

Location embedding based pairwise distance learning for fine-grained diagnosis of urinary stones

Qiangguo Jin^{1,2}, Jiapeng Huang¹, Changming Sun³, Hui Cui⁴, Ping Xuan⁵, Ran Su⁶,
Leyi Wei^{7,8}, Yu-Jie Wu^{9,10}, Chia-An Wu⁹, Henry B.L. Duh¹¹, and Yueh-Hsun
Lu^{9,12}

¹ School of Software, Northwestern Polytechnical University, Shaanxi, China

² Yangtze River Delta Research Institute of NPU, Suzhou, China

³ CSIRO Data61, Sydney, Australia

⁴ Department of Computer Science and Information Technology, La Trobe University, Melbourne, Australia

⁵ School of Engineering, Shantou University, Guangdong, China

⁶ College of Intelligence and Computing, Tianjin University, Tianjin, China

⁷ School of Information, Xiamen University, Xiamen, China

⁸ AIDD, Faculty of Applied Science, Macao Polytechnic University, Macao SAR, China

⁹ Department of Radiology, Shuang-Ho Hospital, Taipei Medical University, Taipei, Taiwan

¹⁰ Master of Public Health Program, National Yang Ming Chiao Tung University, Taipei, Taiwan

¹¹ School of Design, The Hong Kong Polytechnic University, Hong Kong

¹² Department of Radiology, School of Medicine, College of Medicine, Taipei Medical University, Taipei, Taiwan
20001@s.tmu.edu.tw

Abstract. The precise diagnosis of urinary stones is crucial for devising effective treatment strategies. The diagnostic process, however, is often complicated by the low contrast between stones and surrounding tissues, as well as the variability in stone locations across different patients. To address this issue, we propose a novel location embedding based pairwise distance learning network (LEPD-Net) that leverages low-dose abdominal X-ray imaging combined with location information for the fine-grained diagnosis of urinary stones. LEPD-Net enhances the representation of stone-related features through context-aware region enhancement, incorporates critical location knowledge via stone location embedding, and achieves recognition of fine-grained objects with our innovative fine-grained pairwise distance learning. Additionally, we have established an in-house dataset on urinary tract stones to demonstrate the effectiveness of our proposed approach. Comprehensive experiments conducted on this dataset reveal that our framework significantly surpasses existing state-of-the-art methods.

Keywords: Urinary stones diagnosis · Fine-grained classification · Abdominal X-ray image

1 Introduction

Urinary tract stones (UTS) are a primary cause of low back pain, presenting a significant diagnostic challenge in healthcare. Although rarely fatal, their prevalence is remarkably

high, affecting up to 10% of the population in developed countries [7,16]. Non-contrast computed tomography (NCCT) is established as the gold standard for urolithiasis diagnosis, demonstrating a diagnostic accuracy exceeding 92% [9]. However, due to the associated costs and radiation exposure from NCCT, low-dose abdominal X-rays, or KUB (Kidney, Ureter, and Bladder) radiography, are considered a viable, cost-effective initial diagnostic option for UTS diagnosis. Yet, the diagnostic accuracy of KUB exams falls significantly short of NCCT, with reported accuracies ranging from 44% to 77% [18]. Therefore, the development of automated and precise methods for diagnosing UTS from cost-effective KUB images is worth investigating.

However, diagnosing UTS from KUB images presents several distinct challenges. Firstly, distinguishing small-scale UTS from high-density objects like large intestinal feces, vascular calcifications, and phleboliths within KUB images proves to be intricate. Secondly, phleboliths, or small vein wall calcifications commonly found in the pelvis, occurring in approximately 40% of adults [13], are often located near the ureters, making it especially challenging to distinguish them from ureter stones. Hence, the significant variability in stone locations across patients greatly complicates the diagnostic task. Finally, the uneven distribution and limited sample sizes of various stone types (as shown in Fig. 2(a)), alongside the low contrast between them, pose significant obstacles to conducting a fine-grained diagnosis. Thus, the development of more specialized methods is needed to address these diagnostic challenges.

