

High Cost, Low Trust? MSA-PNet Fixes Both for Medical Imaging

Dost Muhammad^{1,*}, Muhammad Salman² and Malika Bendeche³

¹CRT-AI and ADAPT Research Centres, School of Computer Science, University of Galway, Ireland

²Department of Software Engineering, University of Malakand, Pakistan

³ADAPT Research Centre, School of Computer Science, University of Galway, Ireland

Abstract

Deep learning (DL) has advanced medical ultrasound imaging by enabling automatic detection of subtle pathological features, notably in breast cancer diagnostics. However, mainstream architectures namely EfficientNetB7 and VGG19 remain limited by high computational complexity and poor model explainability, which hinders their integration into clinical workflows. This paper proposes MSA-PNet—a Multi-Scale Attention-Enhanced Prototype Network—designed to provide efficient, accurate, and explainable ultrasound-based disease prediction. MSA-PNet introduces an adaptive Feature Pyramid Network (FPN) with learnable scale-aware fusion to capture discriminative features across variable spatial resolutions. A spatial attention module selectively enhances diagnostically relevant regions, while an auxiliary ROI segmentation branch produces spatially aligned tumour masks, reinforcing both localisation accuracy and clinical coherence. For transparency, MSA-PNet incorporates a prototype-based explainability module optimized via metric learning, enabling predictions to be grounded in class-specific prototypical patterns and visual reasoning. Comprehensive evaluations on the BUSI ultrasound dataset demonstrate the superiority of MSA-PNet over state-of-the-art baselines. It achieves a mean Dice coefficient of 79.92%, Jaccard index of 81.07%, and a Hausdorff distance of 26.14, significantly outperforming both EfficientNetB7 and VGG19 across all metrics. Furthermore, MSA-PNet reduces inference time to 21.63 seconds—representing a 5× improvement in computational efficiency—making it highly suitable for real-time diagnostic deployment. By integrating multi-scale attention, prototype-based reasoning, and ROI-aware localisation into a unified architecture, MSA-PNet delivers not only robust diagnostic performance but also high-quality, clinically meaningful explanations. Its outputs exhibit strong alignment with expert-annotated tumour boundaries, thus enhancing trust, explainability, and applicability in high-stakes medical imaging environments. This framework represents a promising step toward the practical deployment of XAI systems in ultrasound-based diagnostics.

Keywords

Efficient Explainable AI, Attention-Guided Tumour Localisation, Prototype-Driven Explanation, Metric-Based Explainability, XAI in Healthcare

1. Introduction

Medical ultrasound imaging plays a critical role in the early detection and diagnosis of life-threatening diseases namely breast cancer, owing to its real-time capability, safety, and accessibility. However, the interpretation of ultrasound images remains a challenging task due to their inherently low contrast, high speckle noise, and variability in lesion morphology [1]. Recent advances in deep learning (DL) have significantly enhanced diagnostic performance in ultrasound imaging, enabling automated identification of subtle pathological features that may elude human observation[2].

While DL models namely EfficientNetB7 [3, 4] and VGG19 [5] have demonstrated remarkable success in medical imaging tasks, their deployment in real-time clinical workflows is constrained by two key limitations: high computational cost and lack of explainability [6, 7]. The computational demands of these models make them impractical for point-of-care settings and edge devices. More importantly, their

EXPLIMED 2025 - Second Workshop on Explainable Artificial Intelligence for the medical domain - 25-30 October 2025, Bologna, Italy

*Corresponding author.

✉ d.muhammad1@universityofgalway.ie (D. Muhammad); salmanuom04@gmail.com (M. Salman); malika.bendeche@universityofgalway.ie (M. Bendeche)

ORCID 0000-0001-9186-6392 (D. Muhammad); 0000-0009-2623-6736 (M. Salman); 0000-0003-0069-1860 (M. Bendeche)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

opaque decision-making processes undermine clinical trust, limiting their integration into diagnostic practice [8].

Explainable Artificial Intelligence (XAI) has emerged as a vital field to address these challenges. Techniques like Grad-CAM, LIME, and SHAP attempt to generate post-hoc visual explanations by highlighting salient input regions [9, 10]. However, these methods often produce diffuse or inconsistent heatmaps and are rarely aligned with annotated pathological structures, particularly in ultrasound images. Additionally, the field lacks standardized, quantitative metrics to rigorously assess explanation quality and reliability, further hindering adoption in clinical domains [8].

