

Pioneering Precision in Lumbar Spine MRI Segmentation with Advanced Deep Learning and Data Enhancement

Istiak Ahmed

North South University

Md. Tanzim Hossain

University of Erlangen-Nuremberg

Md. Zahirul Islam Nahid

North South University

Kazi Shahriar Sanjid

North South University

Md. Shakib Shahariar Junayed

North South University

M. Monir Uddin

`monir.uddin@northsouth.edu`

North South University

Mohammad Monirujjaman Khan

North South University

Research Article

Keywords: Lumbar Spine, Deep Learning, Segmentation, Class Imbalance, Medical Image Analysis

Posted Date: September 30th, 2024








DOI: <https://doi.org/10.21203/rs.3.rs-4999383/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Pioneering Precision in Lumbar Spine MRI Segmentation with Advanced Deep Learning and Data Enhancement

Istiak Ahmed ¹, Md. Tanzim Hossain ²,
Md. Zahirul Islam Nahid ^{1†}, Kazi Shahriar Sanjid ^{1†},
Md. Shakib Shahariar Junayed ^{1†}, M. Monir Uddin ^{3*},
Mohammad Monirujjaman Khan ¹

¹Department of Electrical & Computer Engineering, North South University, Dhaka, 1229, Bangladesh.

²Department of Computer Science, Friedrich-Alexander University, Erlangen, Nürnberg, 91054, Germany.

^{3*}Department of Mathematics & Physics, North South University, Dhaka, 1229, Bangladesh.

*Corresponding author(s). E-mail(s): monir.uddin@northsouth.edu;

Contributing authors: istiak.ahmed1@northsouth.edu;

tanzim.hossain@fau.de; zahirul.nahid@northsouth.edu;

kazi.sanjid@northsouth.edu; shakib.junayed@northsouth.edu;

monirujjaman.khan@northsouth.edu;

[†]These authors contributed equally to this work.

Abstract

This study presents an advanced approach to lumbar spine segmentation using deep learning techniques, focusing on addressing key challenges such as class imbalance and data preprocessing. Magnetic resonance imaging (MRI) scans of patients with low back pain are meticulously preprocessed to accurately represent three critical classes: vertebrae, spinal canal, and intervertebral discs (IVDs). By rectifying class inconsistencies in the data preprocessing stage, the fidelity of the training data is ensured. The modified U-Net model incorporates innovative architectural enhancements, including an upsample block with leaky Rectified Linear Units (ReLU) and Glorot uniform initializer, to mitigate common issues such as the dying ReLU problem and improve stability during training. Introducing

a custom combined loss function effectively tackles class imbalance, significantly improving segmentation accuracy. Evaluation using a comprehensive suite of metrics showcases the superior performance of this approach, outperforming existing methods and advancing the current techniques in lumbar spine segmentation. These findings hold significant advancements for enhanced lumbar spine MRI and segmentation diagnostic accuracy.

Keywords: Lumbar Spine, Deep Learning, Segmentation, Class Imbalance, Medical Image Analysis

1 Introduction

Accurate segmentation of lumbar spine structures in medical images is crucial for diagnosing and treating various spinal conditions, particularly those associated with low back pain [1, 2]. Low back pain is a prevalent medical issue that affects a significant portion of the population, leading to disability and reduced quality of life [3]. With the advent of advanced imaging technologies such as Magnetic resonance imaging (MRI), it is now possible to acquire high-resolution 3D images of the spine [4]. MRI is particularly valuable in this context because it provides detailed images of soft tissues, bones, and nerves [5], making it an essential tool for identifying pathological changes in the lumbar spine that contribute to low back pain [6].

Manual segmentation of these images is labor-intensive and time-consuming [7], necessitating the development of automated segmentation methods. Lumbar spine segmentation is important in diagnosing because it precisely delineates anatomical structures, such as vertebrae, the spinal canal, and intervertebral discs (IVDs) [1]. Accurate segmentation facilitates the identification of abnormalities [8–10], assists in planning surgical interventions [11, 12], and aids in monitoring disease progression and treatment response [13].

Deep learning, particularly Convolutional Neural Networks (CNNs), has shown great promise in this domain, offering substantial improvements in segmentation accuracy and efficiency [14]. Automating the segmentation process using deep learning reduces the burden on radiologists [15–18] and ensures consistency and repeatability in the analysis of medical images. This automation is crucial for handling the large data generated by MRI scans, allowing for quicker and more accurate diagnoses [15].

This paper focuses on the segmentation of lumbar spine structures, specifically targeting four classes: background, vertebrae, spinal canal, and IVDs. These classes are critical for a comprehensive spine analysis and any pathological changes. By automating the segmentation process, radiologists can allocate more time to critical diagnostic tasks [16], reduce human error, and increase throughput in clinical settings [19]. Automated segmentation supports radiologists by providing quantitative measures that can aid in clinical decision-making [17, 20].

Despite significant advancements in deep learning, challenges such as class imbalance and the complexity of spine anatomy remain significant hurdles. Addressing these

challenges effectively is essential for developing robust and reliable segmentation models [21, 22]. Related works in lumbar spine segmentation have utilized machine learning and deep learning approaches. However, many of these methods struggle with class imbalance and do not adequately handle the complex anatomy of the lumbar spine.

Using the SPIDER dataset [23], a well-established and publicly available dataset for lumbar spine segmentation in MR images, provides a realistic representation of clinical conditions with MRI scans from patients with low back pain. Initial data preprocessing revealed several issues, including missing vertebrae in 2D slices and an incorrect classification scheme, which initially resulted in 16 classes instead of the intended four.

To overcome these challenges, a comprehensive data preprocessing pipeline was implemented. First, 2D PNG images were extracted from the 3D MRI MetaImage Medical Format (mha) files. The missing vertebrae issue was identified and rectified during this process, ensuring that each 2D image accurately represented the four target classes. Rigorous data filtration techniques were applied to exclude images with fewer than four classes and those exhibiting an excessive class imbalance ratio (greater than 55%). This step was critical to ensure the quality and reliability of the training data.

A modified U-Net model was employed for the segmentation task, known for its efficacy in medical image segmentation [24]. Modifications included the integration of an upsample block with leaky Rectified Linear Units (ReLU) ($\alpha=0.1$) [25, 26] to mitigate the dying ReLU [27] problem. In this common issue, neurons become inactive and cease learning. The Glorot uniform initializer was also utilized for kernel weights to stabilize the training process [28, 29]. A custom combined loss function was devised to address the class imbalance issue, which is prevalent in medical image segmentation [30], incorporating both focal loss and dice loss. This approach not only penalizes the model for misclassifying minority classes but also ensures smooth gradient flow, enhancing overall performance [21, 31].

Evaluation of the model’s performance was conducted using a suite of metrics including Mean Intersection over Union (Mean IoU), Dice coefficient, Average Surface Distance (ASD), Normalized Surface Distance (NSD), precision, recall, and F1 score for each class. These metrics provided a comprehensive analysis of the model’s segmentation capabilities, ensuring that each class was accurately identified and delineated [32–34].

The results demonstrated significant improvements across all metrics compared to existing methods applied to the same dataset. This highlights the efficacy of the data preprocessing techniques, the robustness of the modified U-Net architecture, and the effectiveness of the custom loss function in handling class imbalance.

This study introduces a novel approach to lumbar spine segmentation, effectively addressing key challenges in data preprocessing and model training. Utilizing a carefully curated dataset and a modified deep learning model, the approach offers a robust solution that advances the state-of-the-art in medical image segmentation. The contributions presented have significant implications for clinical applications, with the potential to enhance diagnostic accuracy and treatment planning for patients experiencing low back pain.

The contributions can be summarized as follows:

- Implemented a robust preprocessing pipeline to extract 2D PNG images from 3D MRI mha files, ensuring accurate representation of four classes: background, vertebrae, spinal canal, and IVDs.
- Developed an advanced data filtering method in data preprocessing and a custom combined loss function (focal loss + dice loss) to effectively handle class imbalance, improving model performance on minority classes.
- Modified U-Net model incorporating an upsample block with leaky ReLU (alpha=0.1) and Glorot uniform initializer to prevent the dying ReLU problem and stabilize training.
- Applied strict data filtration criteria, removing images with fewer than four classes and those with a class imbalance ratio greater than 55%, ensuring high-quality training data.
- Used multiple metrics (Mean IoU, Dice coefficient, ASD, NSD, precision, recall, F1 score) to thoroughly assess model performance for each class, demonstrating significant improvements over existing methods.
- Achieved superior results across all evaluation metrics compared to existing studies using the same dataset, highlighting the effectiveness of the preprocessing techniques and the modified U-Net model.

