Wiley
IET Biometrics
Volume 2025, Article ID 9982355, 6 pages
https://doi.org/10.1049/bme2/9982355

*Research Article*

# FingerUNeSt++: Improving Fingertip Segmentation in Contactless Fingerprint Imaging Using Deep Learning

**Laurenz Ruzicka**,[1] **Bernhard Kohn**,[1] **and Clemens Heitzinger**[2]

[1]*Austrian Institute of Technology, Digital Safety and Security, Giefinggasse 4, Vienna, Austria*
[2]*Center for Artificial Intelligence and Machine Learning (CAIML) and Department of Computer Science, TU Wien, Erzherzog-Johann-Platz 1, Wien, A-1040, Austria*

Correspondence should be addressed to Laurenz Ruzicka; laurenz.ruzicka@ait.ac.at

Received 12 September 2024; Revised 27 March 2025; Accepted 27 June 2025

Academic Editor: Martin Drahansky



Biometric identification systems, particularly those utilizing fingerprints, have become essential as a means of authenticating users due to their reliability and uniqueness. The recent shift towards contactless fingerprint sensors requires precise fingertip segmentation with changing backgrounds, to maintain high accuracy. This study introduces a novel deep learning model combining ResNeSt and UNet++ architectures called FingerUNeSt++, aimed at improving segmentation accuracy and inference speed for contactless fingerprint images. Our model significantly outperforms traditional and state-of-the-art methods, achieving superior performance metrics. Extensive data augmentation and an optimized model architecture contribute to its robustness and efficiency. This advancement holds promise for enhancing the effectiveness of contactless biometric systems in diverse real-world applications.

**Keywords:** biometrics; contactless; deep learning; fingerprint; segmentation

## 1. Introduction

Biometric identification systems play a vital role in modern security and authentication, owing to their reliability and the uniqueness of biometric traits. Among these systems, fingerprint recognition is one of the most widely utilized and trusted methods. Fingerprints provide a distinct pattern of ridges and valleys that remain consistent throughout an individual's life [1], making them an ideal biometric identifier for both personal identification and verification.

Recently, there has been a growing trend toward the use of contactless fingerprint sensors, driven by the demand for more hygienic, user-friendly, and flexible biometric solutions. Contactless systems remove the need for direct contact with the sensor, reducing the risk of spreading infectious diseases. These systems also offer increased versatility, making them suitable for various applications, such as mobile devices, public access points, and high-security environments [2–5], where user convenience and safety are critical.

For contactless fingerprint sensors to function effectively, precise detection and segmentation of the fingertip from captured images is essential. This segmentation step is critical for tasks such as pose correction [4, 6] and feature-based comparison [7], which rely on an accurate mask to distinguish the fingertip area from the background. The success of the segmentation process directly influences the overall accuracy and dependability of the system. Historically, fingertip segmentation has been achieved using a range of techniques, including color- or brightness-based methods [3, 8], machine learning algorithms [9, 10], and shape-based approaches [11].

Color-based methods exploit the distinct skin tone of fingertips, machine learning methods apply advanced algorithms to locate fingertip regions, while shape-based methods emphasize the geometric features of the fingertip.

In this research, we investigate the use of a grayscale sensor for capturing contactless fingerprint images and apply an object detection model to identify fingertips before processing them with a segmentation model.

Our study presents a novel deep learning model called FingerUNeSt++, based on ResNeSt and UNet++, which exceeds the current state-of-the-art in fingertip segmentation for contactless fingerprint images in terms of both segmentation accuracy and inference speed.

*1.1. Related Work.* Significant advancements in fingerprint recognition have transitioned from traditional contact-based methods to sophisticated contactless techniques. In our previous work [4], we enhanced cross-modality compatibility in fingerprint capture using pose correction [6] and unwarping techniques [12], introducing a deep learning–based approach for fingertip segmentation. This was then followed by a novel hand-based fingertip segmentation model, archiving superior performance over SOTA results [13].

Fingerprint segmentation, particularly in contactless images, is a critical step in the recognition process. Kauba et al. [3] examined smartphone-based fingerprint acquisition, focusing on segmentation methods to effectively isolate fingerprints from the background, allowing seamless comparison with contact-based datasets using low-latency color-based techniques.

Priesnitz et al. [14] explored the implementation of contactless fingerprint systems on mobile platforms, utilizing the fast Otsu thresholding method for segmentation.

