MedVis Suite: A Framework for MRI Visualization and U-Net-Based Bone Segmentation with In-Depth Evaluation

Mengyuan Liu[1], *Di* Zhang[1], *Yixiao* Chen[1], *Tianchou* Gong[1], *Hans* Kainz[3], *Seungmoon* Song[2], and *Jeongkyu* Lee[1,*]

[1]Northeastern University, San Jose, USA
[2]Northeastern University, Boston, USA
[3]University of Vienna, Austria

Abstract. This study introduces MedVis Suite, a framework developed to address key challenges in medical image analysis using MRI scans. MedVis Suite integrates advanced machine learning techniques, including U-Net-based segmentation model optimized for bone segmentation, and 3D reconstruction capabilities. An in-depth evaluation of a U-Net-based model for bone segmentation is performed across anatomical planes, optimizing both loss functions and image scales. The axial view showed the highest performance with a Dice score of 0.91 using the baseline model, while the combination of Dice loss and boundary loss produced the best results. MedVis Suite offers significant potential to enhance medical image analysis, improve segmentation accuracy, and provide more comprehensive visualizations for clinical use. Future research will focus on validating MedVis Suite across diverse datasets and clinical applications, with the integration of image preprocessing techniques and fine-tuning strategies to further enhance the U-Net-based segmentation model.

1 Introduction

Translating Magnetic Resonance Imaging (MRI) scans into precise, individualized models is a challenging endeavor due to the inherent variability in bone morphology and the difficulty of distinguishing bones from surrounding tissues. MRI's ability to visualize bone structures is often compromised by low proton density and short transverse relaxation times in bone tissue, leading to poor contrast and nonspecific low signal intensity in conventional sequences. Additionally, MRI scan faces challenges such as susceptibility to motion artifacts, long acquisition times, and susceptibility-related distortions, which can affect the accuracy and reliability of the images [1].

Our research addresses these challenges by developing MedVis Suite, an open-source framework designed to enhance the visualization, segmentation, and 3D reconstruction of medical images, with a particular focus on MRI scans. MedVis Suite integrates advanced machine learning techniques, including the U-Net architecture [2], specifically optimized for femur bone segmentation, automating the segmentation process and increasing the accuracy. Additionally, MedVis Suite offers flexible integration with other segmentation models

*Corresponding author: jeo.lee@northeastern.edu

as needed, making it uniquely adaptable to various clinical workflows. Its user-friendly interface enables clinicians to explore anatomical structures across multiple views intuitively, streamlining medical imaging and analysis. The suite also includes a robust visualization framework that minimizes the impact of artifacts, simplifies complex operations, and accelerates the overall analysis process. As an open-source tool, MedVis Suite encourages collaboration, empowering researchers and clinicians to access, improve, and tailor the tool to specialized medical imaging needs.

To address the relatively underexplored application of U-Net in bone segmentation and to enhance the capabilities of our U-Net-based deep learning model within the MedVis Suite, we conducted a thorough evaluation focusing on several key aspects. We systematically assessed the model's performance across various anatomical planes and explored different image scales and loss functions to optimize segmentation accuracy and model robustness. The insights gained from this evaluation are vital for advancing a more precise and reliable approach to bone segmentation in medical imaging and for further refining the model integrated into the MedVis Suite.

This research makes several key contributions:

- *Development of MedVis Suite*: creating MedVis Suite, integrating 3D rendering with U-Net models for bone segmentation. This platform offers multi-view DICOM image exploration for MRI scans.
- *In-depth evaluation of bone segmentation model*: evaluating U-Net performance across anatomical planes, image scales, and loss functions to optimize segmentation accuracy and model robustness.
- *Innovative data preparation*: introducing preprocessing methods to enhance U-Net's learning efficiency, improving segmentation accuracy and model generalizability across diverse datasets.

By focusing on these contributions, our research aims to advance the field of medical imaging and musculoskeletal modeling, ultimately contributing to improved clinical outcomes in orthopedics and rehabilitation medicine.

Section 2 presents a comprehensive review of prior work in medical imaging analysis. Building on this foundation, section 3 details the methodology employed in this study. With the methods established, section 4 provides an in-depth analysis of the experimental results, while section 5 concludes the paper, summarizing the key contributions and suggesting directions for future work in this area.