2 Related Works

The segmentation process of the Lumbar Spine involves delineating anatomical structures such as vertebrae, spinal canal, and IVDs in MRI scans. Traditional methods for this task, while pioneering, often struggle with the intricate anatomy and variability of the lumbar spine.

Various deep learning methods have been employed to segment medical images, demonstrating remarkable performance. Among these, the U-Net [35] model has been particularly influential, setting a high standard for segmentation tasks. Building on the success of the original U-Net [35], several U-Net-based models, such as Attention U-Net [36], U-Net++ [37], and others, have been developed. These advanced architectures have consistently shown superior performance in the semantic segmentation of medical images, addressing specific challenges and further enhancing segmentation accuracy and efficiency [35–37]. Several studies have applied deep learning to lumbar spine segmentation with varying degrees of success.

Class weights of mask images for the lumbar spine are typically imbalanced, with the background class having a significantly higher weight compared to other classes. Additionally, the class weight of IVDs is much lower than that of vertebrae and the spinal canal. Addressing these imbalances is crucial during data preprocessing and in designing an appropriate loss function. [38] focused on the nn-UNET model and evaluated it using various metrics such as Dice score, ASD, and SD, achieving class-wise Dice coefficients in the range of 0.84 to 0.93. [39] processed the images to enhance the contrast and sharpness of different tissues and labeled the mask images into three categories, including five vertebral bodies. Their study introduced an improved Attention U-Net, featuring two residual modules replacing the original convolutional blocks. An attention module based on multilevel feature map fusion was employed, and a hybrid

loss function was utilized to address class imbalance. The improved Attention U-Net [39] achieved results with an accuracy of 95.50%, a recall of 0.9453, and a Dice similarity coefficient of 0.9501. To analyze lumbar spine segmentation more accurately, it is crucial to implement a variety of metrics to assess class-wise segmentation accuracy. Key metrics include Intersection over Union (IoU), Dice coefficient, ASD, NSD, precision, recall, F1-Score and others. In their study, [40] utilized a comprehensive set of these metrics to evaluate segmentation accuracy, achieving impressive results with an accuracy of 99%, a recall of 0.90, a precision of 0.89, and a Dice similarity coefficient of 0.924. But, In this study [41], Based on the chain structure of the spine, an interactive dual-output vertebrae instance segmentation network was designed to segment the specific vertebrae in CT images, a class-wise evaluation like [38] was implemented to provide a detailed analysis of segmentation performance. For the intervertebral foramen (IVF) region, the results achieved a Dice coefficient of 0.961 and an ASD of 0.29. For the vertebrae, the results were even more impressive, with a Dice coefficient of 0.968 and an ASD of 0.25. In this study [42], an extension of an iterative vertebra segmentation method was utilized, which relies on a 3D fully convolutional neural network to segment the vertebrae one by one. This approach achieved a mean Dice score of $93\% \pm 2\%$ for vertebra segmentation and $86\% \pm 7\%$ for IVDs segmentation. These results demonstrate the efficacy of their method in accurately segmenting vertebrae and IVDs, highlighting its potential for detailed anatomical analysis.

All the mentioned studies have demonstrated commendable performance in segmenting the lumbar spine. However, none have adequately addressed the issue of class imbalance during the data preprocessing stage. Notably, only [39] attempted to mitigate class imbalance by incorporating a hybrid loss function. For a comprehensive comparative analysis of lumbar spine segmentation, employing a diverse set of metrics is imperative. While [40] utilized multiple metrics to evaluate segmentation accuracy, they did not conduct a detailed class-wise evaluation. The proposed model addresses the class imbalance issue during the data preprocessing stage and model training by incorporating a combined loss function. This approach ensures superior performance compared to existing studies, providing a more robust and accurate segmentation of lumbar spine structures.

3 Experiments

3.1 Dataset

The dataset utilized in this study was obtained from the publicly available SPIDER dataset, titled "SPIDER - Lumbar Spine Segmentation in MR Images: A Dataset and a Public Benchmark" [23]. This dataset represents a comprehensive multi-center collection of lumbar spine MRI scans, complete with reference segmentations of vertebrae, IVDs, and the spinal canal.

The SPIDER dataset includes many 3D MRI scans, comprising 447 sagittal T1 and T2 MRI series derived from 218 studies of 218 patients, all of whom have a documented history of low back pain. The data was collected across four hospitals, ensuring a diverse representation of imaging practices and patient demographics. This diversity is crucial for developing robust and generalizable segmentation algorithms that aim

to promote advancements and collaborative efforts in spine segmentation, ultimately enhancing the diagnostic utility of lumbar spine MRI. The structured partitioning into training and validation sets and the detailed demographic and imaging metadata provide an excellent foundation for advancing spine segmentation.

3.2 Data Pre-Processing

3.2.1 Data Extraction

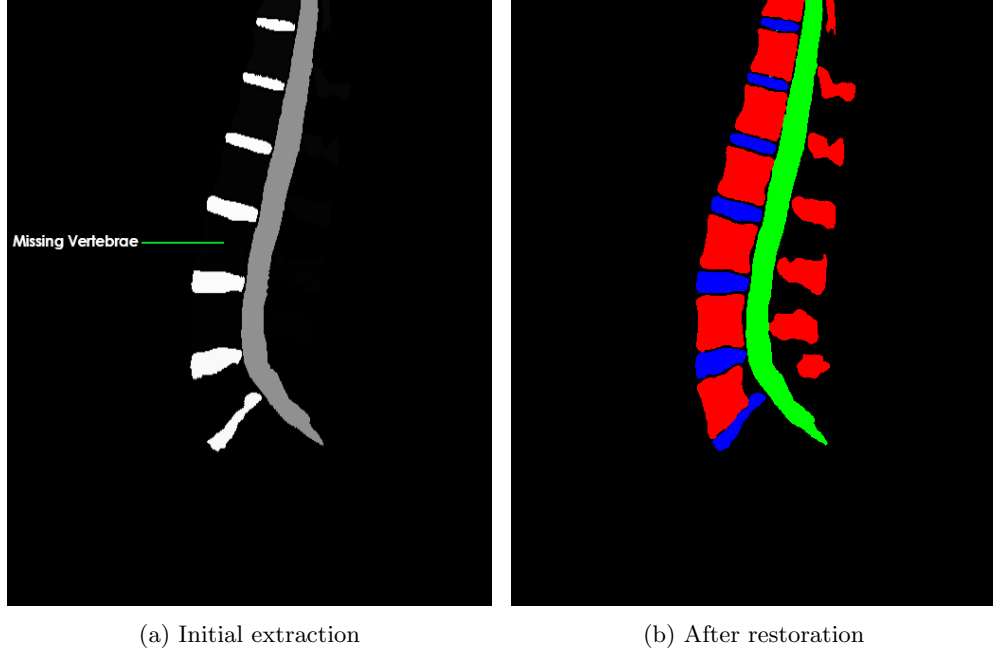


Fig. 1: Vertebrae Restoration by Custom Data Preprocessing Method

The default dataset comprised 3D MRI images in mha and corresponding segmentation masks. These images were categorized into three types: T1, T2, and T2.SPACE. Initial inspection using 3D Slicer [43] revealed the absence of the vertebrae class from all masks. This issue was resolved by adjusting the color lookup table in 3D Slicer, which made the vertebrae visible. For the deep learning model, the 3D images (Sagittal) were converted into 2D PNG format with a resolution of 512 by 640 pixels, including the MRI images and their corresponding masks.