Deep learning has revolutionized contactless fingerprint segmentation. Murshed et al. [15] employed convolutional neural networks (CNNs) to improve segmentation accuracy and robustness, training on large datasets to detect and segment fingerprint regions under diverse conditions.

The U-Net architecture [16], originally designed for medical image segmentation [17, 18], has been adapted for fingerprint segmentation tasks [4]. Variants like Squeeze-Unet [19] and EfficientUNet++ [20] have been tested, offering improved efficiency and performance. Squeeze-Unet incorporates squeeze-and-excitation blocks to enhance channel interdependencies, while EfficientUNet++ applies EfficientNet's scaling techniques to balance accuracy with computational efficiency. UNet++ [21] refines the skip connections in U-Net, creating a densely connected structure that boosts segmentation accuracy.

For feature extraction and classification, ResNeSt [22] and ResNet [23] are key architectures. ResNeSt (ResNet + split-attention networks) introduces split-attention blocks that enhance the model's ability to capture diverse feature representations, making it highly effective for complex image segmentation tasks. The ResNeSt-50d variant strikes a balance between depth and computational efficiency, providing a robust architecture for deep learning applications. ResNet, known for its residual connections, addresses the vanishing gradient problem in deep networks and serves as the backbone for many state-of-the-art models in image analysis. Despite its simplicity, Otsu thresholding remains a benchmark against which more advanced techniques are compared.

Several studies have improved the performance of contactless fingerprint recognition through enhanced segmentation. Labati et al. [24] utilized neural networks to correct perspective distortion and rotational variations, leading to more accurate fingerprint comparison, which relies on precise segmentation. Tan and Kumar [6] improved minutiae extraction and comparison by addressing pose variations, requiring accurate segmentation to calculate finger geometry. Chowdhury and Imtiaz [25] reviewed deep learning approaches, highlighting the critical role of robust segmentation in the recognition process.

*1.2. Contribution*

1. Combination of UNet++ with ResNeSt-50d: We integrate a simplified UNet++-like decoder with a ResNeSt-50d encoder, using different upscaling strategies, ReLU activation, and no dropout, enhancing feature extraction and segmentation accuracy.
2. Extended data augmentation: We augment the dataset to simulate perspective changes and various other changes, improving the model's robustness to variations in contactless fingerprint recordings.
3. Comparison with baseline and SOTA: We compare our model against established and state-of-the-art methods, demonstrating superior segmentation performance in all cases.

## 2. Methods

In this section, we detail the methodology used for fingertip segmentation in contactless fingerprint images. Our approach builds on an established preprocessing pipeline and detection model, described in [3], to precisely extract the fingertip regions from the hand image.

*2.1. Segmentation Using Baseline.* For a baseline segmentation to compare against, we employ Otsu thresholding, a widely used method for image binarization [3, 14]. As described in [14], Otsu's method calculates an optimal threshold value by maximizing the variance between two classes of pixels, effectively separating the foreground (fingertip) from the background.

This approach is straightforward and computationally efficient, making it suitable for real-time applications. However, it relies solely on pixel intensity distributions and does not account for more complex features such as texture or shape, which can limit its effectiveness in varying lighting conditions and diverse backgrounds commonly encountered in contactless fingerprint images. An example of this can be seen in Section 3.

*2.2. Segmentation Using Deep Learning.* The transition to deep learning methods for segmentation leverages the advancements in CNNs to address the limitations of traditional techniques. Unlike Otsu thresholding, which relies on simple pixel intensity distributions, deep learning models can capture complex patterns and features within the images. This capability allows for more robust and accurate segmentation, especially under diverse conditions typical of contactless fingerprint images.

*2.2.1. Preprocessing.* Before feeding the images into the deep learning model, it is crucial to standardize the input to ensure consistent learning of the right features and to aid with the generalization of the model.

Preprocessing of the extracted fingertips for test data involves only rescaling the images to $224 \times 224$ pixels.

For training data, preprocessing includes rescaling to $224 \times 224$ pixels and applying various augmentation techniques to increase the robustness of the model. These augmentations are applied randomly and in a random order to ensure the model learns to handle diverse variations in the data. The augmentation techniques include:

- Resize + crop: The image is resized with a factor between 0.75 and 1 and an aspect ratio of between 0.9 and 1.1, before it is randomly padded to $224 \times 224$ pixels.
- Horizontal flip: The image is flipped horizontally with a probability of 50%.
- Rotation: The image is randomly rotated from $-60°$ to $60°$.
- Perspective change: This technique simulates changes in the viewpoint by randomly distorting the image. This is achieved by transforming the image as if it was viewed from a different angle.
- Gaussian blur: A Gaussian blur is applied to the image, simulating varying degrees of focus and sensor noise.
- Solarize: This technique inverts all pixel values above a certain threshold. This creates high-contrast images, simulating the ridge-valley inversion.
- Posterize: The number of bits used to represent the pixel values is reduced, decreasing the number of shades of gray in the image. This simplification simulates low contrast recording, where the background is hard to separate from the fingertips.
- Histogram equalization: This method adjusts the contrast of the image by spreading out the most frequent intensity values. It is a common enhancing technique used to improve the visual appearance and downstream performance of fingerprints.