To address these limitations, we propose the Multi-Scale Attention-Enhanced Prototype Network (MSA-PNet)—a novel framework tailored for efficient, explainable, and clinically meaningful ultrasound diagnosis. MSA-PNet incorporates an adaptive Feature Pyramid Network (FPN) that dynamically fuses multi-resolution feature maps, guided by spatial attention modules to enhance the representation of diagnostically relevant regions. This allows the model to robustly capture heterogeneous lesion characteristics across multiple scales.

Moreover, MSA-PNet integrates a prototype-based explainability mechanism trained via metric learning, enabling the network to reason about new cases by comparing them to learned prototypical examples. This facilitates case-based explainability aligned with clinical intuition. To further strengthen explainability, we incorporate an explicit Region of Interest (ROI) localisation branch that highlights tumour boundaries with high spatial precision, providing visual justifications closely aligned with radiologist expectations.

The contributions of this work are summarized as follows:

- We present MSA-PNet, a novel architecture that combines adaptive multi-scale attention and prototype-driven reasoning to support efficient and explainable ultrasound-based diagnosis.
- Our model integrates an ROI segmentation head that provides fine-grained, interpretable tumour localisation, promoting transparency and clinical usability.
- MSA-PNet outperforms state-of-the-art models in prediction accuracy, explainability, and inference efficiency—achieving up to $5\times$ faster inference time compared to EfficientNetB7 and VGG19, while producing explanations that closely align with annotated clinical ground truth, as quantitatively validated using Dice coefficient [11], Jaccard index [12], and Hausdorff distance [13].

The remainder of this paper is structured as follows: Section 2 outlines the materials and methodology employed in this study. Section 3 presents the experimental results, followed by an in-depth discussion in Section 4. Finally, Section 5 concludes the paper by summarizing the key findings and outlining future research directions.

2. Materials and Methods

2.1. Dataset

This study utilises the publicly available Breast Ultrasound Images (BUSI) dataset [14], which comprises 780 greyscale ultrasound images categorized into three diagnostic classes: benign, malignant, and normal. Each image is paired with a corresponding binary mask annotated by clinical experts, indicating the precise tumour region. The dataset presents a realistic clinical scenario for evaluating prediction, segmentation, and explainability in breast ultrasound diagnosis.

2.2. Preprocessing

All images and masks were resized to a uniform resolution of 224×224 pixels to ensure consistency across input dimensions. Pixel intensity values were normalized to the range $[0, 1]$ to improve numerical stability during training [15, 16]. To mitigate overfitting and enhance model generalisation, common data augmentation techniques—including horizontal flips, rotations, and random crops—were applied

[17, 18]. The dataset was randomly split into training (80%) and validation (20%) sets for performance evaluation.

2.3. Implementation Environment

The experiments were conducted using Python, chosen for its versatility and comprehensive ecosystem of DL libraries. Model training and evaluation were performed on a computational system featuring an AMD Ryzen 7 5700X eight-core processor and a 16GB NVIDIA GeForce RTX 4080 GPU, providing the necessary computational power for efficient execution of DL workloads.

2.4. Proposed Approach: MSA-PNet

The proposed Multi-Scale Attention-Enhanced Prototype Network (MSA-PNet) is a unified DL architecture designed to enhance both the diagnostic accuracy and explainability of ultrasound-based tumour prediction. The network begins with a convolutional neural network (CNN) [19, 20] backbone to extract hierarchical feature representations from input images. These features are extracted across multiple layers, capturing increasingly abstract information, and are defined recursively as follows:

$$C_i = F_i(C_{i-1}), \quad i \in \{2, 3, 4, 5\} \quad (1)$$

where C_i represents the output feature map at the i -th stage, and $F_i(\cdot)$ denotes the corresponding convolutional operations.

To enable multi-scale analysis, the outputs $\{C_2, C_3, C_4, C_5\}$ are first passed through 1×1 lateral convolutions to project them into a common feature space:

$$L_i = \text{Conv}_{1 \times 1}(C_i), \quad i \in \{2, 3, 4, 5\} \quad (2)$$

These projected features are then fused using an adaptive top-down pathway with learned scalar weights w_i to modulate the contribution from each level:

$$P_i = w_i \cdot \text{Conv}_{3 \times 3}(L_i + \text{Upsample}(L_{i+1})) \quad (3)$$

This adaptive fusion enables the network to emphasize the most diagnostically informative scales for lesion detection.

