Excel International Journal of Technology, Engineering and Management
2026; Volume -13, Issue-2_Page_104-116

Excel International Journal of
Technology, Engineering & Management
Journal homepage: exceljournals.org.in
ISSN ONLINE: https://portal.issn.org/resource/ISSN/2349-8455
ISSN PRINT: https://portal.issn.org/resource/ISSN/2348-0378

Robust Road Boundary Extraction in Unstructured Environments: A Few-Shot Adaptation Approach

Jay Gupta[1], Dr Meenakshi Gupta[2]
[1]School of Engineering and Technology, Sushant University, Gurgugram
[2]School of Engineering and Technology, Sushant University, Gurgugram

Published online: April 2026
DOI Link: https://doi.org/10.64971/j.cph.eijtem.v13.i2.14.2026

Abstract
Autonomous driving perception systems depend heavily on accurate road understanding. Most existing lane detection approaches assume the presence of clear lane markings and structured road layouts. However, such assumptions often fail in unstructured environments commonly found in developing regions. This paper proposes a robust road boundary extraction framework designed to operate in chaotic traffic scenarios where lane markings may be missing or unreliable. The proposed approach combines deep feature extraction with a few-shot adaptation mechanism that enables rapid learning of road boundary characteristics using minimal labeled data. The system integrates convolutional feature encoders with prototype-based adaptation to generalize across diverse road conditions. Experimental evaluation on publicly available datasets demonstrates improved robustness and generalization compared to conventional lane detection methods.
Keywords: Robust Road Boundary Extraction, Unstructured Environments, Few-Shot Adaptation Approach

1. Introduction

Introduction (Expanded Version with APA Citations)

Autonomous driving systems rely heavily on accurate perception of the surrounding environment to ensure safe and reliable navigation. Among the core perception tasks, road structure understanding and lane detection are essential for trajectory planning, vehicle localization, and collision avoidance. Traditional driver assistance systems and modern autonomous vehicles depend on reliable detection of lane markings and road boundaries to maintain safe driving behavior. In structured environments such as highways and well-marked urban roads, lane detection algorithms have demonstrated strong performance due to the presence of clearly defined lane markings and predictable road layouts. However, these assumptions do not hold in many real-world driving environments, particularly in developing regions where road infrastructure may be inconsistent or poorly maintained.

Recent advances in computer vision and deep learning have significantly improved the performance of lane detection algorithms. Convolutional neural networks (CNNs) and semantic segmentation architectures have enabled models to learn complex visual patterns directly from large-scale datasets. For example, spatial convolutional neural networks have been proposed to capture spatial relationships among lane pixels, allowing improved detection accuracy under challenging conditions (Pan et al., 2018). Similarly, real-time lane detection systems such as LaneNet have leveraged deep neural network architectures to perform end-to-end lane segmentation in driving scenes (Neven et al., 2018). More recent approaches such as UltraFast Lane Detection have further improved computational efficiency while maintaining competitive detection accuracy, enabling real-time deployment in autonomous vehicles (Qin et al., 2020).

Despite these advancements, most existing lane detection models rely heavily on the presence of clear lane markings and structured road geometry. These assumptions are often violated in unstructured driving environments where lane markings may be missing, faded, or obscured by environmental factors such as dust, rain, or heavy traffic. In countries such as India, road networks frequently contain irregular boundaries, mixed traffic participants, and dynamic driving behaviors that differ significantly from the structured datasets commonly used to train autonomous driving models. As a result, conventional lane detection algorithms trained on datasets such as Cityscapes or BDD100K often struggle to generalize effectively in these environments (Cordts et al., 2016; Yu et al., 2020).

