatmaps in Fig. 3 further support this observation.

Because no existing CLIP model is specialized for wrist fractures, direct retrieval comparisons have inherent fairness limitations. To provide a more balanced evaluation of feature quality, we additionally perform linear probing. This setting assesses the adaptability of embeddings rather than their raw retrieval ability, offering a fairer comparison of representation strength. While baseline models achieve moderate AUPRC scores (0.800–0.890), their F_1 values remain low, indicating limited fracture sensitivity. WristMIR, in contrast, attains an AUROC of 0.949, an AUPRC of 0.953, and an F_1 of 0.867, demonstrating that its region-aware contrastive learning produces more discriminative embeddings.

Despite WristMIR’s improvements, absolute retrieval numbers remain modest. This reflects the intrinsic difficulty of the task: unlike classical computer vision retrieval benchmarks, where CLIP distinguishes semantically diverse objects, pediatric wrist radiographs differ only in subtle cortical or physeal abnormalities that are easily obscured by overlapping anatomy. The low absolute values thus reflect task complexity rather than model underperformance, underscoring the challenge of fracture-conditioned retrieval in highly homogeneous medical imaging domains.

4.4. Region-Aware Retrieval Evaluation

Table 2 and Fig. 4 compare our two-stage retrieval strategy with a single-stage global baseline. Across all regions and metrics, two-stage retrieval provides consistent improvements.

The gains are most pronounced for the ulnar styloid, which contains the subtlest fracture patterns. Here, the two-stage strategy improves Binary Fracture Matching from 0.374 to 0.522 and Fracture Classification Matching from 0.344 to 0.468, indicating that global embeddings alone miss subtle, region-specific cues, whereas region-aware refinement successfully emphasizes localized morphological cues.

The qualitative results further support this finding. For ulnar styloid, the two-stage system retrieves cases that closely match the query’s fracture type, orientation, and local patterns, whereas the single-stage baseline often surfaces images that are globally similar

