

Semantic Segmentation of Mice Wounds

Bruno Uhlmann Marcato
marcato.2001@alunos.utfpr.edu.br
Federal University of Technology –
Paraná (UTFPR)
Apucarana, Paraná, Brazil

Soraia Mendes Pierotti
soraia.pierotti@uel.br
State University of Londrina (UEL)
Londrina, Paraná, Brasil

Patricia Durgante Ritter
patriciadritter@gmail.com
State University of Londrina (UEL)
Londrina, Paraná, Brasil

Camila Rodrigues Ferraz
camila_ferraz96@hotmail.com
State University of Londrina (UEL)
Londrina, Paraná, Brasil

Waldiceu Aparecido Verri Jr
waverri@uel.br
State University of Londrina (UEL)
Londrina, Paraná, Brasil

Rubia Casagrande
rubiacasa@uel.br
State University of Londrina (UEL)
Londrina, Paraná, Brasil

José Luis Seixas Junior
jlseixasjr@inf.elte.hu
ELTE – Eötvös Loránd University,
Faculty of Informatics
Budapest, Hungary

Rafael Gomes Mantovani
rafaelmantovani@utfpr.edu.br
Federal University of Technology –
Paraná (UTFPR)
Apucarana, Paraná, Brazil

ABSTRACT

Semantic segmentation has been successfully explored in biological studies to handle various applications, such as identifying wounds. This study explores two image segmentation approaches to identify mice wounds, specifically the U-Net and Random Forest algorithms. The latter was combined with features extracted from the first two layers of VGG16, which was used as a feature extractor. Experiments were performed with a real dataset developed by the Pain, Neuropathy, and Inflammation Laboratory at the State University of Londrina with the approval of the University Ethics Committee on Animal Research and Welfare. The experimental results were promising, showing that both alternatives can provide accurate predictions for most images regarding FScore and IoU evaluation measures. Statistical tests were also applied, showing that U-Net obtained statistically better results with an average FScore of 0.72 and IoU of 0.58.

KEYWORDS

Semantic Segmentation, wound identification, U-Net, Random Forest.

1 INTRODUCTION

Image Segmentation is the process of dividing an image into regions [1], often the first step in the image analysis process. It highlights the portion of the image with relevant information, called the Region of Interest (ROI). However, different portions of images can reference the same object (parts), or the same objects can appear in different regions (plurality), thus having different ROI with the same aspect of meaning. For these portions to be understood as the same object or parts of the same region, the Semantic Segmentation adds meaning to the pixel [2], resulting in a pixel-wise classification problem. Badrinarayanan et al. [3] brings up a series of recent studies exploring pixel labeling classification.

This pixel-level precision enables machines to discern objects, their boundaries, and their relationships within the scene. The

applications of Semantic Segmentation are wide-ranging, from autonomous vehicles navigating complex environments to medical image analysis for disease detection and treatment planning, as well as in augmented reality, where it enriches user experiences by accurately identifying and interacting with objects in the real world. Moreover, it facilitates the development of advanced Artificial Intelligence (AI) systems for tasks like object recognition, scene understanding, and even human-computer interaction.

In studies in the biological area, semantic segmentation has been used in skin wounds and is an important research area for evaluating the healing process in animals and humans [4]. The use of Machine Learning (ML) algorithms has been shown to be effective in identifying and delimiting wound areas, allowing more accurate and objective analysis of the healing process. Strategies have been proposed using traditional ML and Deep Learning (DL) algorithms [5].

Although DL models have achieved remarkable results across various domains, it is crucial to acknowledge that they can be computationally intensive and time-consuming, mainly when working with large datasets and complex architectures. This computational burden can hinder rapid prototyping and model iteration. Hence, exploring alternatives with traditional ML algorithms remains valuable. Traditional methods, such as Decision Trees (DTs), Support Vector Machines (SVMs), and Random Forests (RF), often offer faster training times, better interpretability, and require fewer computational resources. Moreover, they can be highly effective for tasks with limited data or when the interpretability of results is a priority. In many practical scenarios, a judicious blend of DL and traditional ML algorithms can yield efficient and effective solutions, balancing performance and computational demands.