When extracting 2D PNG images from the 3D mha files, four primary issues needed to be addressed:

1. Vertebrae class is missing
2. There should be in total 4 classes (Background, Vertebrae, Spinal canal, and IVDs), but 16 classes are present

3. Some of the slices of the images do not represent the sagittal slices
4. Some of the images are rotated or flipped

Issue no. 3 was resolved by manually identifying the images that represented the sagittal axis. The axial axis was extracted to obtain the sagittal slices for those specific 3D images in such cases. Issue no. 4 was addressed by rotating or flipping during the extraction process. In issues 1 and 2, the masks again missed the vertebrae. The extracted 2D masks contained 16 classes instead of the expected 4. This discrepancy necessitated further processing steps to correctly represent the vertebrae and reduce the number of courses to the desired four.

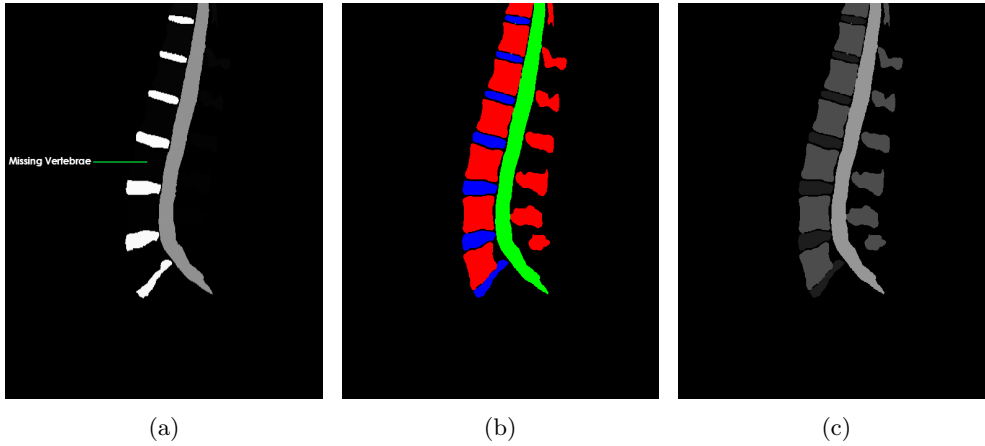


Fig. 2: Data Extraction Process

In Figure 1a, The Vertebra is not visible and the total classes of the image are 16. To overcome these issues, a multi-step preprocessing method was implemented. This method processes the initially extracted images, applying various color transformations and corrections to accurately represent the anatomical structures. It processes initially extracted images (missing vertebrae and 16 classes) from a specified input folder, applying multiple color transformations before saving them to an output folder. Initially, it replaces certain color ranges, such as very dark shades with red, mid-range shades with green, and light shades with blue. It then adjusts each pixel's color to match its neighbors if they share the same color, promoting uniformity. Next, it removes outlines by ensuring pixels that differ significantly from their neighbors adopt the most common surrounding color. The code also specifically addresses border pixels, ensuring their colors align with adjacent pixels. Additionally, it replaces isolated pixels that stand out with the most common neighboring color. If the image contains green or blue but lacks red, these colors are replaced with red to maintain consistency. The processed images are then saved to the output folder, ensuring each step enhances the visual coherence of the images.

In Figure 1b, the missing vertebra issue has been solved. Moreover, it now has exactly four classes, including the background class. Here, the red color represents the restored Vertebrae. Detailed processing steps are following.

Initially, specific color ranges are identified in the extracted images, categorizing them into very dark shades, mid-range shades, and light shades. Dark shades are replaced with red, mid-range shades are replaced with green, and light shades are replaced with blue. This color replacement technique highlights different anatomical structures, making them visually distinct and easier to analyze. Each pixel's color is then adjusted to match its neighboring pixels if they share the same color. This step enhances uniformity across the image and ensures smooth transitions between adjacent regions, reducing noise and enhancing the overall coherence of the image. Outliers, or pixels that differ significantly from their surrounding neighbors are identified next. These outliers are replaced with the most common surrounding color, effectively removing outlines and sharpening the boundaries within the segmentation masks. This process contributes to a cleaner and more precise representation of the structures. Special attention is given to border pixels, ensuring their colors align with those of adjacent pixels. This adjustment prevents discontinuities at the borders, maintaining the integrity of the segmentation masks and ensuring that boundaries are smooth and consistent. Additionally, isolated pixels that stand out from their surroundings are identified and replaced with the most common neighboring color. This step further enhances the uniformity and coherence of the segmentation masks, eliminating any residual noise or artifacts. In cases where an image contains green or blue but lacks red, these colors are replaced with red to maintain consistency across all images. This ensures that all images in the dataset have a uniform color scheme, which is critical for accurate segmentation. Finally, the processed images are saved to an output folder, preserving the enhancements made during preprocessing. Each step in this pipeline is carefully designed to improve the visual coherence of the images, facilitating more accurate and reliable segmentation, which is essential for downstream medical analysis.

Implementing this enhanced preprocessing method effectively addresses the initial issues encountered during data extraction, ensuring that the segmentation masks accurately represent the vertebrae, spinal canal, and IVDs in lumbar spine MRI images.

In Figure 2a, the 3D sagittal masks were missing vertebrae and included 16 confusing classes, making segmentation challenging for the deep learning model. To address this, a detailed preprocessing method APTA was applied, involving color transformations that restored the vertebrae class and refined the segmentation to 4 classes, as shown in Figure 2b. Finally, in Figure 2c, the RGB images were converted to grayscale, preserving class information while simplifying the data for training and effectively preparing the dataset for filtration.

3.2.2 Data Filtration

The data filtration stage focused on refining the extracted images to ensure the dataset's suitability and balance for training the deep learning model. The filtration process involved two main steps: eliminating unnecessary images and addressing class imbalance within the dataset.

Algorithm 1 Adaptive Pixel Transformation Algorithm (APTA)

Require: Input image I

Ensure: Output image O

for each pixel p_{ij} at position (i, j) in I **do**

1: Apply the following operations to each pixel:

1. **Replace pixel colors falling within specified ranges**

2. with red if not already present:

(a) Let $T : \mathbb{Z}^3 \rightarrow \mathbb{Z}^3$ be the color transformation function defined as $T(c) = (255, 0, 0)$ for c in the specified color range.

(b) Apply T to p_{ij} if p_{ij} falls within the specified

(c) color range.

3. **Replace p_{ij} with the color of the first matching adjacent pixel p_{xy} if p_{xy} matches p_{ij} :**

(a) Let $f : \mathbb{Z}^2 \rightarrow \mathbb{Z}^3$ represent the pixel colors of I .

(b) Define a pixel-wise transformation $T : \mathbb{Z}^3 \rightarrow \mathbb{Z}^3$ such that $T(p_{ij}) = p_{xy}$ if p_{xy} is the color of the

(c) first matching adjacent pixel to p_{ij} .

4. **Remove the outline or boundary of objects by replacing p_{ij} with the most common neighboring color:**

(a) Let $C(p_{ij})$ denote the multiset of colors of neighboring pixels to p_{ij} .

(b) Replace p_{ij} with the most common color in

(c) $C(p_{ij})$.

5. **Replace p_{ij} with the most common neighboring**

6. color if both adjacent pixels have different colors:

(a) Let $f : \mathbb{Z}^2 \rightarrow \mathbb{Z}^3$ represent the pixel colors of I .

(b) Define a conditional transformation $T : \mathbb{Z}^3 \rightarrow \mathbb{Z}^3$ such that $T(p_{ij}) = \text{MostCommonNeighbor}(p_{ij})$ if both adjacent pixels to p_{ij} have different colors.

7. **Replace isolated pixels (singletons) with the color of their most common neighboring pixel:**

(a) Define a function $\text{IsSingleton}(p_{ij})$ that determines

(b) if p_{ij} is a singleton based on its neighboring

(c) pixels.

(d) If p_{ij} is a singleton, replace it with the color of its most common neighboring pixel.

8. **Replace green or blue pixels with red if red color is**

9. not already present:

(a) Let $f : \mathbb{Z}^2 \rightarrow \mathbb{Z}^3$ represent the pixel colors of I .

(b) Define a conditional transformation $T : \mathbb{Z}^3 \rightarrow \mathbb{Z}^3$ such that $T(p_{ij}) = (255, 0, 0)$ if the specified condition is met.

