

Scene Classification in Rural Environments for Autonomous Driving Using U-Net and ResNet

Victor M. A. do Nascimento^{1*}, André T. Cunha Lima^{1,2}

¹SENAI CIMATEC University, Salvador, Bahia, Brazil

²Federal University of Bahia, Institute of Physics, Department of Earth and Environmental Physics, Salvador, Bahia, Brazil

This paper addresses the semantic segmentation of road images, a critical task for autonomous vehicle navigation, particularly in non-urban environments that present significant challenges. While much research focuses on well-maintained roads in developed countries, this study confronts the complexities of real-world conditions, such as those prevalent in developing nations, which feature vast networks of unpaved and poorly maintained roads. The core of our methodology is a neural network architecture that synergistically combines the encoder-decoder structure of U-Net with the feature extraction power of a ResNet backbone. The primary objective is the precise classification of each image pixel into one of four essential categories for navigation: background, asphalt, paved, and unpaved road. The model's training regimen involved exploring different ResNet versions (ResNet18, ResNet34, and ResNet50) as the encoder backbone to assess the impact of network depth. A key aspect of our approach was a progressive training strategy, where model versions were trained on images of varying resolutions. The results demonstrated a significant and somewhat counter-intuitive finding: training the ResNet34-U-Net model with images at half the original resolution yielded the best overall performance, achieving the highest Dice and IoU scores. This suggests that reducing image resolution acts as an effective form of regularization, compelling the model to learn more general and robust features by ignoring minor, irrelevant details. This outcome not only enhances the model's generalization capabilities for diverse and imperfect road conditions but also carries a substantial practical advantage by reducing the computational cost of training and inference, a crucial factor for deployment on resource-constrained embedded systems in autonomous vehicles.

Keywords: Semantic Segmentation, Computer Vision, Neural Networks, Rural Roads

Detecting a navigable path is an important function in visual navigation systems, where several challenges must be considered, such as terrain variations, lighting changes, and the presence of potholes or water puddles. Data from the National Transport Confederation of 2022 [1] show that only 15.81% of the Brazilian road network consists of paved roads, highlighting the challenge for the widespread adoption of autonomous vehicles in regions distant from large urban centers. Studies of this kind often utilize road scenarios from developed countries, which typically exhibit little to no variation in terrain surface [2]. Some works involve systematic literature reviews: for instance, one 11-year review analyzed computer vision methods applied to this type of problem, emphasizing

the surface types these methods address, their adaptability to surface changes, and their capability to distinguish potential road defects or changes, such as potholes, shadows, and water puddles [3]. Convolutional Neural Networks (CNNs), the U-Net architecture, and Residual Networks (ResNet) have emerged as prominent deep learning architectures employed in this context. Research has primarily focused on improving semantic segmentation accuracy, acknowledging the ongoing challenges in achieving reliable environmental perception under diverse and complex driving conditions [4]. Convolutional Neural Networks (CNNs) represent a class of deep learning models whose fundamental architecture comprises convolutional layers, pooling layers, and activation functions. Convolutional layers act as feature extractors, learning hierarchical spatial patterns in the input data through the application of convolutional filters [4]. Pooling layers are responsible for reducing the spatial dimensionality of feature maps, providing translational invariance and decreasing sensitivity to minor input perturbations. Activation functions,

Received on 18 March 2026; revised 25 May 2026.

Address for correspondence: Victor do Nascimento. Av. Orlando Gomes, 1845 - Piatã, Salvador – BA – Brazil, Zipcode: 41650-010. E-mail: victor.nascimento@aln.senaicimatec.edu.

J Bioeng. Tech. Health 2026;9(6):572-578

© 2026 by SENAI CIMATEC University. All rights reserved.

such as ReLU (Rectified Linear Unit), introduce nonlinearity into the model, enabling it to learn complex relationships within the data [5]. As shown in Figure 1, the U-Net is a convolutional neural network architecture characterized by its encoder-decoder structure.

The architecture consists of a contracting path (encoder) and an expanding path (decoder). The contracting path follows the typical CNN architecture, comprising multiple convolutional layers followed by max-pooling layers. These progressively reduce the spatial resolution of the feature maps while increasing the number of feature channels. This process allows the encoder to capture high-level contextual information from the input image [5]. Residual Networks (ResNets) address the vanishing gradient problem that typically hinders the training of very deep networks [6]. The main contribution of ResNet is the introduction of residual blocks, which employ skip connections or residual connections, as shown in Figure 2.