This study explores Semantic Segmentation to solve a mice wound pixel classification problem. Experiments were carried out with two different approaches: i) a DL architecture - U-Net, and ii) a traditional ML algorithm - Random Forest (RF) using latent representations of a pre-trained Convolutional Neural Network (Visual Geometry Group - VGG [6]). The image dataset used in the

experiments is composed of wounds in mice, generated by the Laboratory of Research in Pain, Neuropathy and Inflammation of the State University of Londrina (UEL).

This article is organized as follows: Section 2 presents the theoretical concepts necessary for understanding the work and related works found in the literature; Section 3 describes the proposed methodology for comparing and evaluating the segmentation algorithms; Section 4 presents the experimental results; and finally, Section 5 presents conclusions and suggestions for future work.

2 BACKGROUND

Semantic segmentation is a vital aspect of computer vision, involving identifying and classifying individual pixels within an image. Its applications span diverse fields like autonomous driving and medical image analysis. The conventional methods for this task demand substantial amounts of meticulously annotated data, a time-consuming and expensive process. Most semantic segmentation networks commonly utilize cross-entropy as their loss function and assess network performance using the intersection-over-union (IoU) metric [7].

2.1 U-Net

The U-Net algorithm is a Convolutional Neural Network (CNN) architecture developed for image segmentation, which has been shown to be efficient in several applications in the medical context [8, 9]. The network design consists of a symmetric network with descending *pooling* layers to encode low-level information into a high-dimensional representation, ascending convolution (or transposed convolution), and up-sampling layers to reconstruct the segmented image.

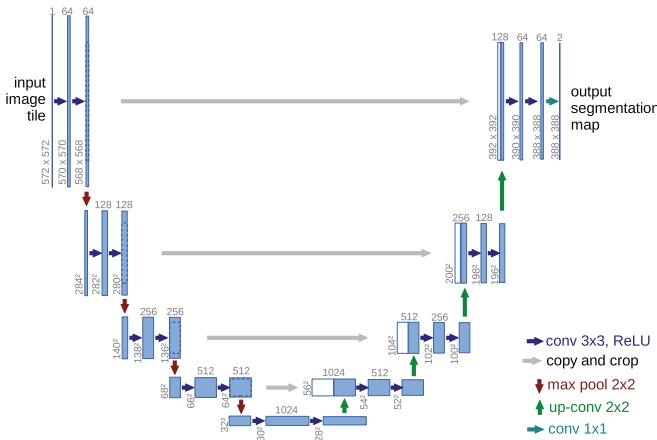


Figure 1: U-Net architecture. The half of the figure (*encoder*) reduces the original images to some latent representation. The second half (*decoder*) reconstructs images using latent features to classify pixel labels. Figure from Ronneberger et al. [10].

The technique of concatenating information from the decoding layer with the corresponding encoding layer allows the combination of low and high-level information, which provides a more

accurate representation of the original image. Figure 1 shows the U-Net architecture, composed by a *encoder* and a *decoder*. The *encoder* is usually a pre-trained classification network like VGG/ResNet models applying convolution blocks followed by a maxpool down-sampling to encode the input image into feature representation at multiple levels.

The decoder is the second part of the architecture. The goal is to semantically project the discriminative (latent) features with lower resolution learned by the encoder onto the pixel space to get a dense classification. The decoder consists of upsampling and concatenation operations followed by regular convolutions. U-Net also uses ReLU-like activation functions for all convolutional layers and a sigmoid activation layer at the end to produce the binary output of the segmentation.

2.2 Random Forest

Random Forest (RF) [11] is a traditional ML algorithm that uses multiple Decision Trees (DTs) [12, 13] to generate predictions. Figure 2 presents a schematic idea of the algorithm. Each tree is trained on a random sample of the training data and a random subset of the attributes, creating different data views. The final prediction is made by polling the predictions of all the trees, making the model more robust and less susceptible to overfitting. RF is widely used in classification and regression applications such as image recognition, sentiment analysis, and fraud detection [14].