Save the modified image O

2: **end for**

Step 1: Elimination of Redundant Images

Initially, in the dataset, masks images that contained less than 4 classes, precisely the black background, were deemed irrelevant for the training process. So, the corresponding original images with the mask images (less than four classes) were filtered out. These images do not contribute useful information for the segmentation task and were thus excluded from the dataset. This step ensures that all remaining images contain significant anatomical structures necessary for the training and validation of the deep learning model.




Fig. 3: Elimination of Redundant Images

In Figure 3a, contains only one class and it is the background class, rendering it redundant for the training process. Consequently, this mask image is excluded from the dataset. Similarly, images in Figure 3b & 3e contain fewer than four classes, which also makes them unsuitable for training and thus they are eliminated. On the other hand, 3c and 3d contain the appropriate number of classes and are therefore deemed suitable for training.

Finally, the Figure 4a and 4b are selected by filtering out Figure 3a, Figure 3b and Figure 3e.

Step 2: Addressing Class Imbalance

The dataset comprised three types of images: T1, T2, and T2.SPACE. To address the class imbalance, the class weights of the dataset were calculated. The class weight for each image type was determined as follows:

$$\text{Class Weight} = \frac{\text{Number of Pixels of a Specific Class}}{\text{Total Number of Pixels in the Image}}$$

Once the class weights were determined, the class imbalance ratio was calculated by dividing the highest class weight by the lowest class weight across the dataset.

The class imbalance ratio was recalculated after filtering by dividing the highest class weight by the lowest class weight:

$$\text{Class Imbalance Ratio} = \frac{\text{Highest Class Weight}}{\text{Lowest Class Weight}}$$

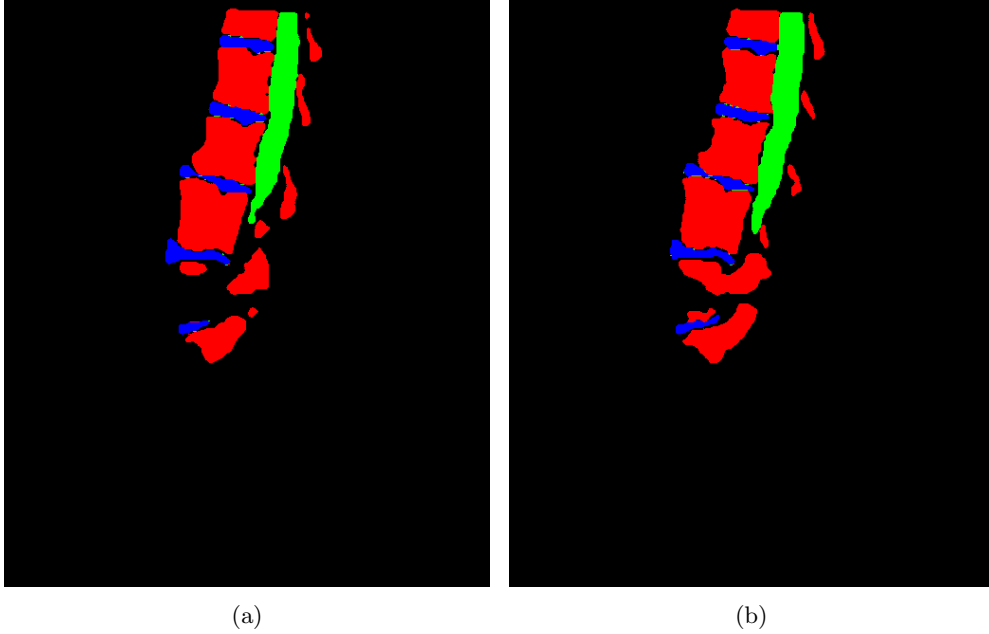


Fig. 4: After Elimination of Redundant Images

Initially, the class imbalance ratios for T1, T2 and T2.SPACE are respectively 57%, 56% and 40%.

Images were filtered based on their class weight ratios to mitigate this imbalance. Specifically, images with a class weight ratio above 55% were considered over-represented and were subsequently removed. This threshold was empirically determined to reduce the over-representation of particular image types without excessively reducing the dataset size.

In Figure 5, Figure 5c significantly contributes to class imbalance. Upon calculating the class imbalance ratio for this image, it was found to exceed 55%. To address this imbalance and improve the dataset's uniformity, Figure 5c should be excluded from the training set.

After applying this filtration step, the class distribution was adjusted to 55% for T1, 55% for T2, and 37% for T2.SPACE images.

In Figure 6, Figure 6a and Figure 6b are evaluated for their suitability in the training dataset. After thorough analysis and verification, these images meet the required criteria and are deemed suitable for inclusion in the training set. Their inclusion ensures a balanced and representative dataset, which is crucial for the practical training of the model. This post-filtration adjustment reduced the class imbalance to 3.51% for T1, 1.79% for T2, and 7.5% for T2.SPACE images. This reduction in class imbalance is critical for training a robust deep learning model, as it prevents the model from becoming biased towards more frequently occurring classes. By ensuring a more

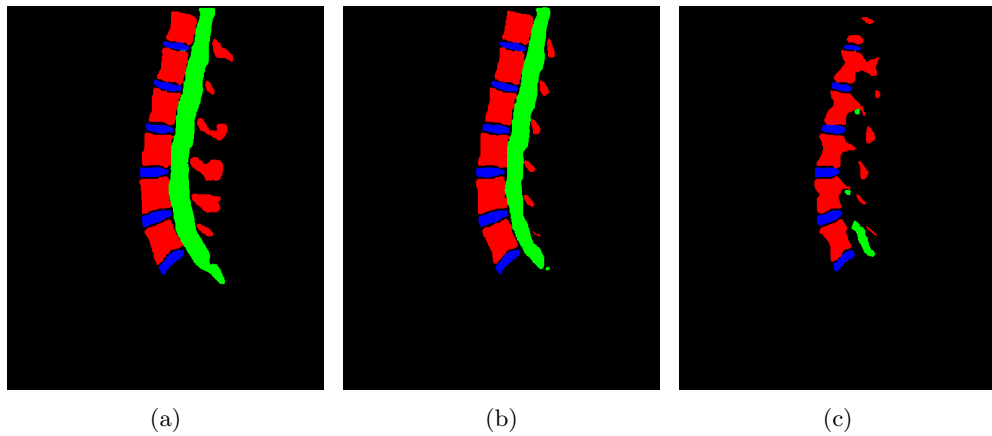


Fig. 5: Handling Class Imbalance

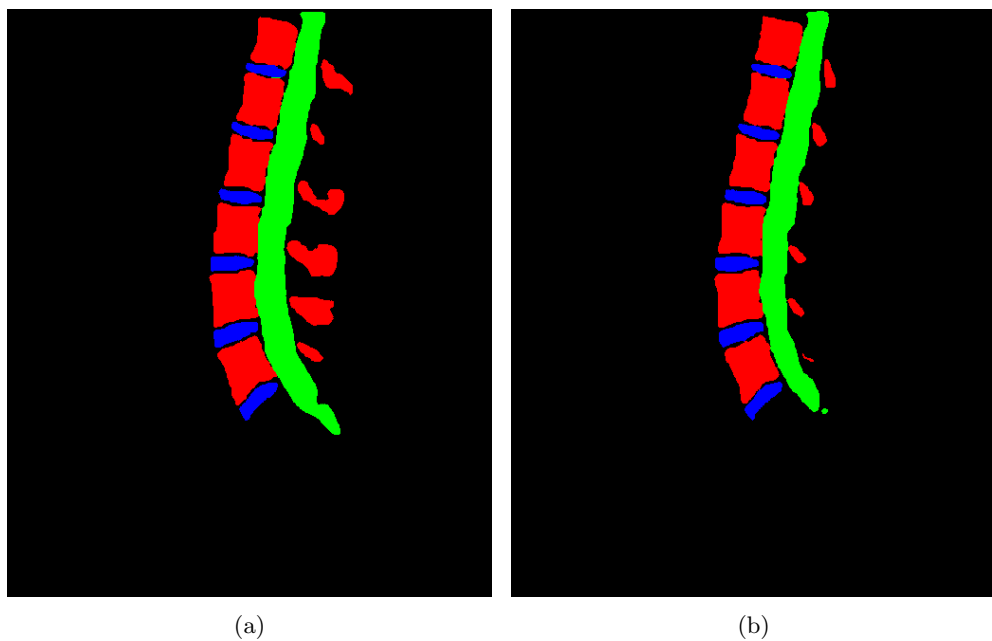


Fig. 6: After Handling Class Imbalance

uniform distribution of T1, T2, and T2_SPACE images, the model can learn features across all image types more effectively, leading to improved segmentation performance.