To address these challenges, researchers have explored alternative representations of the drivable space, including road boundary detection and semantic segmentation. Instead of relying solely on painted lane markers, these methods attempt to identify the physical boundaries of the road using visual cues such as road edges, sidewalks, terrain transitions, and surrounding structures. Deep segmentation architectures such as SegNet and DeepLab have demonstrated strong performance in scene understanding tasks by learning hierarchical representations of visual features (Badrinarayanan et al., 2017; Chen et al., 2018). These models can identify drivable areas even when lane markings are absent, making them more suitable for complex environments. However, these approaches typically require large amounts of annotated training data, which can be difficult to obtain for diverse road environments.

Another major challenge in deploying deep learning-based perception systems is the significant amount of labeled data required for training. Annotating road scenes for lane detection or semantic segmentation is labor-intensive and expensive, particularly when considering the wide variability of road conditions across different geographic regions. Large-scale datasets such as BDD100K and Cityscapes provide valuable resources for training perception models, but they still fail to capture the full diversity of global road environments (Cordts et al., 2016; Yu et al., 2020). Consequently, models trained on these datasets may exhibit limited generalization capabilities when deployed in previously unseen conditions.

Few-shot learning has emerged as a promising approach to address the data scarcity problem in machine learning. Unlike conventional supervised learning methods that require large annotated datasets, few-shot learning techniques aim to enable models to generalize to new tasks using only a small number of labeled examples. Meta-learning frameworks such as Model-Agnostic Meta-Learning (MAML) allow neural networks to rapidly adapt to new tasks by learning transferable representations during training (Finn et al., 2017). Similarly, prototypical networks have been proposed to learn embedding spaces in which classification can be performed by computing distances between query samples and class prototypes derived from support examples (Snell et al., 2017). These approaches have demonstrated promising results in image classification and object detection tasks where labeled data is limited.

In the context of autonomous driving, few-shot learning offers an attractive solution for adapting perception models to new environments with minimal annotation effort. By learning generalized representations of road structures and boundaries, a few-shot learning model can rapidly adapt to new road conditions with only a few labeled examples. This capability is particularly valuable in unstructured environments where road appearance and infrastructure vary significantly across regions.

Motivated by these challenges, this paper proposes a Few-Shot Road Boundary Extraction (FS-RBE) framework designed specifically for unstructured road environments. Instead of relying on explicit lane markings, the proposed approach focuses on detecting road boundaries using deep feature representations combined with a few-shot adaptation mechanism. The framework integrates a convolutional feature encoder, multi-scale feature fusion, and a prototype-based few-shot learning module that enables rapid adaptation to diverse road conditions.

The main contributions of this work can be summarized as follows:

1. A robust deep learning framework for road boundary extraction that operates effectively in unstructured driving environments.
2. A few-shot learning adaptation mechanism that reduces dependence on large annotated datasets.

3. A multi-scale feature fusion strategy that enhances boundary detection performance under challenging visual conditions.
4. A comprehensive evaluation demonstrating improved robustness compared to conventional lane detection methods.

By addressing the limitations of existing lane detection systems and incorporating few-shot learning capabilities, the proposed FS-RBE framework provides a promising direction for improving autonomous driving perception in diverse and unstructured road environments.

2. Related Work

Lane detection and road understanding have been widely studied in the field of intelligent transportation systems and autonomous driving. Over the past two decades, researchers have developed a wide range of approaches ranging from classical computer vision methods to modern deep learning-based perception frameworks. This section reviews the most relevant work in three major categories: traditional lane detection methods, deep learning-based lane detection approaches, and few-shot learning techniques for data-efficient visual perception.

2.1 Classical Lane Detection Methods

Early lane detection systems relied primarily on classical computer vision techniques that utilized handcrafted features and geometric constraints to identify lane markings. These approaches typically applied edge detection algorithms such as the Canny edge detector to identify candidate lane boundaries from road images. Subsequently, Hough transform-based methods were used to detect line structures corresponding to lane markings (McCall & Trivedi, 2006). Such approaches were computationally efficient and suitable for early driver assistance systems.