Instead of directly learning the underlying mapping, the layers within a residual block learn a residual function with respect to the block's input. This residual function is then added to the original input via the skip connection. These skip connections allow gradients to flow more directly


through the network during backpropagation, facilitating the training of networks with hundreds or even thousands of layers [7].

Related Works

Semantic segmentation partitions an image into several semantically meaningful coherent parts and classifies each part into one of the predetermined classes. Existing semantic segmentation methods are unreliable for autonomous driving systems, as they ignore the differing levels of importance of distinct classes for safe driving. For instance, pedestrians, cars, and cyclists in a scene are far more critical than the sky or buildings, so segmentation must be as precise as possible. To incorporate the importance of object class information, the work present in [8] develops an "Importance-Aware Loss" (IAL), which emphasizes objects critical for autonomous driving. The IAL operates under a hierarchical structure where classes with varying importance are located at different levels, thus assigning a distinct weight to them.

Studies like [9] analyze semantic segmentation failures, which are crucial for autonomous driving systems, and detect failure cases in predicted segmentation maps by calculating the

Figure 1. U-Net network architecture [23]



the presence of shadows, other vehicles, people, or objects, and different road geometries. The work by [15] presents a review of articles published between 2005 and 2010 and also considers active vision sensors. It is a study focused on road detection, particularly lanes. Finally, authors in [16] presents work focused on road markings, concentrating on feature-based approaches like lane detection, stereoscopic analysis and edge-based segmentation.

Datasets for Autonomous Vehicle Navigation

Several benchmark datasets for road detection were analyzed: KITTI dataset by [17], CaRINA dataset by [18], CamVid dataset by [19], CityScapes dataset by [20], RTK dataset, and OffRoadScene dataset by [21]. One of the most widely used and cited datasets in path detection articles is the KITTI dataset, from the Karlsruhe Institute of Technology, Germany, featuring rural, urban, and highway scenarios. It presents situations with varying illumination. Most captures are on paved roads and it contains scenarios with many vehicles and pedestrians, as well as scenarios with little movement. For this work, the RTK dataset was used due to its well-curated road samples from Brazil, featuring different surface types such as asphalt variations, other types of pavements, and unpaved roads. It also includes situations with road damage (e.g., potholes).

Materials and methods

Dataset

The proposed method in this study addresses the semantic segmentation of road images using a neural network architecture that combines U-Net and ResNet. The objective is to classify each image pixel into one of the following categories: background, asphalt road, paved road, and unpaved road. The RTK dataset was used, which contains road images with different surface types and conditions. This dataset was chosen because it already has segmentation masks defined for each

class, considerably reducing the time required for creating a database of images and segmentation masks. The dataset was divided into three parts: training, validation, and test, to evaluate the model's performance. This division was performed randomly, ensuring that each part represented the diversity of the original dataset.

Architecture

The U-Net network was combined with a ResNet backbone, which acts as an encoder to extract relevant features from the images. ResNet, proposed by researchers from Microsoft Research, uses residual blocks and skip connections to mitigate the vanishing/exploding gradient problem, enabling the training of deeper and more effective networks [22]. Different versions of ResNet (ResNet18, ResNet34, ResNet50) were explored as the model's backbone, allowing an evaluation of the impact of network depth on road segmentation.

Preprocessing and Data Augmentation

The images were resized to a standard resolution of 512 x 512 pixels, along with data augmentation techniques such as vertical and horizontal flips (at 90°), rotations (at 30°), and {zooms (at 2x). These transformations aim to increase the variability of the training data, making the model more robust to different lighting, orientation, and scale conditions.

Model Training

Model training was initially performed with ResNet18, using the three-cycle training strategy proposed by fast.ai. This technique involves training the model on progressive image resolutions (1:8, 1:4, 1:2, and 1:1) over three phases. This procedure aims to initially train general features and shapes, followed by object features, and finally, object textures. Learning was optimized using the fast.ai's fit one cycle method, which adaptively adjusts the learning rate and momentum. The learning rate was determined using the fast.ai learning rate

find function, which helps identify the optimal learning rate for training. Finally, the loss function used was Cross Entropy Loss, suitable for multi-class classification problems, along with the Adam optimizer, known for its efficiency in training deep neural networks.