Summing up the data extraction and filtration processes, the dataset T1, T2, and T3 weighted images each has 1000 2D MRI images with corresponding masks referring in Table 1.

Dataset	No. of Images	Class Imbalance Ratio
<i>T1</i>	1000	55%
<i>T2</i>	1000	55%
<i>T2_SPACE</i>	1000	37%
<i>Total</i>	3000	—

Table 1: Dataset Summary with Class Imbalance Ratio

Overall, the data filtration process improved the quality and balance of the dataset, thereby enhancing the potential accuracy and generalizability of the deep learning model for lumbar spine segmentation.

3.3 Evaluation Metrics

Several key metrics were utilized to evaluate the performance of the modified U-Net model comprehensively on the task of lumbar spine segmentation. These metrics provide a holistic view of the model’s accuracy, precision, and ability to handle class imbalance. Each metric and its mathematical formulation are detailed below:

Mean Intersection over Union (Mean IoU)

The IoU measures the overlap between the predicted segmentation and the ground truth. For each class c , IoU is defined as:

$$\text{IoU}_c = \frac{|P_c \cap G_c|}{|P_c \cup G_c|}$$

where P_c and G_c are the predicted and ground truth binary masks for class c , respectively. Mean IoU is the average IoU across all classes:

$$\text{Mean IoU} = \frac{1}{C} \sum_{c=1}^C \text{IoU}_c$$

where C is the total number of classes.

Dice Coefficient

The Dice coefficient, also known as the Sørensen-Dice index, is another measure of overlap between the predicted and ground truth segmentations. For each class c , the Dice coefficient is defined as:

$$\text{Dice}_c = \frac{2|P_c \cap G_c|}{|P_c| + |G_c|}$$

The average Dice coefficient is calculated across all classes.

Average Surface Distance (ASD)

The ASD quantifies the average distance between the surfaces of the predicted segmentation and the ground truth. For a given class c , ASD is computed as:

$$\text{ASD}_c = \frac{1}{|S_p| + |S_g|} \left(\sum_{x \in S_p} \min_{y \in S_g} d(x, y) + \sum_{y \in S_g} \min_{x \in S_p} d(y, x) \right)$$

where S_p and S_g are the sets of surface points in the predicted and ground truth segmentations, respectively, and $d(x, y)$ is the Euclidean distance between points x and y .

Normalized Surface Distance (NSD)

NSD measures the surface distance normalized by a tolerance value τ :

$$\text{NSD}_c = \frac{1}{|S_g|} \sum_{y \in S_g} \mathbb{1} \left[\min_{x \in S_p} d(y, x) < \tau \right]$$

where $\mathbb{1}$ is the indicator function that counts the number of ground truth surface points within the tolerance τ from any predicted surface point.

Precision

Precision for each class c measures the proportion of true positive predictions among all predicted positives:

$$\text{Precision}_c = \frac{TP_c}{TP_c + FP_c}$$

where TP_c and FP_c are the true positive and false positive counts for class c , respectively.

Recall

Recall for each class c measures the proportion of true positive predictions among all actual positives:

$$\text{Recall}_c = \frac{TP_c}{TP_c + FN_c}$$

where FN_c is the false negative count for class c .

F1 Score

The F1 score is the harmonic mean of precision and recall for each class c :

$$F1_c = 2 \cdot \frac{\text{Precision}_c \cdot \text{Recall}_c}{\text{Precision}_c + \text{Recall}_c}$$

The overall F1 score is averaged across all classes.

These metrics collectively provide a detailed evaluation of the model’s performance, ensuring that each class is accurately segmented and that the model’s predictions align closely with the ground truth annotations. The modified U-Net model demonstrates significant improvements across all these metrics compared to existing methods, underscoring the approach’s effectiveness in lumbar spine segmentation.

4 Methodology




Fig. 7: Detailed Architecture of Modified U-Net Model with Custom Upsample Block for Lumbar Spine Segmentation

The model architecture in Figure 7 is based on a modified UNET framework tailored for semantic segmentation tasks in Lumbar Spine segmentation. The model has been meticulously crafted through extensive empirical refinement to optimize its performance for medical image segmentation tasks. It consists of a contraction path followed by an expansive path, and central to its design is an innovative upsample block featuring Leaky ReLU activation.

The model’s contractive path begins with convolutional and pooling layers to extract hierarchical features from the input image, with batch normalization and dropout layers included to enhance generalization and prevent overfitting [44, 45]. In the expansive path, upsample blocks with `Conv2DTranspose` layers and Leaky ReLU activation ($\alpha = 0.1$) maintain a small, non-zero gradient even when units are inactive, improving gradient flow and preventing neuron death during training[46]. An additional 512-channel layer boosts the model’s ability to capture intricate details crucial for accurate segmentation. To address the dying ReLU issue, the Glorot uniform initializer is used, promoting stable gradient propagation [47]. Addressing the challenge of class imbalance crucial in medical imaging, a bespoke combined loss function combines Focal Loss and Dice Loss, prioritizing hard-to-classify examples while balancing segmentation precision[48, 49]. A combined loss function is adopted, integrating Focal Loss and Dice Loss.

$$L_{\text{focal}}(y_{\text{true}}, y_{\text{pred}}) = - \sum_i \alpha_i (1 - y_{\text{pred},i})^\gamma y_{\text{true},i} \log(y_{\text{pred},i})$$

Focal loss is designed to address class imbalance by focusing more on hard-to-classify examples, reducing the contribution of easily classified examples.

$$L_{\text{dice}}(y_{\text{true}}, y_{\text{pred}}) = 1 - \frac{2 \sum_i y_{\text{true},i} y_{\text{pred},i} + \epsilon}{\sum_i y_{\text{true},i} + \sum_i y_{\text{pred},i} + \epsilon}$$

Dice loss is a measure of overlap between the predicted segmentation and the ground truth, directly optimizing for the Dice coefficient.

The Focal Loss component, with a gamma ($\gamma = 4.0$) value of 4.0, emphasizes accurate classification of challenging examples, crucial in medical images where anomalies may be subtle yet significant. Prioritizing performance criteria relevant to medical image analysis, the combined loss function’s alpha ($\alpha = .6$) parameter of 0.6 strikes a balance between segmentation precision and minimizing class imbalance.

$$L_{\text{combined}} = \alpha L_{\text{focal}} + (1 - \alpha) L_{\text{dice}}$$

By combining focal and dice loss, the model benefits from both approaches. Alpha (α) and Gamma (γ) Values of the combined loss function were set to 0.6 and 0.4, respectively. The gamma value controls the focus on hard-to-classify regions, and the alpha value controls the balance between focal loss and dice loss functions. This balance ensures that both class imbalance and segmentation accuracy are improved simultaneously[48, 49].

The modified model can capture more complex features and improve training dynamics. Incorporating focal and dice loss, the combined loss function effectively addresses class imbalance and boundary accuracy. At the same time, early stopping and model checkpointing based on validation Mean IoU ensures robust and generalized performance. These improvements collectively lead to better segmentation results, making the modified model superior for lumbar spine segmentation in MRI scans.

5 Result & Analysis

This section presents the results of the proposed model for the semantic segmentation of lumbar spine MRI scans. The model’s performance is evaluated using accuracy, loss, Dice coefficient, Mean IoU, and class-specific metrics.

Figure 8 visually compares the predicted segmentation masks and the labeled ground truth masks, referring to the original MRI images. This comparison highlights the model’s ability to segment the IVDs, Vertebrae, and Spinal Canal accurately.

In Figure 9, The model effectively distinguished each segment, showcasing its precision in accurately delineating the IVDs, Vertebrae, and Spinal Canal. This contrast underscores the model’s proficiency in precisely segmenting these anatomical structures.