2 Related work

In this section, we review MRI scans and their analysis methods, emphasizing the need to evaluate the effectiveness of a U-Net-based model for bone segmentation from MRI scans. Additionally, we discuss related work on medical image visualization software, highlighting the necessity of developing more efficient segmentation models and automated visualization tools.

2.1 Magnetic resonance imaging

MRI scan leverages the principles of Nuclear Magnetic Resonance (NMR) to produce high-resolution images of internal organs and tissues. Through the application of a strong magnetic field and radiofrequency pulses, an MRI scan captures electromagnetic signals from various tissue types, including T1-weighted and FLAIR sequences, enabling detailed visualization of

anatomical structures and physiological processes. In contrast to X-rays, CT scans, and PET scans, MRI scans operate without the use of ionizing radiation [3].

2.2 MRI image analysis

In recent years, advancements have been achieved in the analysis of medical images from MRI scans. Various approaches have been explored to enhance segmentation performance on MRI scans, addressing both 2D and 3D scenarios. Techniques include algorithms for representing segmentation contours to address blurred boundaries, such as those employing Principal Component Analysis [4], Adaptive K-means Clustering algorithms [5], and Markov Random Field (MRF) segmentation methods [6] to reduce noise interference. With the advent of deep learning, which typically involves neural networks with more than five layers [7], and Convolutional Neural Networks (CNNs), significant advancements have been made in image segmentation. The U-Net architecture has emerged as a prominent solution for overcoming common challenges in medical imaging, including blurred boundaries, noise, and limited data. By integrating both low-level and high-level image features using skip connections between low-resolution and high-resolution feature maps, U-Net facilitates accurate delineation of anatomical structures [8]. Additionally, data augmentation techniques enable end-to-end training of U-Net, enhancing the model's robustness [9].

Although U-Net has gained widespread adoption in a variety of MRI image segmentation, such as liver and liver-tumor segmentation, brain and brain-tumor segmentation, optic disc segmentation, cell segmentation, lung segmentation, pulmonary nodule detection, and cardiac image segmentation [8], its application in bone segmentation remains relatively underexplored. This gap suggests a promising area for further investigation.

Our research addresses this gap by focusing on the evaluation and improvement of U-Net's performance in bone segmentation from MRI scans. By incorporating newly labeled datasets and optimizing model parameters, we aim to improve the accuracy and reliability of U-Net for this specific task.

2.3 Visualization software

Among numerous software tools designed for visualization of medical imaging format, 3D Slicer [10] has emerged as a leading open-source platform that offers a broad spectrum of functionalities for medical image processing including visualization and 3D reconstruction for medical imaging. Beyond visualization, 3D Slicer is equipped with a suite of segmentation tools that enable users to manually delineate regions of interest within the medical images. These tools are integral for tasks such as identifying and isolating specific tissues, organs, or pathological areas.

However, despite its robust feature set, 3D Slicer has notable limitations. As shown in Table 1, it offers only a moderate degree of customizability, meaning features cannot be easily tailored to specific needs. It is also not optimized for standard consumer hardware, requiring high-performance systems, which limits accessibility. Additionally, the interface is complex and may pose navigational challenges for users. While 3D Slicer allows for flexible manual segmentation, the process is time-consuming and prone to user variability, potentially leading to inconsistent results. These limitations highlight the need for more flexible, efficient, and user-friendly automated segmentation solutions.

3 Methodology

The MedVis Suite offers a comprehensive and streamlined workflow for medical image analysis, starting with a single-view MRI scan dataset and a pre-trained U-Net model. This

Table 1. Comparative analysis of MedVis and 3D Slicer.

Feature	**MRI Visualization Software**	**3D Slicer**
Extensibility	Highly customizable with integrated AI models, allowing for future enhancements	Moderate (extensions and plugins)
Programming Language	Java with Python integration	Primarily C++ with Python
Hardware Requirement	Optimized for standard hardware; doesn't require high-end systems	Recommended high-end systems for performance
User Interface	Intuitive and user-friendly	Complex with a steeper learning curve
Special Features	Auto-segmentation with Deep Learning and Musculoskeletal feature extractions	Advanced visualization

workflow incorporated five core features, i.e., 3D reconstruction, U-Net-based bone segmentation, MRI scan view transformation, MRI scan intensification, and data format conversion, that together provide an efficient environment for processing and enhancing MRI scans. This seamless integration of tools within the MedVis Suite sets the stage for effective medical image analysis.