To further refine the spatial context, a spatial attention mechanism is applied to the lowest-level fused feature map P_2 . An attention map A is computed using a 7×7 convolution followed by a sigmoid activation:

$$A = \sigma(\text{Conv}_{7 \times 7}(P_2)) \quad (4)$$

This attention map is used to modulate P_2 via element-wise multiplication, yielding an enhanced feature map P_2^* :

$$P_2^* = P_2 \odot A \quad (5)$$

This step ensures that subsequent computations are focused on spatial regions of diagnostic significance.

To provide explicit localisation of tumour regions, an auxiliary ROI segmentation head processes P_2^* through a stack of convolutional layers with ReLU and sigmoid activations. The predicted mask M is computed as:

$$M = \sigma(\text{Conv}_{1 \times 1}(\text{ReLU}(\text{Conv}_{3 \times 3}(P_2^*)))) \quad (6)$$

This mask highlights candidate lesion areas and contributes to both segmentation accuracy and interpretability.

Beyond segmentation, the model incorporates a prototype-based reasoning mechanism for transparent classification. A set of N learnable prototypes $\{p_j\}_{j=1}^N$ in \mathbb{R}^d are used to represent class-specific feature patterns. For a given input, the flattened feature representation $f \in \mathbb{R}^{d \times N_f}$ is compared to each prototype using cosine similarity:

$$\text{sim}(f, p_j) = \frac{f^\top p_j}{\|f\| \cdot \|p_j\|} \quad (7)$$

To reduce sensitivity to local noise and spatial variation, similarity scores are averaged across the top- k highest-scoring spatial locations:

$$S_j = \frac{1}{k} \sum_{i \in \text{top-}k} \text{sim}(f_i, p_j) \quad (8)$$

The resulting vector $S = [S_1, \dots, S_N]$ encodes the degree of alignment between the input and each prototype. Finally, class logits are computed using a linear classifier:

$$y = W_c \cdot S + b_c \quad (9)$$

where W_c and b_c are learnable parameters. This formulation not only provides discriminative predictions but also grounds them in semantically meaningful prototypes, facilitating clinical interpretation and visual traceability.

2.5. Evaluation Metrics for XAI

To evaluate the reliability of explainability generated heatmaps applied to the considered dataset, we employed three key metrics: DC, JI, and HD. These metrics will assess the alignment between generated explanation heatmaps, and ground truth annotations provided by radiologist, ensuring that the explanations are clinically relevant and explainable.

2.5.1. Dice Coefficient (DC)

The DC [21, 22] quantifies the overlap between two sets of pixels, typically the XAI-generated explanation (A) and the ground truth (B). It is defined as in Equation 10:

$$DC(A, B) = \frac{2|A \cap B|}{|A| + |B|} \quad (10)$$

Here, $|A \cap B|$ represents the number of overlapping pixels between the two sets, while $|A|$ and $|B|$ denote the total number of pixels in each set, respectively. The DC ranges from 0 (no overlap) to 1 (perfect overlap), making it an intuitive measure of similarity.

2.5.2. Jaccard Index (JI)

The JI [23, 12], also known as Intersection over Union (IoU), measures the ratio of the intersection to the union of two sets. It is given in Equation 11:

$$JI(A, B) = \frac{|A \cap B|}{|A \cup B|} \quad (11)$$

Here, $|A \cup B| = |A| + |B| - |A \cap B|$ represents the total number of unique pixels in either set. Like DC, the JI ranges from 0 to 1 but is more sensitive to differences in smaller regions, which is critical for detecting subtle discrepancies in medical images.

2.5.3. Hausdorff Distance (HD)

The HD [13] evaluates the maximum distance between the boundary points of two sets, providing a robust measure of how far the predicted explanation is from the ground truth. It is mathematically expressed in Equation 12:

$$HD(A, B) = \max \left\{ \sup_{a \in A} \inf_{b \in B} d(a, b), \sup_{b \in B} \inf_{a \in A} d(b, a) \right\} \quad (12)$$

In this equation, $d(a, b)$ represents the Euclidean distance between points a and b , \sup denotes the maximum distance over all points in the set, and \inf represents the minimum distance to the closest point in the other set. The Hausdorff Distance is particularly useful for evaluating boundary accuracy, capturing the worst-case discrepancy between the edges of predicted and ground truth regions.