In Figure 10, the red circles on the testing label image highlight areas where the segmentation failed to correctly separate the vertebrae from the IVDs, resulting in noticeable misclassifications. However, in the prediction image marked by white circles, the proposed model demonstrates its effectiveness by successfully resolving these




Fig. 8: Comparison between Labelled Mask and Predicted Mask




Fig. 9: Detection of IVDs, Vertebrae and Spinal Canal in Predicted Mask




Fig. 10: Visualization of Predicted Mask Overlaid on Original Image

issues. The model accurately separates the vertebrae from the IVDs and corrects previous misclassifications, showcasing its superior performance. The overall accuracy and reliability of the model are clearly evident in these results.




Fig. 11: Visualization of Predicted Mask Overlaid on Original Image

Figure 11 demonstrates that the predicted mask aligns seamlessly with the expected positions of each class. The overlap highlights the model's ability to accurately outline the IVDs, Vertebrae, and Spinal Canal. This observation highlights the model's skill in segmenting these anatomical structures with accuracy.

The training progress of the model is illustrated in Figures 12, 13, and 14, which show the Accuracy, Dice coefficient and Mean IoU over 100 epochs. The batch size was set to 8. The training and validation curves almost overlap after 100 epochs, indicating that the model does not suffer from overfitting. This is crucial for ensuring that the model generalizes well to unseen data, which is essential for practical deployment in clinical settings. The absence of overfitting ensures that the model maintains its high performance across different datasets, making it a reliable tool for clinical applications.

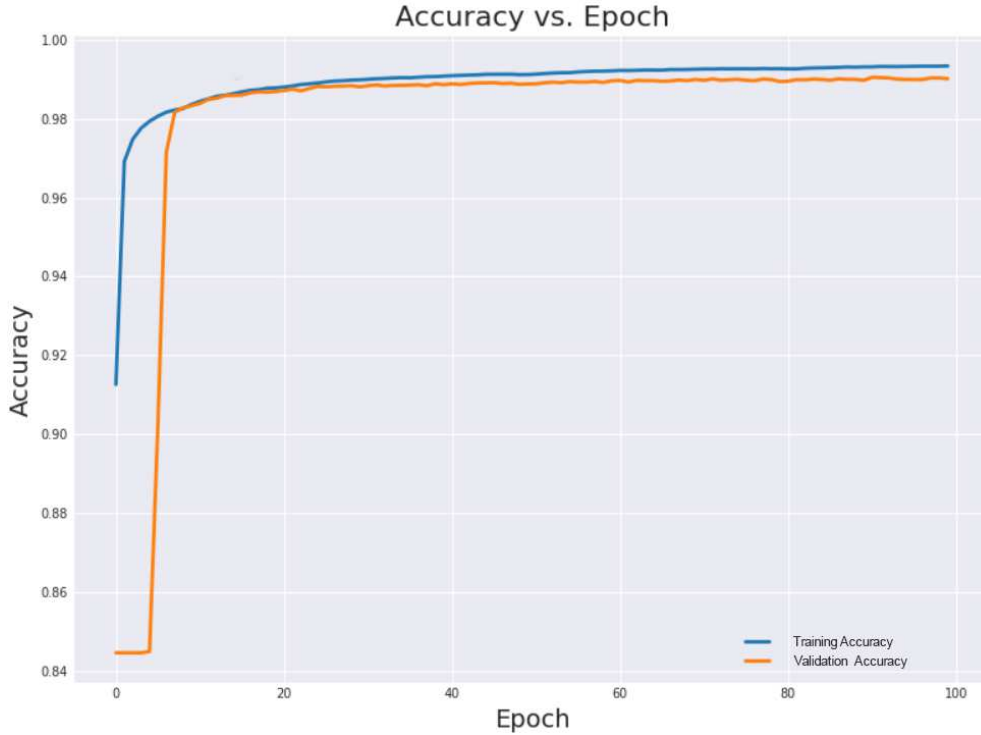


Fig. 12: Accuracy vs Epochs

A comprehensive comparison was conducted between the proposed model, and another model, nn-UNET[38], using the same dataset. Notably, the evaluation of the nn-UNET[38] model lacked several crucial metrics essential for medical image segmentation tasks, such as IoU, NSD, Precision, Recall, and F1-score. In contrast, the proposed model was assessed using a comprehensive set of metrics necessary for such tasks.

As illustrated in Table 3, the scores obtained by the proposed model significantly outperform those of nn-UNET[38] across all metrics. This clear disparity underscores the superior performance of the proposed model in the segmentation task. Specifically, the mean Dice coefficient, a widely used metric for assessing segmentation accuracy, was notably higher for the proposed model compared to nn-UNET[38] in all classes, including IVDs, Vertebrae, and the Spinal Canal. This indicates that the model achieved better overlap and boundary delineation between the predicted and ground truth masks, leading to more accurate segmentation.

Furthermore, the proposed model exhibited lower ASD values for all classes, implying superior boundary detection and closer alignment with the actual anatomical structures. The high NSD values obtained by the proposed model also suggest better overall alignment and spatial correspondence between the segmented regions and the ground truth masks.

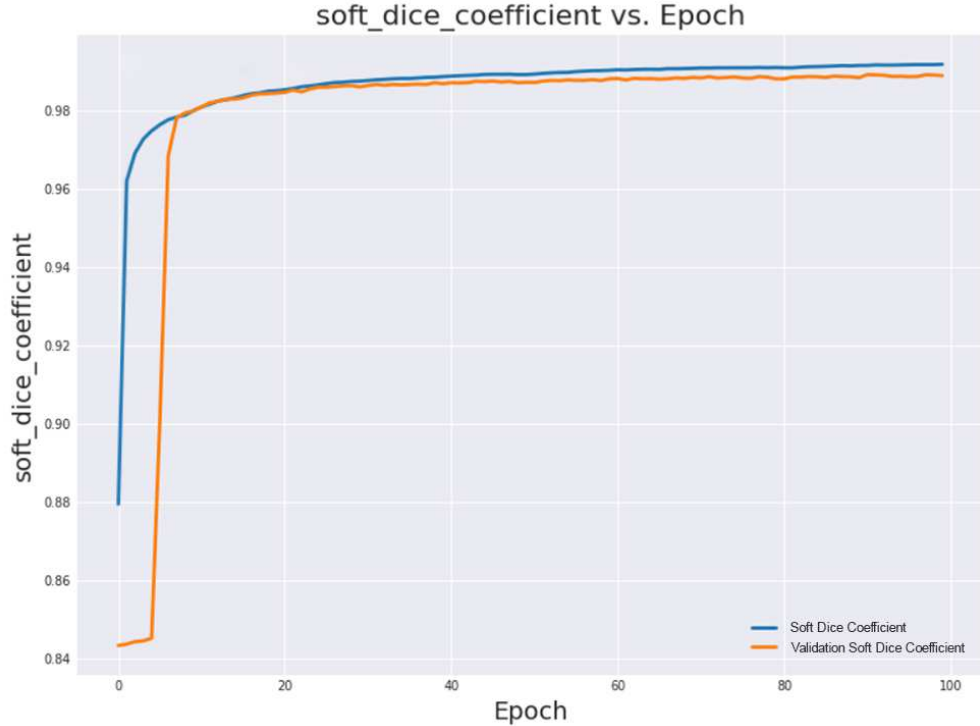


Fig. 13: Dice Coefficient vs Epochs

While the nn-UNET[38] model achieved comparable detection rates for some classes, the lack of comprehensive evaluation metrics limits its utility in assessing segmentation quality. In contrast, the detailed evaluation provided by the proposed model, including IoU, NSD, Precision, Recall, and F1-score, offers a more comprehensive understanding of its segmentation performance.

Overall, the results highlight the effectiveness of the proposed model in accurately segmenting lumbar spine structures from MRI images. Its superior performance, as demonstrated by a comprehensive range of metrics, makes it a promising tool for clinical applications, including disease diagnosis, treatment planning, and medical research.