Deep learning has attracted substantial research interest across various fields of medical diagnosis [2,21,22,12]. For example, Zhang et al. [20] introduced an attention residual learning convolutional neural network (ARLNet) for skin lesion classification. Zhou et al. [21] employed a self-supervised pre-training approach using masked autoencoders (MAE) for medical image analysis tasks, which led to improved performance in chest X-ray disease classification tasks. The integration of additional information has been shown to considerably boost diagnostic accuracy. For instance, Wang et al. [19] incorporated external medical knowledge to guide their training process. Similarly, Han et al. [2] proposed a radiomics-guided transformer (RGT) that combined global image information with local radiomics-guided auxiliary information to enable accurate cardiopulmonary pathology localization and classification. Despite significant advancements in various medical imaging analysis tasks [19,20,2], UTS diagnosis from KUB images is still under study. The most relevant work to date, to our knowledge, is by Liu et al. [10] for the classification of kidney stones. However, the aforementioned methods exhibit certain limitations. Firstly, CNN-based approaches may not effectively incorporate localization information. Secondly, although attempts have been made to include location data, the subtle differences among stones, critical for accurate diagnosis, necessitate further investigation. These gaps highlight the need to develop methods that can capture fine-grained distinctions.

In this work, we propose the location embedding based pairwise distance learning network (LEPD-Net)¹ for the fine-grained diagnosis of urinary stones. The detailed architecture of LEPD-Net, as depicted in Fig. 1, consists of three pivotal components: the context-aware region enhancement (CRE) module, the stone location embedding (SLE) module, and the fine-grained pairwise distance learning (FPD) module. The CRE

¹ <https://github.com/BioMedIA-repo/LEPD-Net.git>




Fig. 1. The overall architecture of LEPD-Net. (a) Context-aware region enhancement module. (b) Stone location embedding module. (c) Fine-grained pairwise distance learning module. It is noted that the fine-grained pairwise distance learning module will be removed during inference.

module identifies stone-related regions, mitigating the potential interference from high-density objects and thus enhancing the model’s capability for image feature representation. Besides, the SLE module encodes textual location information, capturing crucial stone location knowledge. Finally, the FPD facilitates interactions between image pairs, enabling the recognition of fine-grained objects and improving the model’s ability to discriminate between fine-grained stones in scenarios of data scarcity.

The contributions of this study are threefold. Firstly, we propose a novel LEPD-Net that integrates location information into the image modality and strengthens the interactions between image pairs, thereby enhancing the discriminative power with limited data. Second, to overcome the lack of KUB datasets for stone diagnosis, we establish an in-house dataset of 414 patients from Shuang-Ho hospital. Third, our method demonstrates consistent performance improvements over recent state-of-the-art approaches applied to both medical and natural images.

2 Dataset and preprocessing

Dataset: This retrospective study was conducted at Shuang-Ho hospital, where anonymized KUB images from 414 patients treated for UTS were collected. The data collection was approved by the institutional review and ethical board. The resolutions of these images range from 864×924 to $3,311 \times 3,969$ pixels. The dataset contains 974 cases of stones, including 139 cases of ureter stones (US), 448 cases of phleboliths (PS), 296 renal stones (RS), 91 other types of calcifications (OC), and 410 randomly extracted non-stone (NS) patches. All the stone patches were labeled by urologists, with each bounding box centered on the stone, as illustrated in Fig. 2(b).




Fig. 2. (a) Visualization of location distributions of 974 stones with four different types. (b) An example of typical stone patch and the location information.

Preprocessing: We extract patches and location information from KUB images through a two-step process: (1) Stone patches are extracted based on the annotated bounding boxes. (2) The location information includes the coordinates and the organ region of a stone. The coordinates (Position X, Position Y) are the center of a stone’s bounding box. The organ region, i.e., right kidney (RK), left kidney (LK), bladder (BL), and other regions (OR), is identified by using a stone location map generated by all stone data points, as shown in Fig. 2. Following the collection of this information, the location data is then systematically encoded for further analysis.

3 Methodology

As illustrated in Fig. 1, the LEPD-Net consists of context-aware region enhancement (CRE), stone location embedding (SLE), and fine-grained pairwise distance learning (FPD). In the CRE module, global features are first extracted using a ResNet18 [4] backbone, facilitating an initial representation of the stone. Concurrently, a segmentation network utilizes the coarsely annotated bounding box around the stone as the ground truth to extract stone-related regions within the latent space. These regions are then combined with the global features through the CRE module, enhancing the model’s ability to capture comprehensive representations while emphasizing critical stone-specific details. In the SLE module, location information is carefully embedded and subsequently concatenated with the features enhanced by the CRE module. Finally, the FPD module employs pairwise distance learning to identify minor variations across different stone types, thereby enhancing the model’s ability for fine-grained discrimination.