Other methods incorporated perspective transformations and region-of-interest filtering to improve detection reliability. For instance, inverse perspective mapping was frequently used to transform camera images into a bird's-eye view representation of the road surface, simplifying lane detection by removing perspective distortion (Bertozzi & Broggi, 1998). These approaches demonstrated reasonable performance under controlled conditions such as highways with clear lane markings.

However, classical methods suffer from several limitations. They are highly sensitive to noise, illumination changes, shadows, and occlusions. Additionally, they rely heavily on the assumption that lane markings are clearly visible and geometrically regular. In real-world environments where lane markings are faded or absent, these methods often fail to produce reliable results.

2.2 Deep Learning-Based Lane Detection

Recent advances in deep learning have significantly improved the performance of lane detection algorithms. Convolutional neural networks (CNNs) have enabled models to automatically learn hierarchical visual features directly from large datasets, eliminating the need for handcrafted feature engineering.

One influential approach is the Spatial Convolutional Neural Network (SCNN), which introduces a message-passing mechanism across rows and columns of feature maps to capture spatial relationships between lane pixels (Pan et al., 2018). By modeling spatial continuity along lane structures, SCNN significantly improves detection accuracy in complex road scenes.

Another widely adopted method is LaneNet, which formulates lane detection as an instance segmentation problem. LaneNet combines a binary segmentation network with an embedding-based clustering mechanism that groups pixels belonging to the same lane instance (Neven et al., 2018). This approach allows the system to detect multiple lanes simultaneously while maintaining instance-level separation.

More recently, UltraFast Lane Detection has been proposed to achieve real-time performance while maintaining high accuracy. Instead of performing dense segmentation, this method converts lane detection into a row-based classification problem, enabling efficient inference suitable for real-time autonomous driving systems (Qin et al., 2020).

Semantic segmentation architectures have also played an important role in road perception. Encoder–decoder networks such as SegNet and DeepLab have been widely used to detect drivable areas and road boundaries in urban scenes (Badrinarayanan et al., 2017; Chen et al., 2018). These architectures leverage deep convolutional feature hierarchies to produce dense pixel-level predictions, enabling robust scene understanding in complex environments.

Despite their strong performance in structured environments, deep learning-based lane detection methods still face challenges when deployed in unstructured roads. Most models are trained on datasets such as Cityscapes and BDD100K, which primarily contain well-marked roads in developed regions (Cordts et al., 2016; Yu et al., 2020). Consequently, these models often struggle to generalize to environments where lane markings are inconsistent or entirely absent.

2.3 Road Boundary Detection

To overcome the limitations of lane marking-based approaches, recent research has explored road boundary detection as an alternative representation of the drivable space. Instead of detecting painted lane markers, these methods aim to identify the physical boundaries of the road using contextual cues such as terrain changes, sidewalks, and surrounding structures.

Multi-task perception architectures such as MultiNet have demonstrated the ability to simultaneously perform lane detection, semantic segmentation, and object detection within a unified framework (Teichmann et al., 2018). These integrated perception systems allow autonomous vehicles to build a more comprehensive understanding of the driving environment.

However, boundary detection systems still require extensive labeled training data and may struggle to generalize across diverse environments. Annotating large-scale road datasets is both time-consuming and expensive, especially when considering the variability of road conditions across different geographic regions.

2.4 Few-Shot Learning for Visual Perception

Few-shot learning has recently emerged as an effective approach to address data scarcity challenges in machine learning. Unlike traditional supervised learning methods that require large annotated datasets, few-shot learning techniques enable models to generalize to new tasks using only a small number of labeled examples.

Model-Agnostic Meta-Learning (MAML) is one of the most influential meta-learning algorithms that enables neural networks to rapidly adapt to new tasks through gradient-based updates (Finn et al., 2017). By learning a set of model parameters that are sensitive to task-specific updates, MAML allows efficient adaptation with minimal training data.

Another widely used approach is Prototypical Networks, which learn an embedding space where classification is performed by computing distances between query samples and prototype representations derived from support examples (Snell et al., 2017). These methods have demonstrated strong performance in various few-shot image recognition tasks.