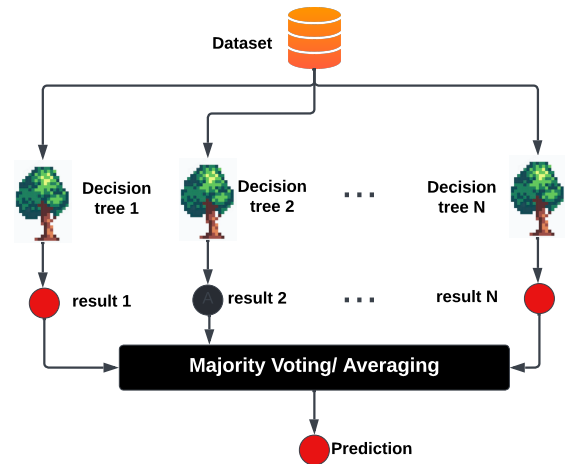


Figure 2: Random Forest ensemble. The *forest* comprises several random decision trees. Each tree is slightly different, handling a subset of features and examples. The prediction returned by the ensemble is the agreement of the committee.

Semantic image segmentation involves identifying and labeling each pixel in an image based on its semantic content, such as identifying each object or region present in the image. This is important in many areas, such as computer vision, robotics, and medical image processing. RF can be trained to classify each image pixel into different classes, such as object, background, or edge,

using texture, color, and shape information. The result is precise image segmentation that can be used in different applications.

2.3 Related works

A recent overview of the field of Semantic Segmentation is provided by Liu et al. [15], emphasizing its significance in image processing and computer vision. In the study, the authors highlight the impact of Deep Neural Networks (DNNs) on the progress of segmentation methods and categorize methods into traditional ML and recent DL-based approaches.

Most of the studies in the literature have been using DL algorithms to solve semantic segmentation of wounds [16, 17]. In Wang et al. [18], the authors performed semantic segmentation of ulcers' images. Experiments were carried out with state-of-the-art models that allow pixelwise classification, such as U-Net, FCN-VGG16, Mask-RCNN, and MobileNetV2. Results reported that U-Net obtained an accurate FScore value (0.915) and the highest recall among all models (0.912).

Fully Connected Networks (FCN) were also studied in Kaymak et al. [19] for segmentation of skin lesions, an experimental study was carried out for segmentation for cancer detection that was compared with other networks such as U-Net and SegNet, same study object as in Mishra and Daescu [20], which, in addition to CNNs, tests the performance of the segmentation using Otsu's threshold, being a much simpler technique than neural networks.

In another study [21], the authors performed the task of segmenting diabetic foot ulcers (DFU) in small datasets. They evaluated different DL models, such as the U-Net, V-Net, and SegNet. They were evaluated by accuracy, IoU, and FScore metrics. The results were positive for U-Net, considering that the model outperformed all the other models on the three metrics (0.949, 0.948, and 0.972, respectively).

Also, aiming to segment DFU images, Wang et al. [22] compared the DL ConvNet model with the classic ML algorithm Support Vector Machine (SVM), using RGB values as features, evaluating pixel-wise accuracy and IoU. Results were positive for ConvNet in both metrics. The model obtained a value of 0.95 and 0.473 for pixel accuracy and IoU, respectively, against 0.776 and 0.264 for SVM.

Innani et al. [23] substitute U-Net encoding blocks with other networks (ResNet, InceptionResNetV2 and EfficientNets) in an encoder-decoder strategy for lesion segmentation also for cancer. However, as the objective of the work is to perform a multiclass classification task, there is a second step after the segmentation. The proposed approaches are evaluated quantitatively in terms of the accuracy, mean IoU, and Dice Similarity Coefficient measures. The proposed cascaded end-to-end DL-based approach was the best overall, with the classification accuracy of the lesion significantly improved because of prior segmentation.

Kang and Nguyen [24] developed a hybrid framework combining RF and DL that learns flexible filters using an iterative optimization algorithm, and segments input images using the learned representations. Experiments were performed in a hand segmentation dataset for hand-object interaction and using two other semantic segmentation datasets. The results show that the proposed method achieves real-time semantic segmentation using limited computational and memory resources.

3 EXPERIMENTAL METHODOLOGY

In this section, we detail the experimental methodology adopted in this article.

An overview of the flow of experiments, including sub-steps, is shown in Figure 3. The following sub-sections give additional details regarding them: the image dataset, data preparation, classification algorithms used, and models' training and evaluations.