3.1 Context-aware region enhancement (CRE) module

Given the challenge of distinguishing between stones and high-density objects with ambiguous characteristics, it is essential to approximate the localization of the stone

region for precise downstream diagnosis. Inspired by the concept that segmentation can enhance diagnostic accuracy [6], we adopt a coarse segmentation network and propose a context-aware region attention to enhance stone context regions, as shown in Fig. 1(a).

Context-aware region attention: Mathematically, for each preprocessed patch I_1 with its corresponding ground truth Y_1 , we concurrently process it through the coarse segmentation network and the pretrained ResNet18 backbone, yielding the stone-related features $\mathbf{x}'$ and global features $\mathbf{x}$, respectively. It is important to note that the ground truth for coarse segmentation is derived from the annotated bounding box, as illustrated in Fig. 2(b). Utilizing $\mathbf{x}'$ and $\mathbf{x}$, we employ a context-aware region attention mechanism to investigate the interaction between stone-related features and global stone features. Specifically, we project $\mathbf{x}'$ into a query vector $\mathbf{Q} = \text{LayerNorm}(\mathbf{x}')$, and $\mathbf{x}$ into both key $\mathbf{K} = \text{LayerNorm}(\mathbf{x})$ and value $\mathbf{V} = \text{LayerNorm}(\mathbf{x})$ vectors, where LayerNorm denotes the layer normalization function. The attention-augmented features $\mathbf{x}^{\text{att}}$ are computed as follows:

$$\mathbf{x}^{\text{att}} = \text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) + \mathbf{x}, \quad \text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V} \quad (1)$$

where $\frac{1}{\sqrt{d}}$ acts as a scaling factor, $\mathbf{K}^T$ is the transpose of $\mathbf{K}$. Following this computation, we employ a set of feature fusion functions to integrate those features, generating the region-enhanced features $\mathbf{x}^{\text{ff}} \in \mathbb{R}^{c \times h \times w}$. Here, c , h , and w represent the channel, height and width, respectively.

3.2 Stone location embedding (SLE) module

Given the variety of stones that may be present in the pelvis, diagnosing urinary tract stones from patch-based images is notably challenging, primarily due to the lack of global location information. We hypothesize that incorporating additional information, specifically the location and region of stones within the body, could significantly enhance visual prediction performance. This additional information, referred to as prior knowledge, is extracted during the preprocessing stage, as illustrated in Fig. 2(b).

The ‘Organ’ attribute is transformed into one-hot encoding matrices, and the numerical attributes, ‘Position X’ and ‘Position Y’, undergo normalization to fit within a $[0,1]$ range. The resulting 6-dimensional vector $\mathbf{y} \in \mathbb{R}^6$ is subsequently embedded through a sequence comprising a 1D convolutional layer, a batch normalization layer, and a sigmoid linear unit (SiLU) layer. This embedded feature is then expanded to match the dimensions of the region-enhanced features $\mathbf{x}^{\text{ff}} \in \mathbb{R}^{c \times h \times w}$, yielding the location-embedded features $\mathbf{y}^{\text{em}} \in \mathbb{R}^{c \times h \times w}$. Afterwards, the location-embedded features $\mathbf{y}^{\text{em}}$ and the region-enhanced features $\mathbf{x}^{\text{ff}}$ are concatenated and fused using a feature fusion module, as depicted in Fig. 1(b). The computation of the final feature $\mathbf{z}$ is as follows:

$$\mathbf{z} = \text{Fusion}(\text{Concat}(\mathbf{x}^{\text{ff}}, \mathbf{y}^{\text{em}})), \quad (2)$$

where Concat is the concatenation operation, the Fusion module consists of two 1×1 convolutional layers, a batch normalization layer, and a rectified linear unit (ReLU) layer.

3.3 Fine-grained pairwise distance learning (FPD) module

The challenge posed by class imbalance can result in the network overfitting to sample-specific features, particularly when distinguishing visually similar classes such as PS, US, and OC. To mitigate this issue, we propose a fine-grained pairwise distance learning module, which aims to reduce the proximity of images belonging to the same stone type while increasing the distance between those of different types.