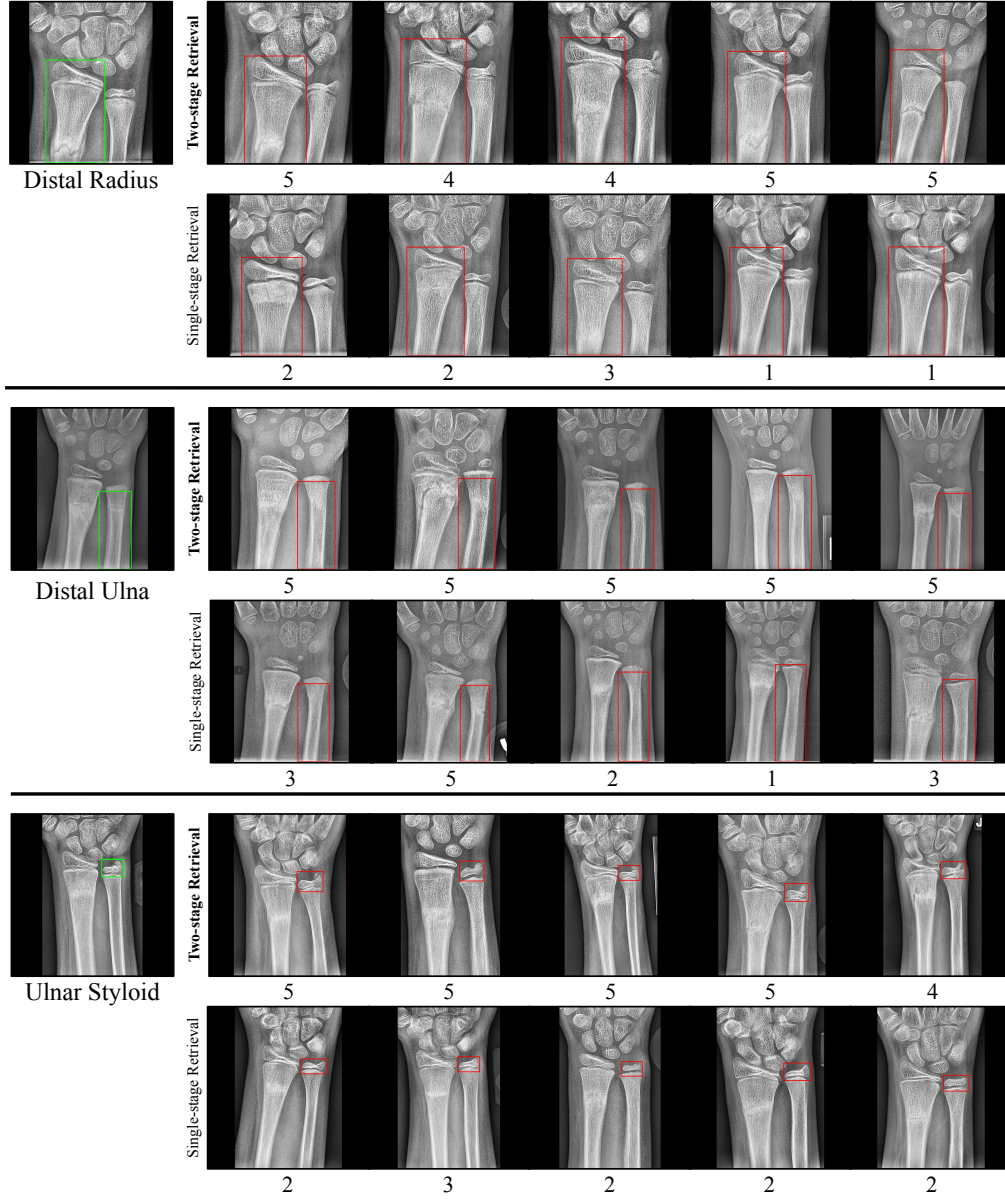


Figure 4: **Comparison of single- and two-stage retrieval.** Region-conditioned reranking retrieves cases that are anatomically and fracture-pattern aligned, whereas single-stage retrieval often surfaces globally similar but pathologically mismatched images. Numbers indicate scores assigned by a pediatric radiologist, showing higher and more clinically relevant retrieval for the proposed two-stage method.

in projection but pathologically mismatched. This demonstrates that reranking constrains retrieval to anatomically relevant evidence rather than broad global similarity.

Radiologist assessments follow the same trend. In a blinded 5-point evaluation on 30 query radiographs (20 retrieved images per query 10 per system), the two-stage method clearly outperforms the single-stage baseline. For ulnar styloid, scores rise from 3.16 to

Table 2: **Single- vs. two-stage retrieval.** Retrieval performance across anatomical regions. The two-stage method consistently improves binary fracture matching, fracture classification, fracture diagnosis, and radiologist-rated clinical relevance compared with single-stage retrieval.

Method	Distal Radius	Distal Ulna	Ulnar Styloid
<i>Binary Fracture Matching</i>			
Single-stage	0.824	0.458	0.374
Two-stage	0.864	0.666	0.522
<i>Fracture Classification Matching</i>			
Single-stage	0.534	0.372	0.344
Two-stage	0.578	0.542	0.468
<i>Radiologist Assesment</i>			
Single-stage	3.47	3.45	3.16
Two-stage	4.22	4.42	4.41
<i>Retrieval-based Fracture Diagnosis</i>			
Single-stage	0.897	0.574	0.233
Two-stage	0.934	0.771	0.554

4.41, indicating that region-conditioned retrieval produces images experts consider more diagnostically meaningful. We hypothesize that the single-stage model frequently retrieves visually similar but clinically irrelevant cases, whereas WristMIR’s reranking step elevates candidates with fracture patterns that accurately correspond to the queried anatomy.

4.5. Retrieval-based Fracture Diagnosis

We further assess whether region-aware retrieval supports diagnosis by aggregating the region-level fracture labels of the top- k retrieved exams. Table 1 reports per-region F_1 scores. The single-stage baseline performs well for the distal radius (0.894) but deteriorates for the distal ulna (0.574) and ulnar styloid (0.233). Incorporating region-conditioned reranking improves performance to 0.934, 0.771, and 0.554, respectively. The largest improvement occurs in the ulnar styloid, where fracture patterns are subtle and easily overshadowed by global appearance, reinforcing that anatomically targeted retrieval better captures subtle fracture patterns.