3.1 Image Dataset

The dataset for this study consists of 71 images of mice with wounds on their backs. This is a real dataset developed by the Pain, Neuropathy, and Inflammation Laboratory at the State University of Londrina (UEL) with the approval of the UEL Ethics Committee on Animal Research and Welfare (process number 15654.2019.33). Data acquisition process had a protocol, but different cameras obtained images. Thus, 36 images have a resolution of 1024×768 pixels, while the remaining 35 images have 4032×3024 pixels.

It is important to mention that assessing the wound progression is outside the scope of this work, as this step requires more information than just detecting the wound region. Hence, this study focuses on pixelwise classification (semantic segmentation) between wounds and non-wounds. Figure 4 shows an example of an image and its corresponding labels and predictions for an animal. The classes (pixel labels) were defined manually using the Labelme tool¹.

3.2 Data Preprocessing and Augmentation

Since images may have different resolutions and neural network input layers require a specific image size, all the images were rescaled to a resolution of 256×256 pixels. The rescale operation was applied to original and mask (label) images. A normalization step followed it, modifying pixel labels from $[0, 255]$ to interval $[0, 1]$.

Due to the low number of images in the dataset (71), a Data Augmentation (DA) process was also required. DA is a technique widely used in Computer Vision to increase the amount and variety of data [25]. The technique involves applying transformations to the original data and generating new instances. The main advantages of using DA are:

- (i) models' precision improvement since it generates more significant variability in the data;
- (ii) cost reduction in data collecting and labeling, given that collecting medical images is challenging and complex work, as well as labeling them, this task is reduced as transformations are applied to the original images and masks simultaneously; and
- (iii) overfitting prevention, since the model tends to learn specific patterns present in the available training data when there is low variability in the data, losing its generalization power to new data.

A total of five transformations were considered in DA: (1) horizontal flip, (2) vertical flip, (3) rotation with a maximum angle of 35 degrees, (4) salt and pepper noise, and (5) image translation. With the original images, the dataset with DA has 426 images available for model induction. It is important to highlight that augmented

¹<https://github.com/wkentaro/labelme>

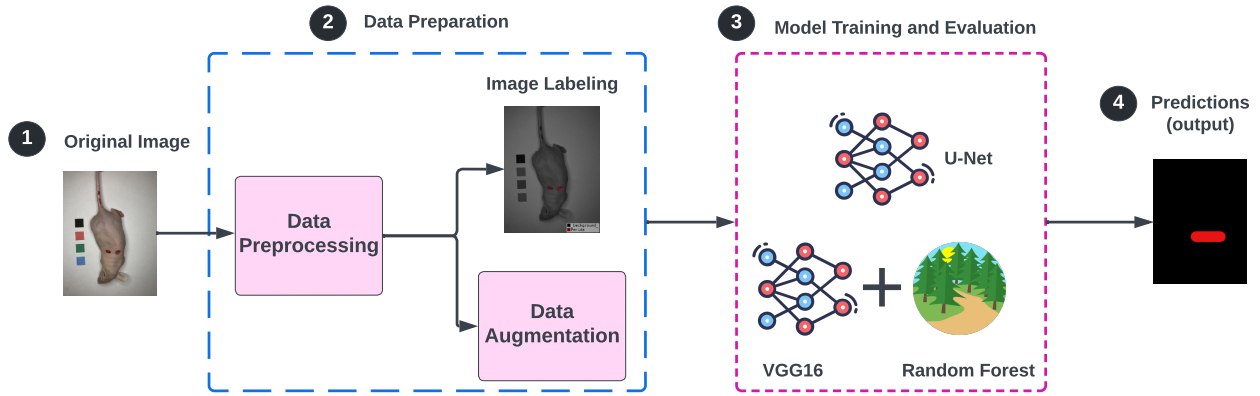


Figure 3: Experimental methodology for semantic segmentation of mice wounds.




Figure 4: Example of a dataset image instance and its correspondent labels.

images are used only in the training fold to induce the predictive models. Only the original images in the testing fold were assessed when evaluating these models.

3.3 Algorithms

We explored two different strategies for wound recognition:

- U-Net: a pure DL architecture widely adopted for Semantic Segmentation problems; and
- RF + VGG16: a hybrid and cheap alternative that extracts VGG16 latent features to train a Random Forest model.