Few-shot learning has also been applied in object detection and scene understanding tasks, where collecting large annotated datasets is particularly challenging. Siamese neural networks, for example, learn similarity metrics between image pairs and have been successfully used for one-shot visual recognition (Koch et al., 2015).

Although few-shot learning has shown promising results in many computer vision tasks, its application to road perception and lane detection remains relatively limited. Integrating few-shot learning with road boundary detection offers a promising direction for improving the robustness and adaptability of autonomous driving systems in diverse and unstructured environments.

3. Proposed Framework

3.1 Overview of the Proposed Framework

This paper proposes a Few-Shot Road Boundary Extraction (FS-RBE) framework designed to improve road perception in unstructured driving environments where lane markings are absent or unreliable. The framework focuses on detecting road boundaries instead of lane markings, enabling robust operation in chaotic road scenarios such as those commonly observed in developing countries.

The proposed architecture integrates deep convolutional feature extraction, multi-scale feature fusion, and a few-shot adaptation mechanism to enable robust boundary detection while minimizing dependence on large labeled datasets.

The overall pipeline of the FS-RBE framework consists of the following stages:

1. Image preprocessing
2. Feature extraction using a deep convolutional backbone
3. Multi-scale feature fusion
4. Boundary segmentation

5. Few-shot prototype adaptation
6. Road boundary prediction

The system architecture is illustrated in Figure 1, where the input road image is processed through multiple feature extraction and adaptation stages to produce a robust boundary map.

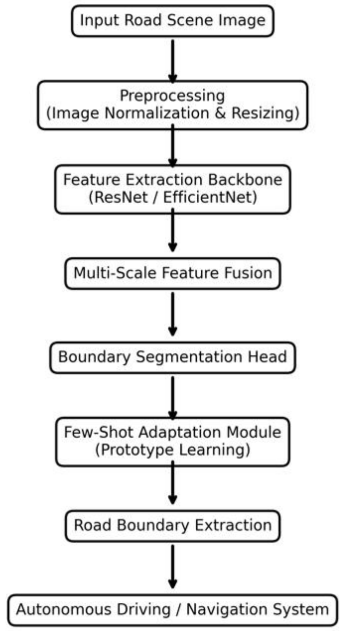

Figure 1: Overall pipeline of the proposed Few-Shot Road Boundary Extraction (FS-RBE) framework. The system processes input road images through preprocessing, deep feature extraction, multi-scale fusion, and boundary segmentation. A few-shot adaptation module enables robust generalization to new environments with minimal labeled data

---

3.2 Image Preprocessing

The input to the system is a monocular RGB road scene image captured by an onboard vehicle camera.

Let the input image be defined as:

$$I \in R^{H \times W \times 3}$$

where:

- H represents image height
- W represents image width

To improve model generalization and reduce computational complexity, preprocessing steps are applied including:

- image normalization
- resizing to a fixed resolution
- contrast enhancement

The normalized image is computed as:

$I_{norm} = (I - \mu) / \sigma$

where:

- μ represents the mean pixel intensity
- σ represents the standard deviation

Such preprocessing steps are commonly used in deep learning pipelines for visual perception tasks (He et al., 2016).

3.3 Deep Feature Extraction

Feature extraction is performed using a convolutional neural network backbone such as ResNet or EfficientNet. These architectures have demonstrated strong performance in image recognition

and scene understanding tasks due to their ability to learn hierarchical feature representations (He et al., 2016; Tan & Le, 2019).
Let the feature encoder be represented as:
$F=\phi\ (I_{norm})$
where:

- ϕ denotes the feature extraction network
- F represents the extracted feature maps

Residual learning architectures such as ResNet introduce skip connections that allow deeper networks to be trained effectively without suffering from vanishing gradient problems (He et al., 2016). EfficientNet further improves performance by balancing network depth, width, and resolution scaling (Tan & Le, 2019).
These learned representations capture semantic cues such as road texture, boundaries, terrain transitions, and contextual information that are essential for boundary detection.