Building on this workflow, this U-Net project aims to optimize bone segmentation by employing a comprehensive approach involving data preprocessing, training, and evaluation. The preprocessing phase includes experimentation with various input sizes and data augmentation techniques to enhance segmentation performance. During the training phase, different loss functions and MRI scan views are evaluated to determine their impact on model accuracy. Once the optimal model is identified, it is tested across multiple data sets to assess its robustness.

3.1 MedVis Suite workflow

The MedVis Suite begins with two primary inputs: one is a single-view MRI dataset, and the other one is a trained U-Net model. Once the inputs are provided, the MedVis Suite offers five core features that can be utilized to process and enhance the MRI scan:

- *3D reconstruction*: The suite offers 3D reconstruction capabilities, enabling the generation of three-dimensional models from 2D MRI scans and segmentation results.
- *U-Net-based bone segmentation*: This feature employs the U-Net model to perform precise segmentation of bone structures from the MRI scans. The result can be viewed either in a single-view plane or as a 3D model.
- *MRI scan view transformation*: This enables the transformation of MRI view across all anatomical planes.
- *MRI scan intensification*: To enhance visibility and contrast of the MRI images.
- *Data format conversion*: The suite converts data imported as DICOM and saves it in PNG and JPEG format. The converted dataset is easily shareable and compatible with other applications.

This modular and integrated workflow within the MedVis Suite allows users to perform a series of operations on MRI scans, from initial segmentation to final visualization and format conversion, all within a single, user-friendly environment.

3.2 U-Net baseline model

In recent studies, variations of the U-Net model have been developed to improve performance in different cases. In this study, the baseline model for in-depth evaluation utilizes a pre-trained ResNet50 as its backbone. This is because of its superior accuracy over ResNet34 and ResNet18, while effectively mitigating degradation issues [11]. Although ResNet101 and ResNet152 could potentially yield higher training accuracy, their evaluation is not pursued in this work due to constraints on computational resources.

Following the default input image size of ResNet50, the image size is set to 224×224 as the default size. For evaluation, we employ the Dice Similarity Coefficient (DSC) [12]. The DSC quantifies the pixel-wise concordance between a predicted segmentation and the corresponding ground truth. The coefficient is calculated as follows in (1):

$$\text{Dice}\,(A, B) = \frac{2|A \cap B|}{|A| + |B|} \tag{1}$$

, where A represents the predicted set of pixels and B denotes the ground truth. The DSC ranges from 0, indicating no overlap, to 1, signifying perfect overlap.

The baseline model utilizes a combined loss function, incorporating both Focal loss and the Dice coefficient since individual loss functions alone do not consistently deliver optimal results in segmentation tasks [13]. Therefore, combined approaches, e.g., Focal loss and Dice loss, are favored.

4 Experimental results

In this section, we highlight the benefits of the MedVis Suite and present the findings from a comprehensive evaluation of the U-Net-based model for bone segmentation.

4.1 MedVis Suite

MedVis Suite effectively visualizes DICOM images, allowing for multi-view exploration with user-friendly interfaces. Automated segmentation, powered by deep learning models such as U-Net, streamlines the segmentation process by allowing users to apply pre-trained models directly to DICOM data. Its 3D reconstruction capabilities were tested using MRI scans from various anatomical regions. The reconstructed 3D models accurately depict the structures within the MRI scan, preserving anatomical details. Users can interactively rotate, zoom, and explore the models within the suite, enabling a thorough examination of the reconstructed anatomy. The high-quality visualizations generated by MedVis suite demonstrate its effectiveness in providing reliable and detailed 3D reconstructions, significantly enhancing the diagnostic process.

4.2 Experiment setup for evaluation

The experiment preparation involves two separate processes: data preparation and training U-Net setup.

In the data preparation stage, data is prepared based on manual annotation for MRI scans. The original Nearly Raw Raster Data (NRRD) [14] files and their corresponding STL file format [15], which include the segmentation data, are converted into Digital Imaging and Communications in Medicine (DICOM) [16] image file format using 3D Slicer. The conversion process requires several meticulous steps to ensure accuracy and consistency. Initially, the corresponding NRRD and STL files are imported, organized, and refined using logical

Table 2. The number of images for each dataset information used in in-depth evaluation.