Initial experiments were done using different Fully Convolutional Network (FCN) [2] architectures, but none of them recognized wound regions in the testing images. We only obtained good results when testing U-Net as the DL pure model. In fact, U-Net is one of the most widely used architectures for Semantic Segmentation of medical images [8].

U-Net was trained with Adam optimizer [26] and a learning rate of $\alpha = 0.0001$. The loss function optimized was the binary cross entropy since the network was used to perform binary classification. The U-Net architecture uses ReLu activation functions between

convolutional layers and the sigmoid activation function as a classifier in its last layer. U-Net was trained for 100 epochs using batch size 2. These values were defined empirically.

In the RF algorithm, the default value for the number of trees is between [100, 500] depending on the coding language and correspondent libraries/packages. However, all of them follow the same idea: the higher the number of trees, the better the generalization power of the induced model. In our experiments, the RF was implemented with the scikit-learn Python library, which uses $t = 100$ trees in the ensemble and 'Gini' index as the attribute evaluation criterion as default values.

Some initial experiments were performed using RGB pixel values for the RF algorithm, but the predictions did not identify any pixel as a wound. Thus, an alternative explored is using 64 feature maps extracted from the first two convolutional layers of the VGG16² neural network [6] to feed the algorithm. This DL model is pre-trained in the ImageNet dataset. VGG16 is well known for accurately performing computer vision tasks, including image recognition [27]. By extracting these feature maps, it is possible to obtain general information about the images that can be used to train machine learning models, such as the RF, with a lower computational cost.

²Visual Geometry Group, 16 layers.

3.4 Experimental Setup

The resampling strategy adopted in experiments was a repeated holdout, with five repetitions using different seeds. The dataset was split into training, testing, and validation sets, with proportions of 60%, 30%, and 10%, respectively. It is worth mentioning that the validation data was extracted as a part of the training set and not used in the testing set.

The results obtained were evaluated for precision and efficiency in the segmentation of wound images by F-Score, an evaluation metric that combines precision and recall to provide an overall measure of the model's accuracy, and IoU [28]. All the code was written in Python, with U-Net coded using PyTorch³. The U-net implementation used no pre-trained weights. The RF+VGG16 strategy used scikit-learn⁴ and Keras⁵ libraries. Experiments were executed on a desktop with a Ryzen 5 5600g processor, using a normal CPU environment, with 16GB of RAM, and motherboard b550m. All the algorithms used their default hyperparameter values, and no hyperparameter tuning process was conducted. The code repository of this study is publicly available⁶.

4 RESULTS

Figure 5 illustrates each algorithm's FScore and Intersection over Union (IoU) values and suggests that both algorithms present similar result distributions. The U-Net median values for FScore (0.75) and IoU (0.60) were slightly higher than RF ones (0.70 and 0.53). The figure also shows that most images have FScore values above 0.6. The overall results indicate that both algorithms could identify wounds regions and estimate the shape and wound filling even with a small dataset.

However, due to the nature of the distributions represented by the violins, some images presented difficulty for the algorithms. U-Net could not perform generalization for some images with tiny wounds. These images have a highly unbalanced pixel distribution, as few pixels compose wound regions. It may affect the FScore values but not the IoU. IoU only measures hits and misses on the object but not on the background, which is why it is usually used as an image segmentation evaluation metric. RF, on the other hand, has more extended distributions for both metrics. Violins indicate a considerable number of images with lower values of IoU and FScore, more than U-Net. A superficial analysis could identify that this occurred in some images when it fails to fill in the wound region or even incorrectly classifies other parts of the animal as a wound, such as regions of the paws or tail.

All the algorithms were executed in the same training and testing partitions and repeated with different seeds (varying from 1 to 5). Thus, the non-parametric Wilcoxon test with $\alpha = 0.05$ (95% of significance) was applied to assess the statistical significance of these results. The null hypothesis states that there is no difference in performance between the algorithms. The test was applied for both metrics (FScore and IoU) and obtained a p-value = 0.00101, which suggests a difference in terms of performance in favor of the U-Net.