During the training stage, we randomly select two images to create an image pair (I_1, I_2) and a corresponding ground truth pair (Y_1, Y_2) . These samples belonging to the same stone type are treated as positive sample pairs, i.e., $Y_1 = Y_2$, whereas the two samples from different types are regarded as negative sample pairs, i.e., $Y_1 \neq Y_2$. To explore the interactions between image pairs, we first compute high-level features $(\mathbf{z}, \hat{\mathbf{z}})$ for the pair (I_1, I_2) through the CRE and SLE modules. Subsequently, we employ a cross-attention mechanism on these pairwise features, which allows us to derive a pair of attention weights $(\mathbf{w}_1, \mathbf{w}_2)$. The attention weights are helpful in determining the relative importance of features within each image in relation to its counterpart. The process is formulated as follows:

$$\mathbf{w}_1 = \text{Softmax}\left(\frac{\hat{\mathbf{z}} \cdot \mathbf{z}}{\sqrt{d}}\right), \mathbf{w}_2 = \text{Softmax}\left(\frac{\mathbf{z} \cdot \hat{\mathbf{z}}}{\sqrt{d}}\right), \quad (3)$$

Finally, we minimize the distance D between the two samples. The distance between the high-level features of the image pair is defined as follows:

$$D = \begin{cases} \text{Distance}(\text{diag}(\mathbf{w}_1), \text{diag}(\mathbf{w}_2)), & \text{if } Y_1 = Y_2 \\ \max(0, 1 - \text{Distance}(\text{diag}(\mathbf{w}_1), \text{diag}(\mathbf{w}_2))), & \text{otherwise,} \end{cases} \quad (4)$$

where the cosine similarity serves as the distance measure, denoted by Distance , and diag returns a square diagonal matrix of weights. Optimizing this attention pair increases the difficulty of network training and reduces the overfitting to sample-specific features. Note that the FPD is only used for training and will be removed for inference without consuming extra computational cost.

3.4 Loss function

For the loss $(\mathcal{L}_{\text{dia}})$ of diagnosis, we integrate the label smoothing strategy [17] with the distribution of stones, aiming to reduce the potential label noise introduced during the annotation process and to prevent overfitting. For coarse segmentation, the Dice coefficient is employed as the loss function $(\mathcal{L}_{\text{seg}})$. Regarding the fine-grained pairwise distance learning, the loss $(\mathcal{L}_{\text{D}})$ is calculated by directly minimizing the cosine distance D . Consequently, the total loss $(\mathcal{L})$ is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{dia}} + \alpha \mathcal{L}_{\text{seg}} + \beta \mathcal{L}_{\text{D}} \quad (5)$$

where α and β are weighting coefficients that balance the contributions of each loss.

4 Experiments and results

4.1 Implementation details and evaluation measures

Our method is implemented in PyTorch using an NVIDIA RTX 3090 graphic card. To optimize our model, we use the Adam optimizer with a polynomial learning rate policy where the initial learning rate 1×10^{-4} is multiplied by $(1 - \frac{epoch}{total_epoch})^{power}$ with *power* as 0.9. We set the batch size to 16, with a total training duration of 200 epochs. The hyper-parameters α and β are empirically set to 0.1. Training patches are resized to 224×224 after applying different online augmentations, including ColorJitter, RandomGrayscale, GaussianBlur, and RandomHorizontalFlip, to improve data variety. It is noted that the encoder and decoder architecture employed for coarse segmentation can be substituted with other models, such as U-Net [15]. To ensure the robustness of our findings, we conduct extensive 5-fold cross-validation across all experiments.

Table 1. Urinary stones diagnosis performance of recently proposed methods.

Method	Acc (%)	Pre (%)	F1 (%)	Sen (%)
ResNet18 [4]	92.81±0.52	76.51±6.83	67.79±5.37	65.20±5.58
ARLNet [20]	93.48±0.35	76.17±2.40	71.38±1.59	68.98±2.03
MobileNetV3 [5]	93.28±0.34	78.34±2.46	70.97±1.54	67.73±2.22
Swin Transformer [11]	91.60±0.28	65.66±6.52	57.08±6.01	56.88±6.55
Conformer [14]	92.58±0.83	72.95±3.68	67.43±5.19	64.98±5.47
MAE [3]	88.66±1.01	60.42±11.33	49.52±7.45	50.44±6.35
RepLKNet-B [1]	90.88±0.64	65.89±1.67	59.84±1.20	57.68±1.65
SMPCConv-T [8]	93.21±0.53	74.67±3.12	69.85±3.29	68.16±3.62
LEPD-Net	94.98±0.46	82.42±1.92	78.58±2.49	76.90±3.42

We employ accuracy (Acc), precision (Pre), F1 score (F1), and sensitivity (Sen) as the evaluation metrics.