6 Conclusion

This study presents a sophisticated approach to lumbar spine segmentation using deep learning techniques, with a focus on addressing key challenges such as class imbalance and data preprocessing. By meticulously preprocessing MRI scans of patients with low back pain and rectifying class inconsistencies, the fidelity of the training data is ensured, laying the foundation for an accurate representation of critical classes: vertebrae, spinal canal, and IVDs. The highly modified U-Net model (proposed model), incorporating innovative architectural enhancements and a custom combined loss

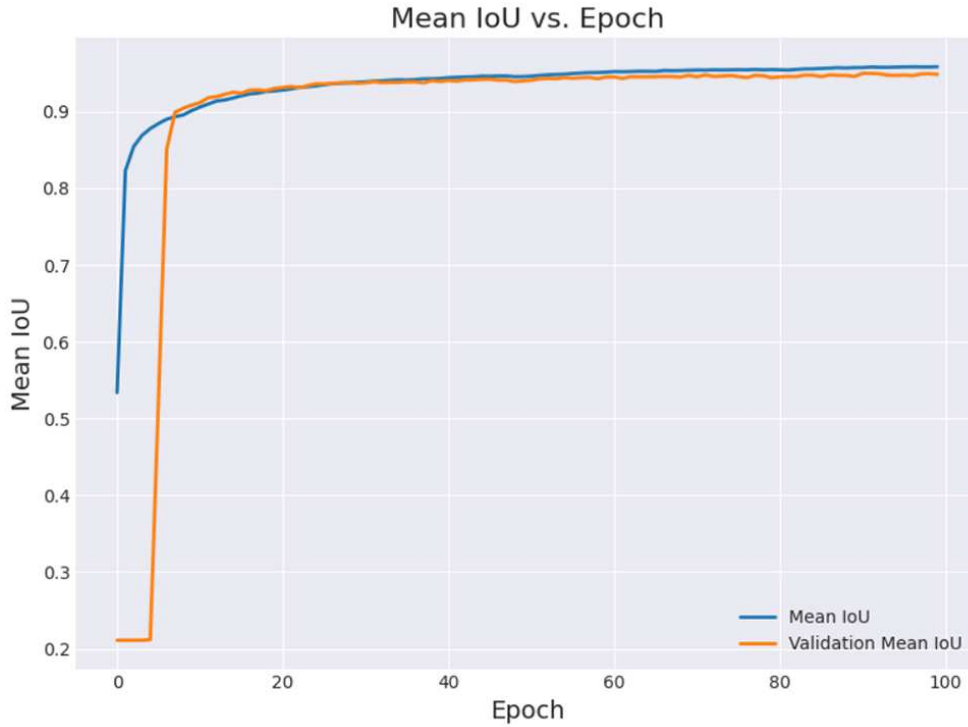


Fig. 14: Mean IoU vs Epochs

function, effectively tackles class imbalance and improves segmentation accuracy. Evaluation using a comprehensive suite of metrics demonstrates the superior performance of this approach, outperforming existing methods and advancing current techniques in lumbar spine segmentation. The comparison with another model nn-UNET, highlights the superiority of the proposed model, with significantly higher scores across all evaluation metrics. Specifically, the proposed method achieves better overlap, boundary delineation, and alignment with ground truth masks, making it a promising tool for clinical applications such as disease diagnosis, treatment planning, and medical research.

Overall, the contributions of this study hold significant implications for enhancing diagnostic accuracy and treatment planning for patients with low back pain, showcasing the potential of deep learning techniques in advancing medical image segmentation.

Acknowledgment

Istiaq Ahmed and Md. Zahirul Islam Nahid would like to thank Dr. Mohammad Monirujjaman Khan and Dr. Mohammad Monir Uddin for their guidance during their CSE499 undergraduate project work at North South University. Some of the results presented in this paper were derived from that project.

References

- [1] Roudsari, B., Jarvik, J.G.: Lumbar spine mri for low back pain: indications and yield. *American Journal of Roentgenology* **195**(3), 550–559 (2010)
- [2] Sheehan, N.: Magnetic resonance imaging for low back pain: indications and limitations. *Postgraduate medical journal* **86**(1016), 374–378 (2010)
- [3] Katz, J.N.: Lumbar disc disorders and low-back pain: socioeconomic factors and consequences. *JbJs* **88**(suppl_2), 21–24 (2006)
- [4] Richardson, J.C., Bowtell, R.W., Mäder, K., Melia, C.D.: Pharmaceutical applications of magnetic resonance imaging (mri). *Advanced drug delivery reviews* **57**(8), 1191–1209 (2005)
- [5] Van Goethem, J., Parizel, P., Jinkins, J.: Mri of the postoperative lumbar spine. *Neuroradiology* **44**, 723–739 (2002)
- [6] Kirkham, A.P., Emberton, M., Allen, C.: How good is mri at detecting and characterising cancer within the prostate? *European urology* **50**(6), 1163–1175 (2006)
- [7] Patil, D.D., Deore, S.G.: Medical image segmentation: a review. *International Journal of Computer Science and Mobile Computing* **2**(1), 22–27 (2013)
- [8] Yang, Z., Wang, Q., Zeng, J., Qin, P., Chai, R., Sun, D.: Rau-net: U-net network based on residual multi-scale fusion and attention skip layer for overall spine segmentation. *Machine Vision and Applications* **34**(1), 10 (2023)
- [9] Yanagisawa, Y., Shido, K., Kojima, K., Yamasaki, K.: Convolutional neural network-based skin image segmentation model to improve classification of skin diseases in conventional and non-standardized picture images. *Journal of Dermatological Science* **109**(1), 30–36 (2023)
- [10] Narayan, V., Faiz, M., Mall, P.K., Srivastava, S.: A comprehensive review of various approach for medical image segmentation and disease prediction. *Wireless Personal Communications* **132**(3), 1819–1848 (2023)
- [11] Warfield, S.K., Jolesz, F.A., Kikinis, R.: Real-time image segmentation for image-guided surgery. In: *SC'98: Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, pp. 42–42 (1998). IEEE
- [12] Li, J., Erdt, M., Janoos, F., Chang, T.-c., Egger, J.: Medical image segmentation in oral-maxillofacial surgery. *Computer-Aided Oral and Maxillofacial Surgery*, 1–27 (2021)
- [13] Aggarwal, P., Vig, R., Bhadoria, S., Dethe, C.: Role of segmentation in medical imaging: A comparative study. *International Journal of Computer Applications*

29(1), 54–61 (2011)

- [14] Zhou, S.K., Greenspan, H., Davatzikos, C., Duncan, J.S., Van Ginneken, B., Madabhushi, A., Prince, J.L., Rueckert, D., Summers, R.M.: A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. *Proceedings of the IEEE* **109**(5), 820–838 (2021)
- [15] Cheng, P.M., Montagnon, E., Yamashita, R., Pan, I., Cadrin-Chênevert, A., Perdigón Romero, F., Chartrand, G., Kadoury, S., Tang, A.: Deep learning: an update for radiologists. *Radiographics* **41**(5), 1427–1445 (2021)
- [16] Wataya, T., Yanagawa, M., Tsubamoto, M., Sato, T., Nishigaki, D., Kita, K., Yamagata, K., Suzuki, Y., Hata, A., Kido, S., *et al.*: Radiologists with and without deep learning-based computer-aided diagnosis: comparison of performance and interobserver agreement for characterizing and diagnosing pulmonary nodules/masses. *European Radiology* **33**(1), 348–359 (2023)
- [17] Neves, J., Hsieh, C., Nobre, I.B., Sousa, S.C., Ouyang, C., Maciel, A., Duchowski, A., Jorge, J., Moreira, C.: Shedding light on ai in radiology: A systematic review and taxonomy of eye gaze-driven interpretability in deep learning. *European Journal of Radiology*, 111341 (2024)
- [18] Liao, Y., Liu, H., Spasić, I.: Deep learning approaches to automatic radiology report generation: A systematic review. *Informatics in Medicine Unlocked* **39**, 101273 (2023)
- [19] Najjar, R.: Redefining radiology: a review of artificial intelligence integration in medical imaging. *Diagnostics* **13**(17), 2760 (2023)
- [20] Patel, J.B.: Improving segmentation pipelines for medical imaging using deep learning. PhD thesis, Massachusetts Institute of Technology (2023)
- [21] Thölke, P., Mantilla-Ramos, Y.-J., Abdelhedi, H., Maschke, C., Dehgan, A., Harel, Y., Kemtur, A., Berrada, L.M., Sahraoui, M., Young, T., *et al.*: Class imbalance should not throw you off balance: Choosing the right classifiers and performance metrics for brain decoding with imbalanced data. *NeuroImage* **277**, 120253 (2023)
- [22] Petrakis, G., Partsinevelos, P.: Lunar ground segmentation using a modified u-net neural network. *Machine Vision and Applications* **35**(3), 50 (2024)
- [23] Graaf, J., *et al.*: Spider-lumbar spine segmentation in mr images: a dataset and a public benchmark. *Zenodo* <https://doi.org/10.5281/zenodo.10159290> (2023)
- [24] Siddique, N., Paheding, S., Elkin, C.P., Devabhaktuni, V.: U-net and its variants for medical image segmentation: A review of theory and applications. *IEEE access* **9**, 82031–82057 (2021)