Evaluation Metrics

The model's performance was evaluated using the Dice and Jaccard metrics, which measure the similarity between the predicted segmentation masks and the ground truth masks. The Jaccard Coefficient, also known as Intersection over Union (IoU), is calculated as the ratio between the area of overlap and the area of union of the predicted and ground truth masks. The Dice metric is similar to IoU but gives more weight to true positives. In semantic segmentation tasks, the objective is to classify each pixel in an image, assigning it a label that represents the class of the object to which it belongs. To evaluate how well the model performs this task, metrics are used that compare the model's predicted segmentation with the actual segmentation (or ground truth).

In the context of image segmentation, these sets are the predicted segmentation mask (the model's output) and the ground truth segmentation mask. Mathematically, the IoU is defined as:

$$IoU = \frac{|A \cap B|}{|A \cup B|} \quad (1)$$

Where:

A represents the set of pixels belonging to a specific class in the predicted mask;

B represents the set of pixels belonging to the same class in the ground truth mask;

$|A \cap B|$ represents the number of pixels present in both A and B (the intersection);

$|A \cup B|$ represents the number of pixels present in A or in B or in both (the union).

The IoU measures the ratio between the area of overlap of the predicted and ground truth segmentations and the total area covered by both segmentations. An IoU of 1 indicates perfect overlap, while an IoU of 0 indicates no overlap.

The Dice metric, also known as the Dice Coefficient, is another commonly used metric to evaluate the similarity between two images. It is very similar to IoU, but differs slightly in its formulation:

$$Dice = \frac{2 \cdot |A \cap B|}{|A| + |B|} \quad (2)$$

Where A , B , $|A \cap B|$, $|A|$ and $|B|$ have the same meaning as in the IoU formula. The main difference between the Dice metric and IoU is that the Dice metric gives more weight to true positives (the correctly classified pixels). This is because the intersection is weighted twice as much, making it tend to be more sensitive to small details, and may be more appropriate when evaluating the segmentation accuracy of small objects.




Results and Conclusion

By comparing U-Net and ResNet architectures, as shown in Table 1, we demonstrated that the combination of a ResNet34 encoder with a U-Net decoder, trained on images at half the original resolution, yielded the best performance, achieving the highest Dice and IoU scores. This result indicates a robust capability for classifying pixels into essential categories such as asphalt, paved, and unpaved roads, as shown in Figures 3, 4 and 5.

The finding that a lower image resolution produced superior results is a significant outcome of this work. By decreasing the level of detail, the model is forced to learn more general and robust features of the road, minor texture variations, or irrelevant details in the vegetation. This not only improves the model's generalization capacity for varied and poorly maintained roads but also offers a considerable practical advantage: it reduces the computational cost of both training and inference, a critical factor for real-time application in embedded systems with limited processing power. Future work will be directed towards two main fronts. First, the cross-validation of the best-performing model (ResNet34 at 1/2 resolution) on a wider range of datasets to ensure its robustness and reliability under different conditions. Second, the

Table 1. Resolutions and performances.

Architecture	Resolution	Dice	IoU
ResNet18	1/8	0.615	0.607
ResNet18	1/4	0.641	0.615
ResNet18	1/2	0.654	0.633
ResNet18	1/1	0.655	0.634
ResNet34	1/8	0.616	0.610
ResNet34	1/4	0.635	0.610
ResNet34	1/2	0.656	0.635
ResNet34	1/1	0.636	0.615
ResNet50	1/8	0.622	0.619
ResNet50	1/4	0.643	0.618
ResNet50	1/2	0.646	0.625
ResNet50	1/1	0.650	0.630

Figure 3. Results for paved road**Figure 4.** Results for paved and unpaved road**Figure 5.** Results for asphalt and unpaved road

optimization of the architecture for deployment on embedded hardware, seeking a balance between accuracy and inference speed. As discussed in the literature, for safety-critical systems, understanding when the model is uncertain or likely to fail is as important as its accuracy. Implementing techniques to predict segmentation failures could significantly increase the safety and reliability of the autonomous navigation system, bringing it closer to practical

and widespread use in challenging environments like those found in Brazil.