4.2 Performance comparison

To demonstrate the effectiveness of our LEPD-Net, we implement several state-of-the-art image classification methods, which include models based on traditional CNN architectures (ResNet18 [4], ARLNet [20], and MobileNetV3 [5]), transformer-based approaches (Swin Transformer [11], Conformer [14], and MAE [3]), and large kernel-based models (RepLKNet-B [1] and SMPCConv-T [8]). To ensure a fair comparison, all models are trained under identical settings.

The results presented in Table 1 show the superior performance of LEPD-Net. Notably, LEPD-Net achieves an outstanding Pre of 82.42%, which represents a 4.08% improvement over the second-best model, MobileNetV3. Furthermore, LEPD-Net achieves an F1 of 78.58%, surpassing MobileNetV3 by 7.61%, demonstrating the method’s balanced precision and sensitivity in classification tasks. Additionally, LEPD-Net demonstrates a significant enhancement in Sen with a score of 76.90%. These improvements highlight the enhanced ability of LEPD-Net to accurately diagnose urinary stones.




Fig. 3. (a) AUC-ROC and Precision-Recall curves for our LEPD-Net and other comparing methods. (b) Visual saliency maps for challenging-to-classify stones.

Table 2. Urinary stones diagnosis performance with ablation studies.

CRE	SLE	FPD	Acc (%)	Pre (%)	F1 (%)	Sen (%)
			92.81±0.52	76.51±6.83	67.79±5.37	65.20±5.58
✓			93.70±0.63	81.46±2.71	71.69±4.83	68.92±4.33
	✓		93.93±0.73	79.40±6.37	72.80±3.29	70.78±4.47
		✓	93.68±0.55	78.88±1.59	71.06±1.93	67.11±2.32
	✓	✓	94.13±0.47	79.64±3.59	71.46±1.71	68.23±1.80
✓	✓		94.36±0.50	81.08±3.51	74.08±3.02	71.96±4.91
✓		✓	94.18±0.46	80.55±1.82	72.99±3.13	70.48±3.43
✓	✓	✓	94.98±0.46	82.42±1.92	78.58±2.49	76.90±3.42

We further visualize the AUC-ROC and Precision-Recall curves to provide an intuitive demonstration of the enhanced performance. As depicted in Fig. 3(a), the proposed LEPD-Net achieves the highest AUC and the highest average precision (AP) scores, thereby validating the effectiveness of our proposed method.

Two observations can be drawn from the results. First, CNN-based methods outperform other approaches. This could be attributed to the relatively small size of the dataset, which may have a performance ceiling that CNN-based methods can more easily reach, while the more heavily parameterized transformer-based methods may be prone to overfitting. Second, the integration of CRE, SLE, and FPD modules in our LEPD-Net leads to superior performance metrics, particularly in the sensitivity score, which is crucial for identifying true positive cases.

4.3 Ablation study

We further perform ablation studies to evaluate the individual contributions of our newly proposed components. As shown in Table 2, the inclusion of the CRE and SLE modules results in a notable enhancement, increasing the average Acc by 1.55% over the baseline model. The addition of FPD elevates the average Sen score from 71.96% to 76.90%. This improvement is critical in a clinical setting, as it indicates the model's ability to correctly identify positive cases, thereby reducing the likelihood of false negatives.

We further employed class activation maps (CAMs) to visualize the class-specific discriminative regions, thereby validating the enhanced diagnostic capabilities of CRE and SLE components. As shown in Fig. 3(b), the integration of CRE and SLE enables the network to incorporate global information, significantly improving its precision in diagnosis.

5 Conclusions

In conclusion, we propose a location embedding based pairwise distance learning model for the fine-grained diagnosis of urinary stones. The proposed model consists of a context-aware region enhancement module, a stone location embedding module, and a fine-grained pairwise distance learning module for improving the feature representation ability of the network. Additionally, we construct an in-house annotated dataset for stone diagnosis. Comprehensive experiments demonstrate the superiority of the proposed method. Our future work includes the extension of our approach to other medical image diagnosis tasks.