- [25] Mastromichalakis, S.: Alrelu: A different approach on leaky relu activation function to improve neural networks performance. arXiv preprint arXiv:2012.07564 (2020)
- [26] Dubey, A.K., Jain, V.: Comparative study of convolution neural network’s relu and leaky-relu activation functions. In: Applications of Computing, Automation and Wireless Systems in Electrical Engineering: Proceedings of MARC 2018, pp. 873–880 (2019). Springer
- [27] Douglas, S.C., Yu, J.: Why relu units sometimes die: Analysis of single-unit error backpropagation in neural networks. In: 2018 52nd Asilomar Conference on Signals, Systems, and Computers, pp. 864–868 (2018). IEEE
- [28] Chumachenko, K., Iosifidis, A., Gabbouj, M.: Feedforward neural networks initialization based on discriminant learning. *Neural Networks* **146**, 220–229 (2022)
- [29] Aguirre, D., Fuentes, O.: Improving weight initialization of relu and output layers. In: Artificial Neural Networks and Machine Learning–ICANN 2019: Deep Learning: 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019, Proceedings, Part II 28, pp. 170–184 (2019). Springer
- [30] Cui, L., Li, D., Yang, X., Liu, C., Yan, X.: A unified approach addressing class imbalance in magnetic resonance image for deep learning models. *IEEE Access* (2024)
- [31] Yeung, M., Sala, E., Schönlieb, C.-B., Rundo, L.: 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)
- [32] Wang, Z., Berman, M., Rannen-Triki, A., Torr, P., Tuia, D., Tuytelaars, T., Gool, L.V., Yu, J., Blaschko, M.: Revisiting evaluation metrics for semantic segmentation: Optimization and evaluation of fine-grained intersection over union. *Advances in Neural Information Processing Systems* **36** (2024)
- [33] Liu, B., Dolz, J., Galdran, A., Kobbi, R., Ayed, I.B.: Do we really need dice? the hidden region-size biases of segmentation losses. *Medical Image Analysis* **91**, 103015 (2024)
- [34] Opitz, J.: A closer look at classification evaluation metrics and a critical reflection of common evaluation practice. *Transactions of the Association for Computational Linguistics* **12**, 820–836 (2024)
- [35] Ronneberger, O., Fischer, P., Brox, T.: 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, pp. 234–241 (2015). Springer

- [36] Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., et al.: Attention u-net: Learning where to look for the pancreas. arXiv preprint arXiv:1804.03999 (2018)
- [37] Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4, pp. 3–11 (2018). Springer
- [38] Graaf, J.W., Hooff, M.L., Buckens, C.F., Rutten, M., Susante, J.L., Kroeze, R.J., Kleuver, M., Ginneken, B., Lessmann, N.: Lumbar spine segmentation in mr images: a dataset and a public benchmark. *Scientific Data* **11**(1), 264 (2024)
- [39] Wang, S., Jiang, Z., Yang, H., Li, X., Yang, Z.: Automatic segmentation of lumbar spine mri images based on improved attention u-net. *Computational Intelligence and Neuroscience* **2022**(1), 4259471 (2022)
- [40] Silvester M, L., Mathusoothana S. Kumar, R.: Efficient segmentation of lumbar intervertebral disc from mr images. *IET Image Processing* **14**(13), 3076–3083 (2020)
- [41] Peng, W., Li, L., Liang, L., Ding, H., Zang, L., Yuan, S., Wang, G.: A convenient and stable vertebrae instance segmentation method for transforaminal endoscopic surgery planning. *International Journal of Computer Assisted Radiology and Surgery* **16**(8), 1263–1276 (2021)
- [42] Graaf, J.W., Hooff, M.L., Buckens, C.F., Lessmann, N.: Segmentation of vertebrae and intervertebral discs in lumbar spine mr images with iterative instance segmentation. In: *Medical Imaging 2022: Image Processing*, vol. 12032, pp. 909–913 (2022). SPIE
- [43] Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.-C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., et al.: 3d slicer as an image computing platform for the quantitative imaging network. *Magnetic resonance imaging* **30**(9), 1323–1341 (2012)
- [44] Santurkar, S., Tsipras, D., Ilyas, A., Madry, A.: How does batch normalization help optimization? *Advances in neural information processing systems* **31** (2018)
- [45] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1), 1929–1958 (2014)

- [46] Hu, Z., Zhang, J., Ge, Y.: Handling vanishing gradient problem using artificial derivative. *IEEE Access* **9**, 22371–22377 (2021)
- [47] Aguirre, D., Fuentes, O.: Improving weight initialization of relu and output layers. In: *Artificial Neural Networks and Machine Learning–ICANN 2019: Deep Learning: 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019, Proceedings, Part II* 28, pp. 170–184 (2019). Springer
- [48] Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2980–2988 (2017)
- [49] Zhao, R., Qian, B., Zhang, X., Li, Y., Wei, R., Liu, Y., Pan, Y.: Rethinking dice loss for medical image segmentation. In: *2020 IEEE International Conference on Data Mining (ICDM)*, pp. 851–860 (2020). IEEE

Dataset	Class	IoU \uparrow	Dice Coefficient \uparrow	ASD \downarrow	NSD \uparrow	Precision \uparrow	Recall \uparrow	F1 Score \uparrow
<i>T2_SPACE</i>	Background	0.9906	0.9950	0.0543	0.9951	0.9951	0.9955	0.9953
	IVDs	0.9476	0.9688	0.0288	0.9994	0.9684	0.9778	0.9731
	Vertebrae	0.9461	0.9712	0.0464	0.9944	0.9745	0.9701	0.9723
	Spinal Canal	0.9501	0.9671	0.0361	0.9963	0.9761	0.9727	0.9744
<i>T1</i>	Background	0.9891	0.9893	0.0578	0.9901	0.9899	0.9921	0.9917
	IVDs	0.9378	0.9592	0.0389	0.9954	0.9556	0.9696	0.9663
	Vertebrae	0.9367	0.9634	0.0489	0.9912	0.9711	0.9697	0.9678
	Spinal Canal	0.9445	0.9628	0.0372	0.9893	0.9711	0.9692	0.9725
<i>T2</i>	Background	0.9867	0.9877	0.0533	0.9921	0.9856	0.9933	0.9867
	IVDs	0.9317	0.9527	0.0401	0.9934	0.9524	0.9658	0.9650
	Vertebrae	0.9331	0.9647	0.0488	0.9919	0.9703	0.9671	0.9666
	Spinal Canal	0.9439	0.9625	0.0371	0.9856	0.9689	0.9697	0.9709

Table 2: Comprehensive Performance Metrics across each Class and Dataset. \uparrow denotes higher is better and \downarrow denotes lower is better

Class	Metric	Proposed	nn-UNET[38]
IVDs	Mean Dice Coefficient	0.9688	0.86
	IoU	0.9476	—
	ASD	0.0288	0.58
	NSD	0.9994	—
Vertebrae	Mean Dice Coefficient	0.9712	0.92
	ASD	0.0464	0.51
	Detection	100%	99.7%
Spinal Canal	Mean Dice Coefficient	0.9671	0.92
	ASD	0.0361	0.46
	Detection	100%	100%

Table 3: Performance Comparison of Proposed Model vs. nn-UNET Across Multiple Segmentation Metrics for Lumbar Spine MRI