3. Results and Analysis




Figure 1: Visual comparison of model-generated heatmaps from EfficientNetB7, VGG19, and MSA-PNet, highlighting more accurate tumour localisation achieved by MSA-PNet.

The diagnostic performance and explainability of the proposed MSA-PNet framework were assessed through a comprehensive comparative study involving two widely adopted DL architectures: EfficientNetB7 [3] and VGG19 [5]. Experiments were carried out using the BUSI ultrasound dataset, employing evaluation metrics namely the Dice coefficient, Jaccard index, Hausdorff distance, and overall computational inference time.

Quantitative results are summarized in Table 1. MSA-PNet achieved a significantly higher mean Dice coefficient ($79.92\% \pm 1.75\%$) and Jaccard index ($81.07\% \pm 1.53\%$) compared to EfficientNetB7 (48.27% Dice, 49.11% Jaccard) and VGG19 (46.43% Dice, 47.07% Jaccard). These metrics indicate superior spatial overlap between the predicted and ground truth tumor regions, demonstrating MSA-PNet’s enhanced ability to delineate lesion boundaries with greater precision.

Furthermore, the Hausdorff distance—measuring the maximum deviation between predicted and annotated contours—was substantially lower for MSA-PNet (26.14 ± 9.71) compared to EfficientNetB7 (237.52 ± 58.90) and VGG19 (256.82 ± 77.24), reflecting finer edge localisation and more anatomically consistent segmentations. Importantly, MSA-PNet required only 21.63 seconds to process the complete validation set, representing a five-fold improvement in inference efficiency over EfficientNetB7 (107.46 s) and VGG19 (108.43 s). These results confirm MSA-PNet’s computational scalability and clinical viability.

In addition to the numerical results, Figure 1 presents a qualitative comparison of the explanation heatmaps generated by each model. EfficientNetB7 and VGG19 tend to produce broad, diffuse saliency regions that often extend beyond the pathological areas, lacking precise anatomical alignment. In contrast, MSA-PNet consistently generates tight, well-localized activation maps that closely correspond to the tumour regions, thereby enhancing the clinical interpretability of its predictions. The improved spatial focus of MSA-PNet’s heatmaps stems from its built-in ROI segmentation branch and prototype-

DL Method	Dice coefficient (\uparrow)	Jaccard index (\uparrow)	Hausdorff distance (\downarrow)	Time (sec. (\downarrow))
EfficientNetB7	48.27 ± 3.72	49.11 ± 3.43	237.52 ± 58.90	107.46
VGG19	46.43 ± 4.30	47.07 ± 3.58	256.82 ± 77.24	108.43
MSA-PNet	79.92 ± 1.75	81.07 ± 1.53	26.14 ± 9.71	21.63

Table 1

Mean performance ± 1 standard deviation and computational time for EfficientNetB7, VGG19, and MSA-PNet across three ultrasound images (Figure 1).

based decision logic, both of which explicitly encode spatial alignment with known tumour patterns.

4. Discussion

The empirical performance of MSA-PNet can be directly attributed to a series of integrated architectural innovations, each addressing a known limitation of prior DL models in medical ultrasound imaging. First, traditional CNNs namely VGG19 or EfficientNetB7 operate with fixed-scale feature hierarchies, making them suboptimal for capturing ultrasound-specific pathology, which can present with highly variable lesion sizes, shapes, and textures. MSA-PNet overcomes this limitation through its adaptive Feature Pyramid Network (FPN), which employs learnable fusion weights (Equation 3) to dynamically prioritize spatial resolutions most relevant to the diagnostic task.

Second, generic CNN models often rely solely on the global receptive field of their deeper layers for context aggregation, which can result in the loss of fine-grained spatial cues essential for accurate tumour localisation. MSA-PNet’s spatial attention mechanism (Equation 4) reweights the low-level feature maps, forcing the model to selectively attend to diagnostically significant regions while suppressing background noise. This attention-driven localisation improves the focus and precision of both classification and segmentation outputs.

Third, the inclusion of an explicit ROI segmentation head (Equation 6) allows MSA-PNet to generate clinically actionable binary masks without requiring a separate post-processing pipeline. This auxiliary output not only strengthens the spatial alignment of predicted tumour regions but also contributes to overall model robustness during training via multi-task learning.