Dataset	Axial	Coronal	Sagittal
TD01_S1_MRI	585	166	516
TD02_S1_MRI	570	158	546
TD03_S1_MRI	572	149	548
TD04_S1_MRI	547	141	546
TD05_S1_MRI	549	155	545
TD06_S1_MRI	544	144	544
TD07_S1_MRI	544	128	544
TD08_S1_MRI	545	177	551
TD09_S1_MRI	547	156	569
TD10_S1_MRI	549	163	548

Table 3. Dataset sizes and Dice scores for MRI scan views (baseline model).

Views	Training Dataset Size	Dice Score
Coronal View	95	0.6975
Sagittal View	216	0.7374
Axial View	445	0.9095

operators to ensure that the resulting DICOM file correctly represents the segmentation of all relevant structures. The segmented data is then exported as a binary label map and converted into DICOM format, ensuring alignment with the original MRI scan across axial, coronal, and sagittal planes. The generated DICOM images are further converted into PNG images using MedVis Suite. Table 2 shows detailed information about the dataset used in our in-depth evaluation.

Below is the computing environment used in the training model setup:

- CUDA version: 11.8.
- GPU: NVIDIA GeForce RTX 2080 SUPER.
- PyTorch version: 1.13 or later.

4.3 In-depth evaluation results

In this section, we assess the performance of the U-Net model on dataset TD01_S1_MRI in Table 2 in terms of different views, input size, loss functions, and data augmentation. In addition, we evaluate the model across ten datasets introduced in Table 2 and compare them with the Dice scores to demonstrate its robustness.

4.3.1 Three views of MRI scan

First, we evaluate our baseline model, as detailed in Section 3.2, on three distinct MRI scan views: axial, sagittal, and coronal. Table 3 shows the number of images for each view after preprocessing, i.e., removing non-informative images, and its performance result using Dice

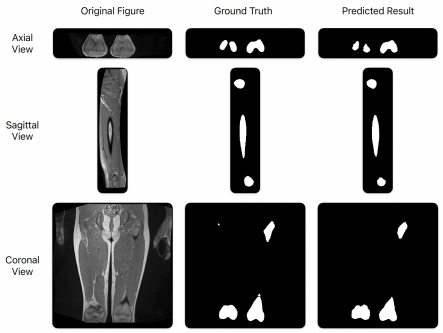

Figure 1. Examples of three different anatomical planes of MRI scans with corresponding predictions and ground truths.

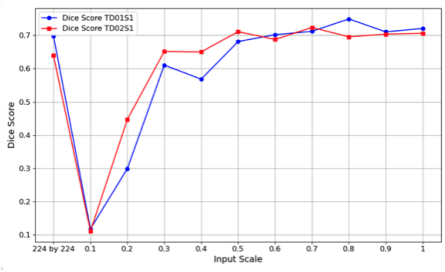

Figure 2. Evaluation results of input scale using Dice score for 2 datasets.

score. As shown in the table, the model performs best on the axial view, followed by the sagittal view, with the coronal view exhibiting the poorest results.

The smaller size of sagittal and coronal datasets, compared to axial datasets, leads to overfitting, which adversely affects model performance [17]. In cases with limited data and considerable good learning capability of ResNet50, the model fits the noise and specific patterns of the training set closely, resulting in poor generalization to new or unseen data. This phenomenon manifests as a decrease in evaluation metrics such as the Dice score, which measures the overlap between predicted and ground truth segmentations.

Figure 1 shows the example of segmentation results for each view, i.e., 1st row for axial view, 2nd row for sagittal view, and 3rd row for coronal view. In the third row, it is evident that the segmentation targets occupy a smaller portion of the images compared to the dominant background class. This imbalance means that the model is exposed to a disproportionate amount of background data relative to the structures of interest. Consequently, the model tends to become biased towards predicting the majority class, i.e., background due to its prevalence in the training dataset [18].

4.3.2 Different input size

In order to explore how varying the size of input images affects the performance of the U-Net model, input scale analysis is conducted for the coronal view of TD01_S1_MRI and TD02_S1_MRI in Table 1. By systematically adjusting the input scale on the chosen dataset,

Table 4. Baseline model Dice scores across different loss functions.