Acknowledgments. This work was supported by the National Natural Science Foundation of China [Grant No. 62201460, No. 62222311, and No. 62322112], the Basic Research Programs of Taicang [Grant No. TC2023JC22], and the Fundamental Research Funds for the Central Universities.

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

References

1. Ding, X., Zhang, X., Han, J., Ding, G.: Scaling up your kernels to 31x31: Revisiting large kernel design in CNNs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11963–11975 (2022)
2. Han, Y., Holste, G., Ding, Y., Tewfik, A., Peng, Y., Wang, Z.: Radiomics-Guided Global-Local Transformer for Weakly Supervised Pathology Localization in Chest X-Rays. *IEEE Transactions on Medical Imaging* **42**(3), 750–761 (2022)
3. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16000–16009 (2022)
4. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
5. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for MobileNetv3. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1314–1324 (2019)
6. Jin, Q., Cui, H., Sun, C., Meng, Z., Su, R.: Cascade knowledge diffusion network for skin lesion diagnosis and segmentation. *Applied Soft Computing* **99**, 106881 (2021)
7. Khan, S.R., Pearle, M.S., Robertson, W.G., Gambaro, G., Canales, B.K., Doizi, S., Traxer, O., Tiselius, H.G.: Kidney stones. *Nature Reviews Disease Primers* **2**(1), 1–23 (2016)

8. Kim, S., Park, E.: Smpconv: Self-moving point representations for continuous convolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10289–10299 (2023)
9. Kobayashi, M., Ishioka, J., Matsuoka, Y., Fukuda, Y., Kohno, Y., Kawano, K., Morimoto, S., Muta, R., Fujiwara, M., Kawamura, N., et al.: Computer-aided diagnosis with a convolutional neural network algorithm for automated detection of urinary tract stones on plain X-ray. *BMC Urology* **21**(1), 1–10 (2021)
10. Liu, Y.Y., Huang, Z.H., Huang, K.W.: Deep learning model for computer-aided diagnosis of urolithiasis detection from kidney–ureter–bladder images. *Bioengineering* **9**(12), 811 (2022)
11. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
12. Lu, M., Wang, T., Zhu, H., Li, M.: HACL-Net: Hierarchical Attention and Contrastive Learning Network for MRI-Based Placenta Accreta Spectrum Diagnosis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 304–314. Springer (2023)
13. Luk, A.C.O., Cleaveland, P., Olson, L., Neilson, D., Srirangam, S.J.: Pelvic phlebolith: a trivial pursuit for the urologist? *Journal of Endourology* **31**(4), 342–347 (2017)
14. Peng, Z., Huang, W., Gu, S., Xie, L., Wang, Y., Jiao, J., Ye, Q.: Conformer: Local features coupling global representations for visual recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 367–376 (2021)
15. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer Assisted Intervention. pp. 234–241. Springer (2015)
16. Scales Jr, C.D., Smith, A.C., Hanley, J.M., Saigal, C.S., in America Project, U.D., et al.: Prevalence of kidney stones in the United States. *European Urology* **62**(1), 160–165 (2012)
17. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2818–2826 (2016)
18. Türk, C., Petřík, A., Sarica, K., Seitz, C., Skolarikos, A., Straub, M., Knoll, T.: EAU guidelines on diagnosis and conservative management of urolithiasis. *European Urology* **69**(3), 468–474 (2016)
19. Wang, K., Zhang, X., Huang, S.: KGZNet: Knowledge-guided deep zoom neural networks for thoracic disease classification. In: 2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). pp. 1396–1401. IEEE (2019)
20. Zhang, J., Xie, Y., Xia, Y., Shen, C.: Attention residual learning for skin lesion classification. *IEEE Transactions on Medical Imaging* **38**(9), 2092–2103 (2019)
21. Zhou, L., Liu, H., Bae, J., He, J., Samaras, D., Prasanna, P.: Self pre-training with masked autoencoders for medical image classification and segmentation. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–6. IEEE (2023)
22. Zhou, Y.J., Liu, W., Gao, Y., Xu, J., Lu, L., Duan, Y., Cheng, H., Jin, N., Man, X., Zhao, S., et al.: A Novel Multi-task Model Imitating Dermatologists for Accurate Differential Diagnosis of Skin Diseases in Clinical Images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 202–212. Springer (2023)