Most notably, MSA-PNet introduces a novel prototype-based explainability module (Equations 7–9). Unlike post-hoc methods such as Grad-CAM or LIME, which provide approximate and potentially unstable explanations, the prototype mechanism grounds each classification decision in a concrete, learned visual concept. These prototypes function as case-based reasoning anchors—comparing each input with stored representations of prototypical benign, malignant, or normal patterns—offering clinicians a more transparent and traceable decision path.

The reduction in Hausdorff distance (over 90% improvement compared to baselines) demonstrates that MSA-PNet excels not only in semantic prediction but also in precise boundary localisation. Moreover, the significant drop in inference time underscores the model’s suitability for real-time clinical deployment, an essential requirement for diagnostic imaging systems in fast-paced environments.

Together, these innovations enable MSA-PNet to bridge the long-standing gap between model performance and clinical trust. It delivers explainable predictions with quantitative and qualitative alignment to radiologist expectations—unlike traditional black-box models whose outputs are difficult to validate or interpret. These results strongly position MSA-PNet as a practical, accurate, and trustworthy solution for AI-assisted ultrasound diagnostics.

Beyond technical performance, the proposed MSA-PNet framework holds substantial clinical impact by addressing core barriers to AI adoption in medical imaging—namely trust, transparency, and deployment efficiency. By delivering fast, accurate, and explainable predictions, this approach has the potential to enhance radiologist workflow, support early cancer detection, and extend diagnostic capabilities to resource-constrained settings where expert availability is limited. Thus, MSA-PNet not only advances

the state of the art in XAI but also lays a practical foundation for integrating explainable DL models into real-world clinical practice.

5. Conclusion

This study introduced MSA-PNet, a novel DL framework tailored for explainable and efficient ultrasound-based disease diagnosis. By integrating adaptive multi-scale attention mechanisms, a dedicated ROI localisation module, and a prototype-based explainability layer, MSA-PNet addresses key limitations of conventional DL models in terms of computational efficiency, spatial precision, and clinical transparency. The model consistently outperformed state-of-the-art architectures namely EfficientNetB7 and VGG19, achieving higher diagnostic accuracy, more precise tumour localisation, and substantially reduced inference time.

In addition to its predictive capabilities, MSA-PNet provides visually and quantitatively robust explanation maps that closely align with expert-annotated ground truths. These explainable outputs enhance model trustworthiness and support clinical decision-making by offering clear justifications for each prediction. The model's computational efficiency further facilitates real-time application in clinical workflows, including deployment in low-resource settings. Overall, MSA-PNet not only advances the methodological landscape of XAI in medical imaging but also demonstrates strong potential for real-world integration in diagnostic radiology, improving both speed and reliability in ultrasound-based disease detection.

Acknowledgments

This research was supported by Taighde Éireann – Research Ireland under grant numbers GOIPG/2025/8471, 18/CRT/6223 (RI Centre for Research Training in Artificial Intelligence), 13/RC/2106/P_2 (ADAPT Centre), 13/RC/2094/P_2 (Lero Centre) and College of Science and Engineering, University of Galway. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Declaration on Generative AI

OpenAI's ChatGPT, Grammarly etc, were not used in the preparation of this manuscript. All content, analysis, and writing were entirely conceived, developed, and validated by the authors.

References

- [1] C. A. Duarte-Salazar, A. E. Castro-Ospina, M. A. Becerra, E. Delgado-Trejos, Speckle noise reduction in ultrasound images for improving the metrological evaluation of biomedical applications: an overview, *IEEE Access* 8 (2020) 15983–15999.
- [2] D. Muhammad, I. Ahmed, M. O. Ahmad, M. Bendeche, Randomized explainable machine learning models for efficient medical diagnosis, *IEEE Journal of Biomedical and Health Informatics* (2024) 1–10. doi:10.1109/JBHI.2024.3491593.
- [3] M. Tan, Q. Le, Efficientnet: Rethinking model scaling for convolutional neural networks, *arXiv preprint arXiv:1905.11946* (2019).
- [4] D. Muhammad, M. Salman, A. Keles, M. Bendeche, All diagnosis: Can efficiency and transparency coexist? an explainable deep learning approach, *Scientific Reports* 15 (2025) 12812.
- [5] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, *arXiv preprint arXiv:1409.1556* (2014).