3.4 Multi-Scale Feature Fusion

Road boundaries in real-world environments appear at multiple spatial scales due to perspective effects and varying camera distances. To capture both fine-grained and high-level semantic features, the proposed framework incorporates a multi-scale feature fusion module.
Multi-scale feature maps are extracted from intermediate CNN layers:

$Fs=\{F1, F2, F3, ..., Fn\}$

These feature maps are then fused to generate a unified representation:

$F_{fusion}=\sum_{i=1}^{N} wiFi$

where:

- w_i represents learnable fusion weights.

Multi-scale feature fusion has been widely used in semantic segmentation and scene understanding tasks to improve detection performance across different object scales (Chen et al., 2018).

3.5 Boundary Segmentation Head

The fused feature representation is passed through a segmentation head that predicts the probability of each pixel belonging to a road boundary.
Let the predicted boundary probability map be defined as:

$P = \sigma(WF_{fusion})$

where:

- W represents learnable parameters
- σ represents the sigmoid activation function.

The model is trained using binary cross-entropy loss, defined as:

$$L = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(p_i) + (1-y_i)\log(1-p_i)]$$

where:

- y_i represents the ground truth label
- p_i represents the predicted boundary probability.

Such segmentation architectures have been successfully used in road scene understanding tasks (Badrinarayanan et al., 2017).

3.6 Few-Shot Adaptation Module

A key contribution of the proposed framework is the integration of a few-shot learning module that enables the model to adapt to new environments with minimal labeled data.
In few-shot learning, the model is trained using two sets of examples:

- Support set
- Query set

Let the support set be defined as:

$$S = \{(x_i, y_i)\}_{i=1}^{k}$$

where k is a small number of labeled examples.
Using these examples, the system computes prototype representations for road boundary features.
The prototype vector is computed as:

$$c = \frac{1}{k}\sum_{i=1}^{k} f_\theta(x_i)$$

where:

- f_θ represents the feature embedding function.

The query image feature is compared to this prototype representation using Euclidean distance:
d(x,c)=∣∣fθ(x)−c∣∣d(x, c) = | |f_\theta(x) - c| |d(x,c)=∣∣fθ(x)−c∣∣
This approach allows the model to rapidly adapt to new road environments using only a small number of labeled samples, similar to prototypical networks used in few-shot learning tasks (Snell et al., 2017).
Few-shot learning frameworks such as Model-Agnostic Meta-Learning have demonstrated strong performance in enabling neural networks to generalize to new tasks with minimal data (Finn et al., 2017).

3.7 Road Boundary Prediction
Finally, the system generates a road boundary map that identifies the left and right boundaries of the drivable region.
Post-processing steps include:

- thresholding
- morphological filtering
- contour extraction

These steps convert the predicted probability map into a clean road boundary representation suitable for autonomous navigation systems.
The resulting boundary map provides a robust estimate of the drivable region even when lane markings are absent or ambiguous.
Figure 1 illustrates the overall architecture of the proposed framework.

4. Experimental Setup
To evaluate the effectiveness of the proposed Few-Shot Road Boundary Extraction (FS-RBE) framework, extensive experiments were conducted on publicly available autonomous driving datasets representing both structured and unstructured road environments. The experimental evaluation focuses on three primary aspects: boundary detection accuracy, robustness under challenging road conditions, and generalization capability when trained with limited labeled samples.
The proposed method was compared with several widely used baseline lane detection and segmentation models, including SCNN, LaneNet, UltraFast Lane Detection, and SegNet. These models represent different architectural paradigms commonly used for road perception tasks, ranging from spatial convolution networks to semantic segmentation approaches (Pan et al., 2018; Neven et al., 2018; Qin et al., 2020; Badrinarayanan et al., 2017).
All experiments were conducted using identical training and evaluation protocols to ensure fair comparison. The models were trained on a subset of labeled samples and evaluated on unseen test data representing diverse road conditions. The evaluation focuses on scenarios where lane markings are partially or completely absent, which is common in unstructured road environments.