³<https://segmentation-models.pytorch.readthedocs.io/en/latest/docs/api.html#unet>

⁴<https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html>

⁵<https://keras.io/api/applications/vgg/>

⁶<https://github.com/BrunoMarcato/MiceWoundSegmentation>

Although U-Net is statistically better than RF regarding predictions, the computational cost for training both algorithms differs. RF took 20 minutes in total to execute the five repetitions, while U-Net needed 9 hours. That is, the 64 feature maps extracted from the VGG16 and fed into the RF can be a viable alternative to speed the process of wound recognition rather than using a deep architecture, such as U-Net.

4.1 Prediction analysis

Figure 6 presents a heatmap with the obtained predictions of each algorithm in each image. The X-axis lists all the images, while the y-axis lists the evaluated algorithms (U-Net and RF). Heatmaps were generated for both performance measures (FScore and IoU). The higher the values, the more blue the cell is (and better). The lower the values, the more red the cell is (and worse). It is desired heatmaps with as much blue as possible.

Both measures show the same behavior: FScore values are higher in the same images where the correspondent IoU values are high. Good predictive values were obtained with images with ids = {1, 2, 3, 5, 6, 7, 66, 68, 69, 70, 71}. All of them are images with bigger wounds. An example can be seen in Figures 7a and 7b. Those images illustrated the U-Net predictions on the image with id = 5. Obtaining predictions are close to the desired area identified in the class mask (Figure 7b).

On the other hand, predictions were not accurate for images with ids = {28, 29, 38, 50, 52}, to name a few. Wounds in these images are tiny, with fewer pixels, and thus more challenging to identify. An example is depicted in Figures 7c and 7d, with the RF predictions on the image with id = 50 and its desired mask. Here it is possible to see that besides being unable to recognize the small regions of wounds, the RF algorithm wrongly identifies other parts of the animal as a wound.

The fact that RF uses initial layers of VGG ends up using high-level features such as color, thus giving importance to red regions, which is expected for wounds. However, this color is also present in other regions, as shown in the image (7c) with prediction generated by RF, which can be one of the causes of the loss of performance concerning U-Net.

5 CONCLUSIONS

This study investigated U-Net and RF with VGG latent features to solve a semantic segmentation problem of mice wounds. Experiments were carried out with an image dataset generated by the Pain, Neuropathy, and Inflammation Laboratory at the State University of Londrina, composed of 71 images showing wounds in mice. U-Net and RF+VGG were executed five times with different seeds splitting data into training and testing folds. Predictions were evaluated by F-Score and IoU metrics.

Results were promising: U-Net obtained F-Score and IoU average values of 0.75 and 0.60, respectively. Most images generated accurate predictions of the wound regions, with some exceptions in images with tiny wounds. The RF+VGG alternative was faster regarding the computational cost but statistically lower in performance, with lower average F-Score and IoU (0.70 and 0.53). However, it can be improved with some post-processing techniques to fill regions not wholly identified.

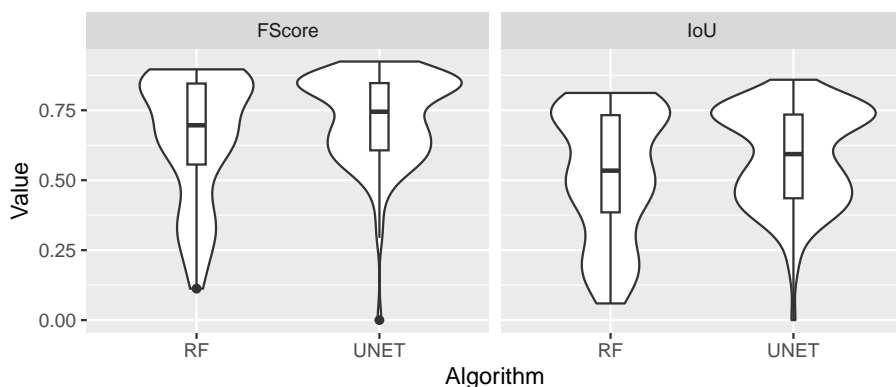


Figure 5: Overall results with FScore values obtained on test images.