5. Conclusion

We introduced WristMIR, a region-aware retrieval framework for pediatric wrist radiographs that integrates structured report mining with joint global–local representation learning. By using both global and bone-level representations, WristMIR enables efficient, anatomy-specific retrieval of clinically analogous cases, improving diagnostic confidence, treatment planning, and education. Its two-stage retrieval design delivers fine-grained accuracy while remaining computationally practical for real-time use in clinical workflows. More broadly, WristMIR illustrates how structured radiology reports and anatomically grounded reasoning can overcome key limitations of global CLIP-style models for medical images, offering a scalable foundation for next-generation clinical decision support tools. Clinical deployment considerations and methodological limitations are discussed in Appendix A.

References

- 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, Noel C Codella, and Ivan Tarapov. 3d-mir: A benchmark and empirical study on 3d medical image retrieval in radiology, 2023. URL <https://arxiv.org/abs/2311.13752>.
- Jooae Choe, Hye Jeon Hwang, Joon Beom Seo, Sang Min Lee, Jihye Yun, Min-Ju Kim, Jewon Jeong, Youngsoo Lee, Kiok Jin, Rohee Park, Jihoon Kim, Howook Jeon, Namkug Kim, Jaeyoun Yi, Donghoon Yu, and Byeongsoo Kim. Content-based image retrieval by using deep learning for interstitial lung disease diagnosis with chest CT. *Radiology*, 302(1):187–197, October 2021.
- Samuel Colvin. Pydantic, 2017. URL <https://github.com/pydantic/pydantic>. Python data validation using Python type annotations.
- Shiv Ram Dubey. A decade survey of content based image retrieval using deep learning. *IEEE Trans. Cir. and Sys. for Video Technol.*, 32(5):2687–2704, May 2022. ISSN 1051-8215. doi: 10.1109/TCSVT.2021.3080920. URL <https://doi.org/10.1109/TCSVT.2021.3080920>.
- Brian Hu, Bhavan Vasu, and Anthony Hoogs. X-mir: Explainable medical image retrieval. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1544–1554, 2022. doi: 10.1109/WACV51458.2022.00161.
- Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Han-naneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021. URL <https://doi.org/10.5281/zenodo.5143773>.
- Glenn Jocher, Jing Qiu, and Ayush Chaurasia. Ultralytics YOLO, January 2023. URL <https://github.com/ultralytics/ultralytics>.
- Alistair E W Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren, Chih-Ying Deng, Roger G Mark, and Steven Horng. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific Data*, 6(1):317, December 2019.
- Curtis P Langlotz et al. Radlex: a new method for indexing online educational materials. *RadioGraphics*, 26(6):1595–1597, 2006.
- Ho Hin Lee, Alberto Santamaria-Pang, Jameson Merkow, Ozan Oktay, Fernando Pérez-García, Javier Alvarez-Valle, and Ivan Tarapov. Region-based contrastive pretraining for medical image retrieval with anatomic query, 2023. URL <https://arxiv.org/abs/2305.05598>.
- Weixiong Lin, Ziheng Zhao, Xiaoman Zhang, Chaoyi Wu, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-clip: Contrastive language-image pre-training using biomedical documents. In *Medical Image Computing and Computer Assisted Intervention – MICCAI*