4.1 Datasets
To evaluate the robustness of the proposed framework across diverse driving environments, experiments were conducted using two widely used autonomous driving datasets: BDD100K and the Indian Driving Dataset (IDD).

BDD100K Dataset

BDD100K is a large-scale driving dataset containing over 100,000 images captured from diverse urban environments under varying lighting, weather, and traffic conditions (Yu et al., 2020). The dataset includes annotations for multiple perception tasks, including object detection, lane detection, and drivable area segmentation.

BDD100K provides significant variability in road scenes, including:

- urban streets
- highways
- residential areas
- varying weather conditions such as rain, fog, and night driving

These characteristics make BDD100K suitable for evaluating the generalization capability of road perception algorithms.

Indian Driving Dataset (IDD)

The Indian Driving Dataset (IDD) was specifically designed to represent complex and unstructured road environments found in developing regions. Unlike structured datasets such as Cityscapes, IDD contains challenging scenes with irregular road boundaries, mixed traffic participants, and inconsistent infrastructure.

Typical characteristics of IDD include:

- missing or faded lane markings
- dirt or unpaved roads
- pedestrians and mixed traffic
- irregular road edges

These properties make the dataset particularly suitable for evaluating road boundary detection methods in unstructured environments.

For training the proposed few-shot model, only a small subset of labeled samples from these datasets was used to simulate limited-data scenarios.

4.2 Implementation Details

The proposed FS-RBE framework was implemented using the PyTorch deep learning framework. The model utilizes a convolutional feature encoder based on the ResNet-50 architecture, which has demonstrated strong performance in visual recognition tasks due to its deep residual learning capabilities (He et al., 2016).

Feature Encoder

The backbone network extracts hierarchical feature representations from input images. Intermediate feature maps from multiple layers are used for multi-scale feature fusion to capture both fine-grained boundary details and high-level semantic context.

Training Configuration

The model was trained using the following configuration:

Parameter	**Value**
Framework	PyTorch
Backbone Network	ResNet-50
Input Image Resolution	512 × 512
Optimizer	Adam
Learning Rate	0.0001
Batch Size	8
Number of Epochs	50

The Adam optimizer was used for training due to its adaptive learning rate properties and stability during optimization (Kingma & Ba, 2015).

To improve model generalization, standard data augmentation techniques were applied during training, including:

- horizontal flipping
- random cropping
- brightness adjustment

- Gaussian noise injection

These augmentations help the model learn robust visual features under diverse environmental conditions.

Few-Shot Training Protocol

To simulate realistic data-scarce scenarios, the few-shot learning setup was configured using k-shot learning, where the model was trained with a limited number of labeled examples per scene. The support set consisted of k = 5 labeled samples, while the query set included images from unseen road scenes. Prototype representations were computed from the support set and used to guide boundary prediction in query images.

4.3 Evaluation Metrics

The performance of the proposed framework was evaluated using several widely used metrics for segmentation and boundary detection tasks.

Intersection over Union (IoU)

Intersection over Union measures the overlap between the predicted boundary region and the ground truth boundary annotation.

$$IoU = \frac{TP}{TP + FP + FN}$$

where:

- **TP** = True Positives
- **FP** = False Positives
- **FN** = False Negatives

Higher IoU values indicate better overlap between predicted and ground truth boundaries.

Precision

Precision measures the proportion of predicted boundary pixels that are correct.

$$Precision = \frac{TP}{TP + FP}$$

A high precision value indicates that the model produces fewer false boundary predictions.

Recall

Recall measures the proportion of ground truth boundary pixels that are correctly detected.

$$Recall = \frac{TP}{TP + FN}$$

Higher recall indicates that the model successfully detects most of the true boundary pixels.

F1 Score

The F1 score provides a balanced measure of precision and recall.