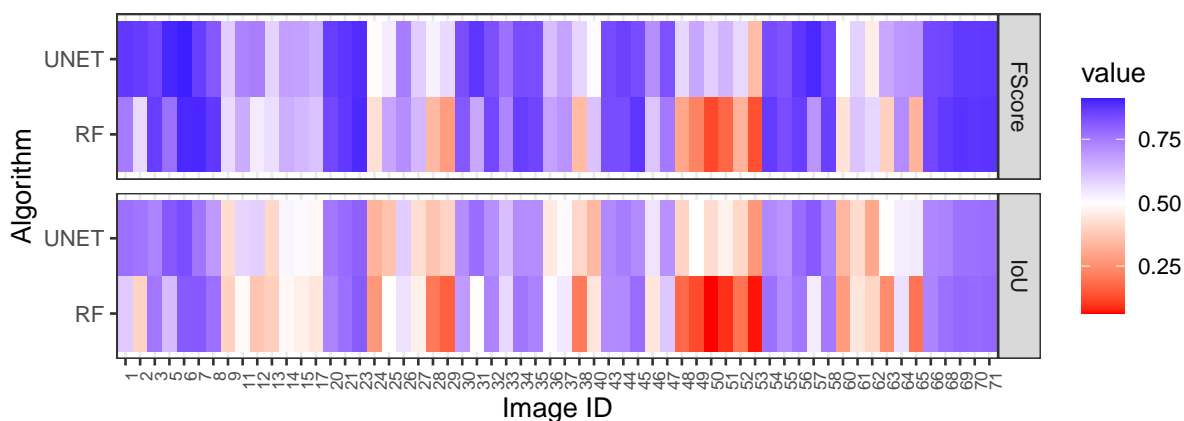


Figure 6: Predictions obtained by each algorithm in each image.




Figure 7: Examples of good and bad segmentation wound areas obtained in experiments.

An alternative might be replacing VGG16 with a more robust DL architecture, such as ResNet, or setting up the U-Net and RF with DL architectures with pre-trained weights by Transfer Learning. Other alternatives for future works include evaluating more and different DL architectures, ML algorithms and automating the entire pipeline.

ACKNOWLEDGMENTS

The authors would like to thank the Federal University of Technology – Parana (UTFPR), campus of Apucarana, and the Tempus Public Foundation through the Stipendium Hungaricum Scholarship Programme, for the financial support.