Loss Functions	**Dice Score**
Dice loss with Focal loss [21]	0.6975
Dice loss with HD loss [22]	0.6899
Dice loss with BD loss [23]	0.7036
Dice loss [24]	0.6654
Tversky loss [20]	0.5791
Cross-entropy [25]	0.6140
IoU (Jaccard) loss [26]	0.3655
Lovász loss [27]	0.6800

this experiment aims to identify the point at which the model achieves the best performance. The results are shown in Figure 2.

The Dice scores show significant improvement up to an input scale of 0.6 ~ 0.7, where near-peak performance is reached. Beyond this point, the margin of gain decreased, with scores stabilizing between 0.6 and 1.0. This suggests that using an input scale of 0.6 is optimal, as it maintains a high Dice score while reducing computational resource usage. Although larger input scales provide more detailed features, they also increase the risk of overfitting, where the model may start learning noise rather than meaningful patterns. Using a smaller input size like 0.6 allows for a good balance between accuracy and efficiency.

4.3.3 Different loss function

According to Ma et al. [13], compound loss functions, such as combinations of Dice loss with Focal loss, Dice loss with Hausdorff Distance (HD) loss, and Dice loss with Boundary (BD) loss, have demonstrated superior performance in medical image segmentation tasks compared to single loss functions. To determine the most effective approach for our bone segmentation, we evaluate these three compound loss functions. In addition, we assess the performance of commonly used single loss functions, including Dice loss, cross-entropy loss, Intersection over Union (IoU) loss, and its optimization loss function Lovász loss, which are known for their efficacy across various segmentation tasks [13, 19]. Furthermore, we incorporate Tversky loss, as it is specifically designed to handle class imbalance [20], which is a prevalent issue in the segmentation task. Table 4 shows the list of loss functions with evaluation results in terms of Dice score.

As demonstrated in Table 4, combined loss functions, i.e., Dice loss with Focal loss, Dice loss with HD loss, and Dice loss with BD loss, exhibit superior performance compared to single loss functions. The enhanced performance of Dice loss with HD loss and Dice loss with BD loss can be attributed to the integration of region-based and boundary-based loss components. This combination not only reduces the overall mismatch between the predicted segmentation and the ground truth but also introduces a boundary penalty, improving the accuracy of object delineation [13]. Figure 3 shows the examples of segmentation results for each loss function mentioned in Table 4.

Dice loss combined with Focal loss is particularly well-suited for addressing the challenges posed by imbalanced segmentation tasks. Given the significant class imbalance in the coronal view data, this combination enhances the model's ability to accurately segment minority classes, i.e., bone.

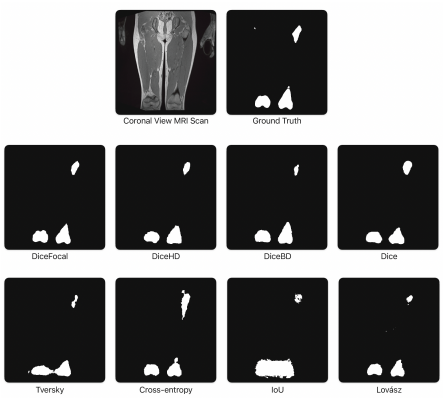

Figure 3. Segmentation results with different loss functions.

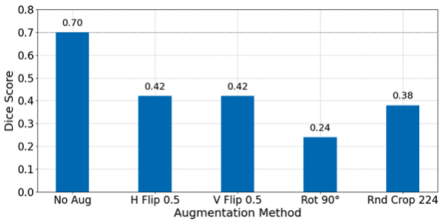

Figure 4. Evaluation scores of augmentation methods.

It is important to note that Dice loss, even without being combined with other loss functions, can yield desirable results, which explains its widespread use in medical image segmentation. Dice loss is particularly effective because it emphasizes the overlap between predicted and ground truth regions rather than individual pixel accuracy, making it well-suited for handling class imbalance in medical images.

4.3.4 Data augmentation

Data augmentation is crucial for improving the performance of deep learning models, especially when dealing with limited datasets. In the context of U-Net, data augmentation techniques such as horizontal flipping, vertical flipping, rotation, and random cropping are commonly employed. This test explores the different methods to examine its effects on segmentation accuracy. This will be particularly useful when acquiring large, professionally annotated datasets is challenging. The results are shown in Figure 4.