These augmentation techniques aim at enhancing the model's ability to generalize from the training data.

While the segmentation model is trained and applied to images rescaled to $224 \times 224$ pixels for computational efficiency and training stability, this does not limit the resolution of the final fingerprint data used for the next steps. The output of the segmentation network is a binary mask representing the fingertip region. This mask, which is initially at the $224 \times 224$ resolution, can be easily upsampled (e.g., using nearest-neighbor or bilinear interpolation) to match the dimensions of the original and high-resolution input image. This upsampled mask is then applied to the original image, effectively extracting the fingertip region at its full, native resolution (e.g., 500 ppi or higher).

*2.2.2. Architecture.* Our segmentation model FingerUNeSt++ combines elements from UNet++ and ResNeSt-50d to achieve high accuracy in fingertip segmentation. Here, we describe the key components of this architecture.

Table 1: Number of parameters and floating-point operations (FLOPs) for the different model parts and the whole model.

| Model part | Parameters | FLOPs |
|---|---|---|
| Encoder | $25.434 \times 10^6$ | $65.204 \times 10^9$ |
| Decoder | $25.478 \times 10^6$ | $0.48 \times 10^{12}$ |
| Segmentation head | 145 | $86.704 \times 10^6$ |
| Total | $50.911 \times 10^6$ | $0.545 \times 10^{12}$ |

*2.2.2.1. UNet and UNet++.* The UNet architecture consists of an encoder–decoder structure with skip connections. The encoder progressively reduces the spatial dimensions while increasing feature depth and the decoder reconstructs the image to its original size, using the skip connections to fuse high-resolution features from the encoder. UNet++ extends this concept by adding dense skip connections, creating a more complex network of connections between encoder and decoder. This helps capture more detailed and nuanced features.

*2.2.2.2. ResNeSt Family.* ResNeSt (ResNet with split attention) enhances the traditional ResNet by incorporating split attention blocks, which allow the network to focus on different feature subsets simultaneously. This results in better feature representation and improved performance on complex tasks. We chose ResNeSt over ResNet because it provides a more powerful feature extraction mechanism, crucial for handling the variability in contactless fingerprint images.

*2.2.2.3. Simplifications of Our UNet++ Implementation.* In our implementation, we simplify UNet++ by using the nearest-neighbor upscaling methods, ReLU activation functions instead of LeakyReLU, and removing dropout layers. These changes streamline the architecture while maintaining high performance, reducing computational complexity and training time. We utilize PyTorch and build on the code provided by Iakubovskii [26].

*2.2.2.4. Model Parameters and FLOPs.* The model's components, including the encoder, decoder, and segmentation head, each contribute to the overall parameter count and computational complexity. The details are given in Table 1.

*2.2.3. Training and Hyperparameter.* Our model is trained using the Jaccard loss function [27], also known as the intersection over union (IoU) loss, which is particularly effective for segmentation tasks. The Jaccard loss is defined as:

$$\text{Jaccard loss} = 1 - \left| \frac{A \cap B}{A \cup B} \right|, \tag{1}$$

where $A$ is the predicted segmentation mask and $B$ is the ground truth mask.

We utilize stochastic gradient descent (SGD) as the optimizer for training the model. The hyperparameters for SGD include a momentum of 0.9 to accelerate convergence and a learning rate of $8 \times 10^{-5}$.

Table 2: Comparison of final loss during training and validation.

| Model | Training loss | Validation loss |
|---|---|---|
| Baseline | — | $3.82 \times 10^{-1}$ |
| FingerUNeSt++ | $2.075 \times 10^{-2}$ | $2.095 \times 10^{-2}$ |

*Note:* The baseline is not trained, therefore, no training loss.