REFERENCES

- [1] R.C. Gonzalez and R.E. Woods. *Processamento Digital De Imagens*. Pearson, São Paulo, 2010. ISBN 9788576054016.
- [2] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, Boston, MA, USA, Jun 2015. IEEE. doi: 10.1109/CVPR.2015.7298965.
- [3] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12):2481–2495, 2017. doi: 10.1109/TPAMI.2016.2644615.
- [4] Eli Stevens, Luca Antiga, and Thomas Viehmann. *Deep Learning with PyTorch*. Manning Publications Co, Shelter Island, NY, 2020.
- [5] Charu C. Aggarwal. *Neural Networks and Deep Learning: A Textbook*. Springer International Publishing, Cham, 2018. ISBN 978-3-319-94463-0. doi: 10.1007/978-3-319-94463-0.
- [6] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, San Diego, CA, USA, 2015. CoRR.
- [7] Yifeng Huang, Zhirong Tang, Dan Chen, Kaixiong Su, and Chengbin Chen. Batching soft iou for training semantic segmentation networks. *IEEE Signal Processing Letters*, 27:66–70, 2020. doi: 10.1109/LSP.2019.2956367.
- [8] Narinder Singh Punn and Sonali Agarwal. Modality specific u-net variants for biomedical image segmentation: a survey. *Artificial Intelligence Review*, 55: 5845–5889, 2022. doi: 10.1007/s10462-022-10152-1.
- [9] Reza Azad, Ehsan Khodapanah Aghdam, Amelie Rauland, Yiwei Jia, Atlas Had-dadi Avval, Afshin Bozorgpour, Sanaz Karimijafarbigloo, Joseph Paul Cohen, Ehsan Adeli, and Dorit Merhof. Medical image segmentation review: The success of u-net, 2022.
- [10] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015.
- [11] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [12] L. Breiman, J.H. Friedman, R.A. Olshen, and C.J. Stone. *Classification and Regression Trees*. Chapman & Hall (Wadsworth, Inc.), 1984.
- [13] J. Ross Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993. ISBN 1558602402.
- [14] Gérard Biau and Erwan Scornet. A random forest guided tour. *TEST*, 25:197–227, 2016. doi: 10.1007/s11749-016-0481-7.
- [15] Xiaolong Liu, Zhidong Deng, and Yuhan Yang. Recent progress in semantic image segmentation. *Artificial Intelligence Review*, jun 2018. ISSN 1573-7462. doi: 10.1007/s10462-018-9641-3. URL [10.1007/s10462-018-9641-3](https://doi.org/10.1007/s10462-018-9641-3).
- [16] Zahra Mirikharaji, Kumar Abhishek, Alceu Bissoto, Catarina Barata, Sandra Avila, Eduardo Valle, M. Emre Celebi, and Ghassan Hamarneh. A survey on deep learning for skin lesion segmentation. *Medical Image Analysis*, 88:102863, 2023. ISSN 1361-8415. doi: 10.1016/j.media.2023.102863. URL <https://www.sciencedirect.com/science/article/pii/S1361841523001238>.
- [17] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghahfoorian, Jeroen A.W.M. van der Laak, Bram van Ginneken, and Clara I. Sánchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42:60–88, 2017. ISSN 1361-8415. doi: 10.1016/j.media.2017.07.005. URL <https://www.sciencedirect.com/science/article/pii/S1361841517301135>.
- [18] Chuanbo Wang, DM Anisuzzaman, Victor Williamson, Mrinal Kanti Dhar, Behrouz Rostami, Jeffrey Niezgoda, Sandeep Gopalakrishnan, and Zeyun Yu. Fully automatic wound segmentation with deep convolutional neural networks, 2020.
- [19] Ruya Kaymak, Cagri Kaymak, and Aysegul Ucar. Skin lesion segmentation using fully convolutional networks: A comparative experimental study. *Expert Systems with Applications*, 161:113742, 2020. ISSN 0957-4174. doi: 10.1016/j.eswa.2020.113742.
- [20] Rashika Mishra and Ovidiu Daescu. Deep learning for skin lesion segmentation. In *International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 1189–1194. IEEE, 2017. doi: 10.1109/BIBM.2017.8217826.
- [21] Rania Niri, Douzi Hassan, Lucas Yves, and Sylvie Treuillet. Semantic segmentation of diabetic foot ulcer images: Dealing with small dataset in dl approaches, 07 2020.
- [22] Changhan Wang, Xinchun Yan, Max Smith, Kanika Kochhar, Marcie Rubin, Stephen M. Warren, James Wrobel, and Honglak Lee. A unified framework for automatic wound segmentation and analysis with deep convolutional neural networks. In *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2415–2418, Milan, Italy, 2015. IEEE. doi: 10.1109/EMBC.2015.7318881.
- [23] Shubham Innani, Prasad Dutande, Bhakti Baheti, Ujjwal Baid, and Sanjay Talbar. Deep learning based novel cascaded approach for skin lesion analysis. In Deep Gupta, Kishor Bhurchandi, Subrahmanyam Murala, Balasubramanian Raman, and Sanjeev Kumar, editors, *Computer Vision and Image Processing*, pages 615–626, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-31407-0.
- [24] Byeongkeun Kang and Truong Q. Nguyen. Random forest with learned representations for semantic segmentation. *IEEE Transactions on Image Processing*, 28(7), 2019.
- [25] Suorong Yang, Weikang Xiao, Mengcheng Zhang, Suhan Guo, Jian Zhao, and Furao Shen. Image data augmentation for deep learning: A survey, 2022.
- [26] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [27] Alashiri Olaitan, Adeyinka Adewale, Sanjay Misra, Akshat Agrawal, Ravin Ahuja, and Jonathan Oluranti. Face recognition using vgg16 cnn architecture for enhanced security surveillance—a survey. In Pradeep Kumar Singh, Slawomir T. Wierchoń, Jitender Kumar Chhabra, and Sudeep Tanwar, editors, *Futuristic Trends in Networks and Computing Technologies*, pages 1111–1125, Singapore, 2022. Springer Nature Singapore. ISBN 978-981-19-5037-7.
- [28] Francisco J García-López, P B García-Allende, Olga M Conde, and José M López-Higuera. Comparison of evaluation metrics for hyperspectral image segmentation. *Sensors*, 19(7):1501, 2019.