For the TD01_S1_MRI coronal view, the baseline Dice score without augmentation at a scale of 1.0 was 0.70. However, applying augmentation techniques such as random horizontal flip, vertical flip, or 90-degree rotation at 50% of the input did not improve the score; instead, it worsens. This suggests that these augmentations disrupted the spatial relationships in the data, which are crucial for accurate segmentation. Random cropping also led to a lower Dice score. These results highlight the importance of selecting appropriate data augmentation methods to enhance the performance and robustness of deep learning models in medical image segmentation.

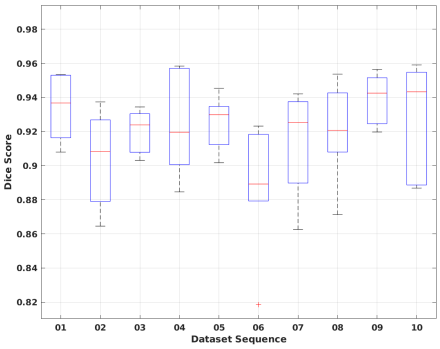

Figure 5. Dice scores for optimal settings across datasets.

4.3.5 Robustness and stability

To evaluate the robustness and stability of the U-Net model, we apply optimal parameters from previous experiments across ten distinct datasets, running each dataset ten times to verify stability. The optimal settings used include the axial view dataset, a 1.0 scale input size, no augmentation, and the use of Dice loss and BD loss functions. As illustrated in Figure 5, the model's segmentation accuracy remains robust even as it is applied to different MRI scans. While slight variability exists across datasets, the majority of scores stay well above 0.9, suggesting that the model handles inter-sample variability effectively. Such scores showcase the model's potential for accurate anatomical delineation in clinical imaging tasks.

5 Conclusion

This research presents MedVis Suite, an innovative framework designed to address key challenges in medical image analysis. Through the integration of advanced visualization techniques, U-Net-based segmentation models, and 3D reconstruction capabilities, MedVis Suite offers a comprehensive solution for analyzing MRI scans. This research makes several significant contributions to the field of medical imaging. (i) by incorporating 3D rendering with U-Net models, the MedVis Suite enables a comprehensive exploration of DICOM images, granting clinicians a detailed view of anatomical structures in MRI scans. (ii) the extensive assessment of bone segmentation, which involved testing different anatomical planes and experimenting with diverse input image scales, and loss functions, has led to improvements in segmentation precision and model reliability. and (iii) experimental results indicate that data augmentation techniques are not suitable for medical image preparation [28], suggesting the need for a more cautious approach in selecting data augmentation methods. While MedVis Suite shows promise in advancing medical image analysis, further research is needed to validate its performance across a broader range of datasets and clinical applications. To enhance femur bone segmentation performance, future investigations could focus on improving data quality through preprocessing techniques, alongside fine-tuning U-Net-based segmentation models to achieve higher accuracy.

This project was funded in part by the Northeastern TIER 1 seed grant.

References

[1] M.C. Florkow, K. Willemsen, V.V. Mascarenhas, E.H. Oei, M. van Stralen, P.R. Seevinck, Magnetic resonance imaging versus computed tomography for three-dimensional bone imaging of musculoskeletal pathologies: a review, Journal of Magnetic Resonance Imaging **56**, 11 (2022).

[2] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18* (Springer, 2015), pp. 234–241

[3] N.S. Punn, S. Agarwal, Modality specific u-net variants for biomedical image segmentation: a survey, Artificial Intelligence Review **55**, 5845 (2022).

[4] A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, W.E. Grimson, A. Willsky, A shape-based approach to the segmentation of medical imagery using level sets, IEEE transactions on medical imaging **22**, 137 (2003).

[5] T.N. Pappas, N.S. Jayant, An adaptive clustering algorithm for image segmentation, in *International Conference on Acoustics, Speech, and Signal Processing,* (IEEE, 1989), pp. 1667–1670

[6] K. Held, E.R. Kops, B.J. Krause, W.M. Wells, R. Kikinis, H.W. Muller-Gartner, Markov random field segmentation of brain mr images, IEEE transactions on medical imaging **16**, 878 (1997).

[7] Z. Akkus, A. Galimzianova, A. Hoogi, D.L. Rubin, B.J. Erickson, Deep learning for brain mri segmentation: state of the art and future directions, Journal of digital imaging **30**, 449 (2017).