Table 3: Comparison of accuracy, mean intersection over union (MIoU), F1, and F2 scores for the different model architectures.

| Model | Accuracy (%) | MIoU (%) | F1 (%) | F2 (%) | Parameters |
|---|---|---|---|---|---|
| Baseline | 85.31 | 85.31 | 92.07 | 96.67 | — |
| EfficientUNet++[a] | 87.56 | 49.72 | — | — | $6.3 \times 10^6$ |
| Squeeze U-Net[a] | 96.42 | 85.73 | — | — | $\mathbf{2.5 \times 10^6}$ |
| Adapted U-Net[a] | 97.85 | 91.38 | — | — | $1.0 \times 10^7$ |
| FingerUNeSt++ | **99.45** | **99.31** | **99.65** | **99.52** | $5.1 \times 10^7$ |

*Note:* Bold indicates best performing values.
[a]The values were taken from Ruzicka et al. [4].

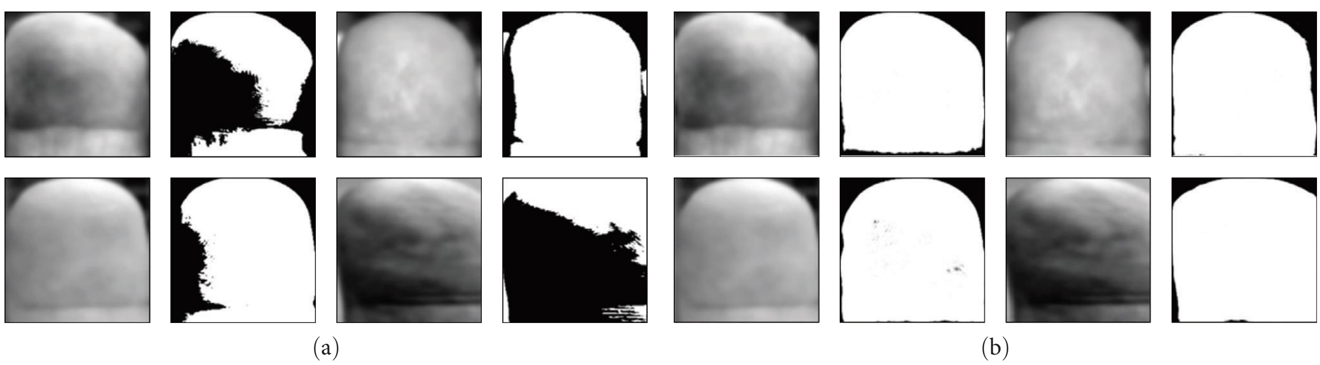

Figure 1: (a) Baseline segmentation results and input images. (b) UNet++ segmentation results, input images are blurred for privacy reasons.

*2.3. Experiment.* We utilize the same dataset as in [4], which consisted of 5828 manually annotated hand images from two different recording setups. The data is split into 1457 images for validation, 1822 images for testing, and 2549 images for training. All metrics are calculated on the test set, which was not shown to the model during training.

## 3. Results

This section presents the results of our segmentation model, comparing its performance against a baseline method as well as SOTA segmentation models.

The training and validation losses for the FingerUNeSt++ model and the validation loss for the baseline method are summarized in Table 2. The baseline method, which is not trained, shows a validation loss of 0.382. In contrast, the FingerUNeSt++ model achieves much lower training and validation losses.

Table 3 provides a detailed comparison of the performance metrics for different model architectures, including accuracy, mean IoU (MIoU), F1, and F2 scores. The FingerUNeSt++ model outperforms the baseline and all other architectures across all metrics, achieving an accuracy of 99.45%, MIoU of 99.31%, F1 score of 99.65%, and F2 score of 99.52%.

We also measured the inference time over a batched input and found a processing time of 12 ms on a laptop with a Quadro M3000M and a time of 0.2 ms on a workstation with a Nvidia GeForce RTX 3090.

Figure 1a,b shows qualitative comparisons of the segmentation results. The baseline method produces segmentation masks with noticeable errors, while the FingerUNeSt++ model delivers precise and accurate segmentation of the fingertip regions.

## 4. Discussion

Our study presents a significant advancement in the segmentation of fingertips in contactless fingerprint imaging by leveraging a novel deep learning approach. In this discussion, we will analyze the implications of our findings, the performance of our proposed model, and potential limitations and future directions.

The results demonstrate that our FingerUNeSt++ model with a ResNeSt-50d encoder substantially outperforms both the baseline Otsu thresholding method and other advanced segmentation models. Specifically, our model achieves an accuracy of 99.45%, a MIoU of 99.31%, an F1 score of 99.65%, and an F2 score of 99.52%. These metrics underscore the model's robustness and precision in segmenting fingertip regions from contactless fingerprint images.