- 2023: 26th International Conference, Vancouver, BC, Canada, October 8–12, 2023, *Proceedings, Part VIII*, page 525–536, Berlin, Heidelberg, 2023. Springer-Verlag. ISBN 978-3-031-43992-6. doi: 10.1007/978-3-031-43993-3_51. URL https://doi.org/10.1007/978-3-031-43993-3_51.
- Henning Müller, Nicolas Michoux, David Bandon, and Antoine Geissbuhler. A review of content-based image retrieval systems in medical applications-clinical benefits and future directions. *Int J Med Inform*, 73(1):1–23, February 2004.
- Eszter Nagy, Michael Janisch, Franko Hržić, Erich Sorantin, and Sebastian Tschauner. A pediatric wrist trauma x-ray dataset (GRAZPEDWRI-DX) for machine learning. *Scientific Data*, 9(1):222, May 2022.
- S.M. Pizer, R.E. Johnston, J.P. Ericksen, B.C. Yankaskas, and K.E. Muller. Contrast-limited adaptive histogram equalization: speed and effectiveness. In *[1990] Proceedings of the First Conference on Visualization in Biomedical Computing*, pages 337–345, 1990. doi: 10.1109/VBC.1990.109340.
- Adnan Qayyum, Syed Muhammad Anwar, Muhammad Awais, and Muhammad Majid. Medical image retrieval using deep convolutional neural network. *Neurocomputing*, 266: 8–20, 2017. ISSN 0925-2312. doi: <https://doi.org/10.1016/j.neucom.2017.05.025>. URL <https://www.sciencedirect.com/science/article/pii/S0925231217308445>.
- Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cían Hughes, Charles Lau, Justin Chen, Fereshteh Mahvar, Liron Yatziv, Tiffany Chen, Bram Sterling, Stefanie Anna Baby, Susanna Maria Baby, Jeremy Lai, Samuel Schmidgall, Lu Yang, Kejia Chen, Per Bjornsson, Shashir Reddy, Ryan Brush, Kenneth Philbrick, Mercy Asiedu, Ines Mezerreg, Howard Hu, Howard Yang, Richa Tiwari, Sunny Jansen, Preeti Singh, Yun Liu, Shekoofeh Azizi, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Riviere, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Elena Buchatskaya, Jean-Baptiste Alayrac, Dmitry Lepikhin, Vlad Feinberg, Sebastian Borgeaud, Alek Andreev, Cassidy Hardin, Robert Dadashi, Léonard Hussenot, Armand Joulin, Olivier Bachem, Yossi Matias, Katherine Chou, Avinatan Hassidim, Kavi Goel, Clement Farabet, Joelle Barral, Tris Warkentin, Jonathon Shlens, David Fleet, Victor Cotruta, Omar Sanseviero, Gus Martins, Phoebe Kirk, Anand Rao, Shravya Shetty, David F. Steiner, Can Kirmizibayrak, Rory Pilgrim, Daniel Golden, and Lin Yang. Medgemma technical report, 2025. URL <https://arxiv.org/abs/2507.05201>.
- Serge Didenko Vasylechko, Andy Tsai, Onur Afacan, and Sila Kurugol. Enhancing wrist fracture detection through LLM-powered data extraction and knowledge-based ensemble learning. In *Medical Imaging with Deep Learning*, 2025. URL <https://openreview.net/forum?id=Rd93NNCOnW>.

- Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. Medclip: Contrastive learning from unpaired medical images and text, 2022. URL <https://arxiv.org/abs/2210.10163>.
- Ke Yan, Xiaosong Wang, Le Lu, and Ronald M Summers. DeepLesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *J Med Imaging (Bellingham)*, 5(3):036501, July 2018.
- Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, Cliff Wong, Andrea Tupini, Yu Wang, Matt Mazzola, Swadheen Shukla, Lars Liden, Jianfeng Gao, Angela Crabtree, Brian Piening, Carlo Bifulco, Matthew P. Lungren, Tristan Naumann, Sheng Wang, and Hoifung Poon. Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs, 2025a. URL <https://arxiv.org/abs/2303.00915>.
- Tengfei Zhang, Ziheng Zhao, Chaoyi Wu, Xiao Zhou, Ya Zhang, Yanfeng Wang, and Weidi Xie. RadIR: A Scalable Framework for Multi-Grained Medical Image Retrieval via Radiology Report Mining . In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*, volume LNCS 15964. Springer Nature Switzerland, September 2025b.
- Aoxiao Zhong, Xiang Li, Dufan Wu, Hui Ren, Kyungsang Kim, Younggon Kim, Varun Buch, Nir Neumark, Bernardo Bizzo, Won Young Tak, Soo Young Park, Yu Rim Lee, Min Kyu Kang, Jung Gil Park, Byung Seok Kim, Woo Jin Chung, Ning Guo, Ittai Dayan, Manudeep K. Kalra, and Quanzheng Li. Deep metric learning-based image retrieval system for chest radiograph and its clinical applications in covid-19. *Medical Image Analysis*, 70: 101993, 2021. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2021.101993>. URL <https://www.sciencedirect.com/science/article/pii/S1361841521000396>.
- Şaban Öztürk. Class-driven content-based medical image retrieval using hash codes of deep features. *Biomedical Signal Processing and Control*, 68:102601, 2021. ISSN 1746-8094. doi: <https://doi.org/10.1016/j.bspc.2021.102601>. URL <https://www.sciencedirect.com/science/article/pii/S1746809421001981>.