References

- Confederação Nacional do Transporte. Malha rodoviária total [Internet]. Brasília: CNT; 2025.
- Kyem DA, Denteh I, Asamoah J, Tutu AA, Aboah C. Advancing pavement distress detection in developing countries: a novel deep learning approach with locally-collected datasets [Internet]. arXiv:2408.05649; 2024.
- Rateke T, Justen K, Chiarella V, Sobieranski A, Comunello E, Wangenheim A. Passive vision region-based road detection: a literature review. *ACM Comput Surv.* 2020;52(31):1-34.
- Feng D, Haase-Schütz C, Rosenbaum L, Hertlein H, Gläser C, Timm F, et al. Deep multi-modal object detection and semantic segmentation for autonomous driving: datasets, methods, and challenges. *IEEE Trans Intell Transp Syst.* 2021;22(3):1341-1360. doi:10.1109/TITS.2020.2972974.
- Ronneberger O, Fischer P, Brox T. U-Net: convolutional networks for biomedical image segmentation. In: Navab N, Hornegger J, Wells WM, Frangi AF, editors. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Cham: Springer; 2015. p. 234-241. doi:10.1007/978-3-319-24574-4_28.
- Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. In: Pereira F, Burges CJC, Bottou L, Weinberger KQ, editors. *Advances in Neural Information Processing Systems 25*. Red Hook (NY): Curran Associates; 2012. p. 1097-1105.
- He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*; 2016 Jun; Las Vegas, NV, USA. p. 770-778. doi:10.1109/CVPR.2016.90.
- Chen B, Gong C, Yang J. Importance-aware semantic segmentation for autonomous vehicles. *IEEE Trans Intell Transp Syst.* 2019;20(1):137-147.
- Song J, Ahn W, Park S, Lim M. Failure detection for semantic segmentation on road scenes using deep learning. *Appl Sci.* 2021;11(4):1870.
- Cheng C, Knoll A, Liao H. Safety metrics for semantic segmentation in autonomous driving. In: *2021 IEEE International Conference on Artificial Intelligence Testing (AITest)*; 2021. p. 57-64.
- Borba F. Redes neurais profundas e ensemble de classificadores: uma aplicação em imagens [dissertação]. Campinas (SP): Universidade Estadual de Campinas; 2021.
- Shengyan Z, Jianwei G, Huiyan C, Iagnemma K. Road detection using support vector machine based on online learning and evaluation. In: *IEEE Intelligent Vehicles Symposium*; 2010. p. 256-261.
- McCall J, Trivedi M. Video-based lane estimation and tracking for driver assistance: survey, system, and evaluation. *IEEE Trans Intell Transp Syst.* 2006:20-37.
- Yenikaya S, Yenikaya G, Düven E. Keeping the vehicle on the road: a survey on on-road lane detection systems. *ACM Comput Surv.* 2013:1-43.
- Hillel A, Lerner R, Levi D, Raz G. Recent progress in road and lane detection: a survey. *Mach Vis Appl.* 2014:727-745.
- Shin B, Xu Z, Klette R. Visual lane analysis and higher-order tasks: a concise review. *Mach Vis Appl.* 2014:1519-1547.
- Geiger A, Lenz P, Urtasun R. Vision meets robotics: the KITTI dataset. *Int J Robot Res.* 2013:1231-1237.
- Shinzato PY, Grassi V, Osorio FS, Wolf DF. Fast visual road recognition and horizon detection using multiple artificial neural networks. In: *IEEE Intelligent Vehicles Symposium*; 2012. p. 1090-1095.
- Brostow GJ, Julien F, Cipolla R. Road area detection based on texture orientations estimation and vanishing point detection. In: *SICE Annual Conference*; 2013. p. 1138-1143.
- Cordts M, Omran M, Ramos S, Rehfeld T, Enzweiler M, Benenson R, et al. The Cityscapes dataset for semantic urban scene understanding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; 2016. p. 3213-3223.
- Shang E, Zhao J, Li J, An X, Wu T. Offroadscene: an open database for unstructured road detection algorithms. In: *International Conference on Computer Sciences and Applications*; 2013. p. 779-783.
- He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition [Internet]. arXiv:1512.03385; 2015.
- NVIDIA. Deep learning fundamentals course [Internet]. [cited 2025 May 1]. Available from: <https://www.nvidia.com/pt-br/training/instructor-led-workshops/deep-learning-fundamentals/>