The improvements can be attributed to several factors:

- Enhanced feature extraction: The ResNeSt-50d encoder effectively captures complex features due to its split attention blocks, which allow the network to focus on different feature subsets simultaneously. This results in superior feature representation, crucial for handling the variability in contactless fingerprint images.
- Dense skip connections: UNet++'s densely connected skip pathways facilitate the capture of fine-grained details by effectively combining high-resolution spatial information from the encoder with the decoder, enhancing segmentation accuracy.
- Extended data augmentation: The comprehensive augmentation techniques, including perspective change, Gaussian blur, and histogram equalization, significantly improve the model's generalizability to varying recording conditions and perspectives.

The accuracy of fingertip segmentation directly influences the performance of downstream tasks such as pose correction, feature-based comparison, and overall fingerprint recognition. While improved segmentation can enhance comparison by providing a cleaner input, its role extends beyond simply removing low-quality regions. The precise delineation of the fingertip boundary is essential for preprocessing steps, most notably pose correction. As described in [4, 6], pose correction algorithms rely heavily on the accurate segmentation mask to estimate the 3D geometry of the finger. The mask defines the contour used to construct a 3D model of the fingertip. Inaccuracies in the segmentation, such as including background pixels or excluding parts of the true fingertip, will directly translate to errors in the 3D model, leading to incorrect pose normalization and degraded comparison performance. Therefore, the segmentation is not primarily a quality filter, but a crucial geometric foundation. This is particularly important for contactless fingerprint systems, where variations in perspective and environmental conditions can introduce additional challenges.

While our model shows promising results, its inference time is a critical factor for real-time applications. The reported inference time of 12 ms on a laptop with a Quadro M3000M GPU and 0.2 ms on a workstation with an Nvidia GeForce RTX 3090 GPU indicates its potential for real-time use, though deployment on devices with less computational power may require further optimization.

Despite its high performance, the model's complexity and parameter count (50.9 million parameters) may pose challenges for deployment in resource-constrained environments. Future work could focus on model compression techniques, such as pruning and quantization, to reduce the computational load without compromising accuracy.

## 5. Conclusion

Our study introduces a novel deep learning approach for fingertip segmentation in contactless fingerprint imaging called FingerUNeSt++, combining the strengths of UNet++ and ResNeSt-50d. The proposed model demonstrates superior performance in terms of segmentation accuracy, making it a promising solution for enhancing the reliability and effectiveness of contactless fingerprint systems. Future work will be focused on optimizing the model for real-time applications and further improving its generalizability.

## Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

## Disclosure

This work was performed as part of the employment of Laurenz Ruzicka and Bernhard Kohn at the Austrian Institute of Technology.

## Conflicts of Interest

The authors declare no conflicts of interest.

## Funding

No funding was received for this manuscript.

## References

[1] S. Yoon and A. K. Jain, "Longitudinal Study of Fingerprint Recognition," *Proceedings of the National Academy of Sciences* 112, no. 28 (2015): 8555–8560.

[2] P. Birajadar, M. Haria, P. Kulkarni, et al., "Towards Smartphone-Based Touchless Fingerprint Recognition," *Sadhana* 44, no. 7 (2019): 161.

[3] C. Kauba, D. Söllinger, S. Kirchgasser, et al., "Towards Using Police Officers' Business Smartphones for Contactless Fingerprint Acquisition and Enabling Fingerprint Comparison Against Contact-Based Datasets," *Sensors* 21, no. 7 (2021): 2248.

[4] L. Ruzicka, D. Söllinger, B. Kohn, C. Heitzinger, A. Uhl, and B. Strobl, "Improving Sensor Interoperability Between Contactless and Contact-Based Fingerprints Using Pose Correction and Unwarping," *IET Biometrics* 2023 (2023): 7519499.

[5] A. Weissenfeld, R. Schmid, B. Kohn, B. Strobl, and G. Dominguez, "Case Study of the Acquisition of Contactless Fingerprints in a Real´ Police Setting," in *2022 International Conference of the Biometrics Special Interest Group (BIOSIG*, (IEEE, 2022): 1–5.

[6] H. Tan and A. Kumar, "Towards More Accurate Contactless Fingerprint Minutiae Extraction and Pose-Invariant Matching," *IEEE Transactions on Information Forensics and Security* 15 (2020): 3924–3937.

[7] S. A. Grosz, J. J. Engelsma, E. Liu, and A. K. Jain, "C2CL: Contact to Contactless Fingerprint Matching," arXiv: 2104.02811 [cs, eess] (2021).