- [6] D. Muhammad, M. Bendeche, Can ai be faster, accurate, and explainable? spikenet makes it happen, in: Annual Conference on Medical Image Understanding and Analysis, Springer, 2025, pp. 43–57.
- [7] M. Ali, D. Muhammad, O. I. Khalaf, R. Habib, Optimizing mobile cloud computing: A comparative analysis and innovative cost-efficient partitioning model, SN Computer Science 6 (2025) 1–25.
- [8] D. Muhammad, M. Bendeche, Unveiling the black box: A systematic review of explainable artificial intelligence in medical image analysis, Computational and Structural Biotechnology Journal 24 (2024) 542–560. URL: <https://www.sciencedirect.com/science/article/pii/S2001037024002642>. doi:<https://doi.org/10.1016/j.csbj.2024.08.005>.
- [9] M. I. Hossain, G. Zamzmi, P. R. Mouton, M. S. Salekin, Y. Sun, D. Goldgof, Explainable ai for medical data: current methods, limitations, and future directions, ACM Computing Surveys 57 (2025) 1–46.
- [10] D. Muhammad, A. Keles, M. Bendeche, Towards explainable deep learning in oncology: Integrating efficientnet-b7 with xai techniques for acute lymphoblastic leukaemia, in: Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)(Spain, 2024), 2024.
- [11] R. R. Shamir, Y. Duchin, J. Kim, G. Sapiro, N. Harel, Continuous dice coefficient: a method for evaluating probabilistic segmentations, arXiv preprint arXiv:1906.11031 (2019).
- [12] J. Bertels, T. Eelbode, M. Berman, D. Vandermeulen, F. Maes, R. Bisschops, M. B. Blaschko, Optimizing the dice score and jaccard index for medical image segmentation: Theory and practice, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22, Springer, 2019, pp. 92–100.
- [13] D. P. Huttenlocher, G. A. Klanderman, W. J. Rucklidge, Comparing images using the hausdorff distance, IEEE Transactions on pattern analysis and machine intelligence 15 (1993) 850–863.
- [14] W. Al-Dhabyani, M. Gomaa, H. Khaled, A. Fahmy, Dataset of breast ultrasound images, Data in brief 28 (2020) 104863.
- [15] S. Zheng, Y. Song, T. Leung, I. Goodfellow, Improving the robustness of deep neural networks via stability training, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 4480–4488.
- [16] D. Muhammad, I. Ahmad, M. I. Khalil, W. Khalil, M. O. Ahmad, A generalized deep learning approach to seismic activity prediction, Applied Sciences 13 (2023). URL: <https://www.mdpi.com/2076-3417/13/3/1598>. doi:10.3390/app13031598.
- [17] C. Shorten, T. M. Khoshgofaar, A survey on image data augmentation for deep learning, Journal of big data 6 (2019) 1–48.
- [18] D. Muhammad, null Rafiullah, M. Bendeche, Improving diagnostic trust: an explainable deep learning framework for genitourinary cancer prediction, IET Conference Proceedings 2024 (2024) 47–54. URL: <https://digital-library.theiet.org/doi/abs/10.1049/icp.2024.3275>. doi:10.1049/icp.2024.3275. arXiv:<https://digital-library.theiet.org/doi/pdf/10.1049/icp.2024.3275>.
- [19] K. O’shea, R. Nash, An introduction to convolutional neural networks, arXiv preprint arXiv:1511.08458 (2015).
- [20] D. Muhammad, I. Ahmed, K. Naveed, M. Bendeche, An explainable deep learning approach for stock market trend prediction, Heliyon 10 (2024).
- [21] L. R. Dice, Measures of the amount of ecologic association between species, Ecology 26 (1945) 297–302.
- [22] T. Sorensen, A method of establishing group of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on danish commons. i kommission hos e, Biologiske Skrifter, Kongelige Danske Videnskabernes Seleskab 5 (1948) 1–34.
- [23] T. Eelbode, J. Bertels, M. Berman, D. Vandermeulen, F. Maes, R. Bisschops, M. B. Blaschko, Optimization for medical image segmentation: theory and practice when evaluating with dice score or jaccard index, IEEE transactions on medical imaging 39 (2020) 3679–3690.