Table 3: **Dataset composition and fracture distribution.** Summary of dataset used for WristMIR training and evaluation.

<i>Dataset Overview</i>	
Total examinations	7,540
Total region-level fractures	8,637
Normal cases (0 fractures)	2,209
Fracture-positive cases (≥ 1)	5,331
<i>Fractures per Case</i>	
0	2,209
1	2,264
2	2,869
3	159
4	37
5	2
<i>Fracture Location</i>	
Distal radius	5,369
Distal ulna	2,030
Ulnar styloid	1,238
<i>Fracture Morphology</i>	
Salter–Harris	1,621
Buckle	1,007
Transverse	1,924
Healing fractures	3,955

Appendix A. Limitations and Clinical Deployment

WristMIR relies on detector accuracy and structured report quality; failures in either can affect region-level retrieval. Hardware and casts may still bias the global encoder despite our region-aware design. The current evaluation is limited to a single institution and does not assess cross-domain generalization. From a clinical perspective, the two-stage retrieval pipeline is computationally efficient and compatible with workstation-integrated decision-support tools, but prospective validation is required before real-world deployment.

Appendix B. Dataset Details

Our pediatric wrist radiograph dataset contains 7540 examinations paired with structured radiology reports produced through our VLM-based report-mining pipeline. Across all studies, we identify 8637 region-level fractures. A total of 2209 cases are normal (0 fractures), while 5331 contain one or more fractures. Most fracture-positive cases involve one or two regions, with only a small number involving multiple sites.

Fractures most commonly involve the distal radius (5369 instances), followed by the distal ulna (2030) and ulnar styloid (1238). The dataset includes a diverse range of fracture types, notably 1621 Salter–Harris, 1007 buckle, and 1924 transverse fractures, alongside 3955 healing fractures, providing broad clinical variability useful for training region-specific embeddings. Table 3 includes the key statistics.




Figure 5: **WristMIR bone-level attention maps.** For each anatomical region (distal radius, distal ulna, and ulnar styloid), the model concentrates its attention on localized morphological cues that align with fracture-relevant structures. The dashed bounding boxes are included only to guide the reader by indicating the approximate fracture locations.

Appendix C. CLIP Training Details

WristMIR’s encoders are fine-tuned from BIOMEDCLIP (PUBMEDBERT-ViT-B/16) using the OPENCLIP framework (Ilharco et al., 2021) on 4 NVIDIA A100 GPUs. The model is initialized from `microsoft/BiomedCLIP-PubMedBERT_256-vit_base_patch16_224` on HUGGINGFACE, with a ViT-B/16 visual encoder and BIOMEDBERT text encoder projected into a shared 512-dimensional space. Consistent with our region-aware design, only the final eight visual transformer blocks are unfrozen, and standard CLIP preprocessing (RGB 224×224 , bicubic interpolation, mean/std normalization) is applied.

Training uses AdamW ($\text{lr}=1 \times 10^{-5}$, `weight decay`=0.01, $\beta_1/\beta_2=0.9/0.98$) with a cosine schedule and 50 warmup steps over 30 epochs. A global batch size of 2048 (512×4) is used with gradient clipping at 1.0. Because many radiographs share similar report-derived captions, a multi-positive contrastive loss is employed, treating all identical captions as valid positives to better match clinical supervision patterns.

Figure 5 shows additional attention maps from the fine-tuned image encoder, demonstrating WristMIR’s improved anatomical specificity across the distal radius, distal ulna, and ulnar styloid. After training, attention consistently localizes to fracture-relevant regions, cortical margins, metaphyseal interfaces, and subtle irregularities, illustrating how multi-granular supervision reshapes the embedding space toward clinically interpretable cues and supports the gains observed in region-aware retrieval.