[8] A. M. Bazen and S. H. Gerez, "Segmentation of Fingerprint Images," (2001).

[9] M. G. S. Murshed, K. Bahmani, S. Schuckers, and F. Hussain, "Deep Age-Invariant Fingerprint Segmentation System," arXiv: 2303.03341 [cs] (2023).

[10] J. Priesnitz, C. Rathgeb, N. Buchmann, et al., "Deep Learning-Based Semantic Segmentation for Touchless Fingerprint Recognition," in *Pattern Recognition. ICPR International Workshops and Challenges. Lecture Notes in Computer Science*, eds. Alberto Del

Bimbo, Rita Cucchiara, and Stan Sclaroff, et al., (Springer International Publishing, Cham, S., 2021): 154–168.

[11] X. Zheng, Y. Wang, and X. Zhao, "Fingerprint Image Segmentation Using Active Contour Model," in *Fourth International Conference on Image and Graphics*, (IEEE, 2007): 437–441.

[12] D. Söllinger and A. Uhl, "Optimizing Contactless to Contact-Based Fingerprint Comparison Using Simple Parametric Warping Models," in *2021 IEEE International Joint Conference on Biometrics (IJCB)*, (Shenzhen, China: IEEE, 2021), 1–7.

[13] L. Ruzicka, B. Kohn, and C. Heitzinger, "TipSegNet: Fingertip Segmentation in Contactless Fingerprint Imaging," *Sensors* 25, no. 6 (2025): 1824.

[14] J. Priesnitz, R. Huesmann, C. Rathgeb, N. Buchmann, and C. Busch, "Mobile Contactless Fingerprint Recognition: Implementation, Performance and Usability Aspects," *Sensors* 22, no. 3 (2022): 792.

[15] M. G. S. Murshed, S. K. Abbas, S. Purnapatra, D. Hou, and F. Hussain, "Deep Learning-Based Approaches for Contactless Fingerprints Segmentation and Extraction," arXiv: 2311.15163 [cs] (2023).

[16] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and ComputerAssisted Intervention – MICCAI 2015*, eds. Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F., Hrsg. Frangi, (Springer International Publishing, 2015): 234–241.

[17] H. Lu, Y. She, J. Tie, and S. Xu, "Half-UNet: A Simplified U-Net Architecture for Medical Image Segmentation," *Frontiers in Neuroinformatics* 16 (2022): 911679.

[18] X.-X. Yin, L. Sun, Y. Fu, R. Lu, and Y. Zhang, "U-Net-Based Medical Image Segmentation," *Journal of Halthcare Engineering* 2022 (2022): 4189781.

[19] N. Beheshti and L. Johnsson, "Squeeze U-Net: A Memory and Energy Efficient Image Segmentation Network," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, (Seattle, WA, USA: IEEE, 2020), 1495–1504.

[20] J. L. Silva, M. N. Menezes, T. Rodrigues, B. Silva, F. J. Pinto, and A. L. Oliveira, "Encoder-Decoder Architectures for Clinically Relevant Coronary Artery Segmentation," (2021).

[21] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, "UNet++: A Nested U-Net Architecture for Medical Image Segmentation," arXiv: 1807.10165 [cs, eess, stat] (2018).

[22] H. Zhang, C. Wu, Z. Zhang, et al., "ResNeSt: Split-Attention Networks," arXiv: 2004.08955 [cs] (2020).

[23] K. He, X. Zhang, S. Ren, and J. Sun, (Deep Residual Learning for Image Recognition, arXiv: 1512.03385 [cs], 2015).

[24] R. D. Labati, A. Genovese, V. Piuri, and F. Sotti, "Contactless Fingerprint Recognition: A Neural Approach for Perspective and Rotation Effects Reduction," in *2013 IEEE Symposium on Computational Intelligence in Biometrics and Identity Management (CIBIM)*, (IEEE, 2013): 22–30.

[25] A. M. M. Chowdhury and M. H. Imtiaz, "Contactless Fingerprint Recognition Using Deep Learning—A Systematic Review," *Journal of Cybersecurity and Privacy* 2, no. 3 (2022): 714–730.

[26] P. Iakubovskii, "Segmentation Models Pytorch," (2019).

[27] Z. Wang, X. Ning, and M. B. Blaschko, "Jaccard Metric Losses: Optimizing the Jaccard Index With Soft Labels," arXiv: 2302.05666 [cs] (2024).