[8] R. Wang, T. Lei, R. Cui, B. Zhang, H. Meng, A.K. Nandi, Medical image segmentation using deep learning: A survey, IET image processing **16**, 1243 (2022).

[9] X. Wei, T. Feng, Q. Huang, Q. Chen, C. Zuo, H. Ma, Deep learning-powered biomedical photoacoustic imaging, Neurocomputing p. 127207 (2023).

[10] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka et al., 3d slicer as an image computing platform for the quantitative imaging network, Magnetic resonance imaging **30**, 1323 (2012).

[11] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778

[12] L.R. Dice, Measures of the amount of ecologic association between species, Ecology **26**, 297 (1945).

[13] J. Ma, J. Chen, M. Ng, R. Huang, Y. Li, C. Li, X. Yang, A.L. Martel, Loss odyssey in medical image segmentation, Medical Image Analysis **71**, 102035 (2021). https://doi.org/10.1016/j.media.2021.102035

[14] Wikipedia contributors, Nrrd — Wikipedia, the free encyclopedia (2022), [Online; accessed 27-August-2024], https://en.wikipedia.org/w/index.php?title=Nrrd&oldid=1085264080

[15] Wikipedia contributors, Stl (file format) — Wikipedia, the free encyclopedia (2024), [Online; accessed 27-August-2024], https://en.wikipedia.org/w/index.php?title=STL_(file_format)&oldid=1230514338

[16] Wikipedia contributors, Dicom — Wikipedia, the free encyclopedia (2024), [Online; accessed 27-August-2024], https://en.wikipedia.org/w/index.php?title=DICOM&oldid=1236385307

[17] X. Liu, L. Song, S. Liu, Y. Zhang, A review of deep-learning-based medical image segmentation methods, Sustainability **13**, 1224 (2021).

[18] M. Yeung, E. Sala, C.B. Schönlieb, L. Rundo, Unified focal loss: Generalising dice and cross entropy-based losses to handle class imbalanced medical image segmentation, Computerized Medical Imaging and Graphics **95**, 102026 (2022).

[19] S. Minaee, Y. Boykov, F. Porikli, A. Plaza, N. Kehtarnavaz, D. Terzopoulos, Image segmentation using deep learning: A survey, IEEE transactions on pattern analysis and machine intelligence **44**, 3523 (2021).

[20] S.S.M. Salehi, D. Erdogmus, A. Gholipour, Tversky loss function for image segmentation using 3d fully convolutional deep networks (2017), 1706.05721, https://arxiv.org/abs/1706.05721

[21] W. Zhu, Y. Huang, L. Zeng, X. Chen, Y. Liu, Z. Qian, N. Du, W. Fan, X. Xie, Anatomynet: deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy, Medical physics **46**, 576 (2019).

[22] D. Karimi, S.E. Salcudean, Reducing the hausdorff distance in medical image segmentation with convolutional neural networks, IEEE Transactions on medical imaging **39**, 499 (2019).

[23] C. Wang, Y. Zhang, M. Cui, P. Ren, Y. Yang, X. Xie, X.S. Hua, H. Bao, W. Xu, Active boundary loss for semantic segmentation, Proceedings of the AAAI Conference on Artificial Intelligence **36**, 2397 (2022). 10.1609/aaai.v36i2.20139

[24] M. Drozdzal, E. Vorontsov, G. Chartrand, S. Kadoury, C. Pal, The Importance of Skip Connections in Biomedical Image Segmentation, in *Deep Learning and Data Labeling for Medical Applications*, edited by G. Carneiro, D. Mateus, L. Peter, A. Bradley, J.M.R.S. Tavares, V. Belagiannis, J.P. Papa, J.C. Nascimento, M. Loog, Z. Lu et al. (Springer International Publishing, Cham, 2016), pp. 179–187

[25] M. Thomas, A.T. Joy, Elements of information theory (Wiley-Interscience, 2006)

[26] 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**, 3679 (2020).

[27] M. Berman, A.R. Triki, M.B. Blaschko, The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2018), pp. 4413–4421

[28] D. Vasu, S. Song, H. Kainz, J. Lee, MRI Segmentation of Musculoskeletal Components Using U-Net: Preliminary Results, in *Proceedings of the 2024 14th International Conference on Bioscience, Biochemistry and Bioinformatics* (2024), pp. 30–35