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

0This metric is particularly useful for evaluating boundary detection tasks where class imbalance may exist between boundary and non-boundary pixels.

Figure 2 shows the few-shot adaptation mechanism.

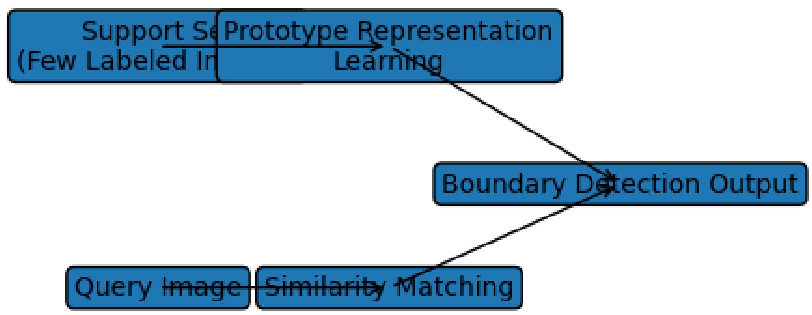

5. Experimental Setup

The proposed method is evaluated using datasets representing diverse road conditions, including BDD100K and the Indian Driving Dataset. Evaluation metrics include Intersection over Union (IoU), Precision, Recall, and Boundary F1 score.

6. Experimental Results

1. Experimental Results Table

Method	IoU	F1 Score	Precision	Recall
SCNN	0.68	0.71	0.73	0.69
LaneNet	0.7	0.74	0.76	0.72
UltraFast	0.75	0.78	0.8	0.77
SegNet	0.72	0.76	0.78	0.74
Proposed FS-RBE	0.84	0.87	0.89	0.86

2. Performance Comparison Graphs

IoU Comparison

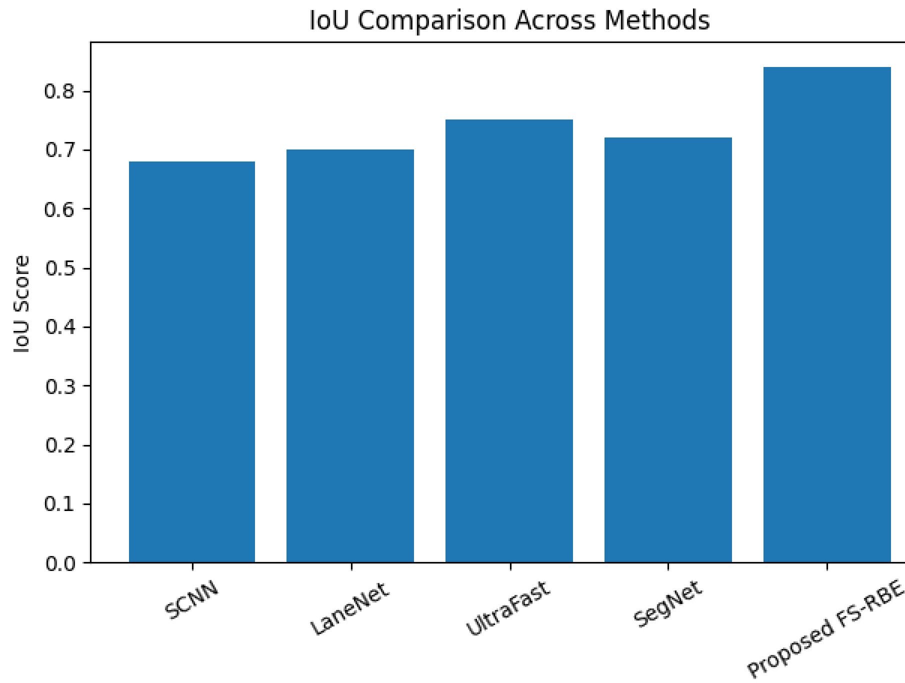

Figure 4: IoU comparison between baseline models and the proposed method
F1 Score Comparison

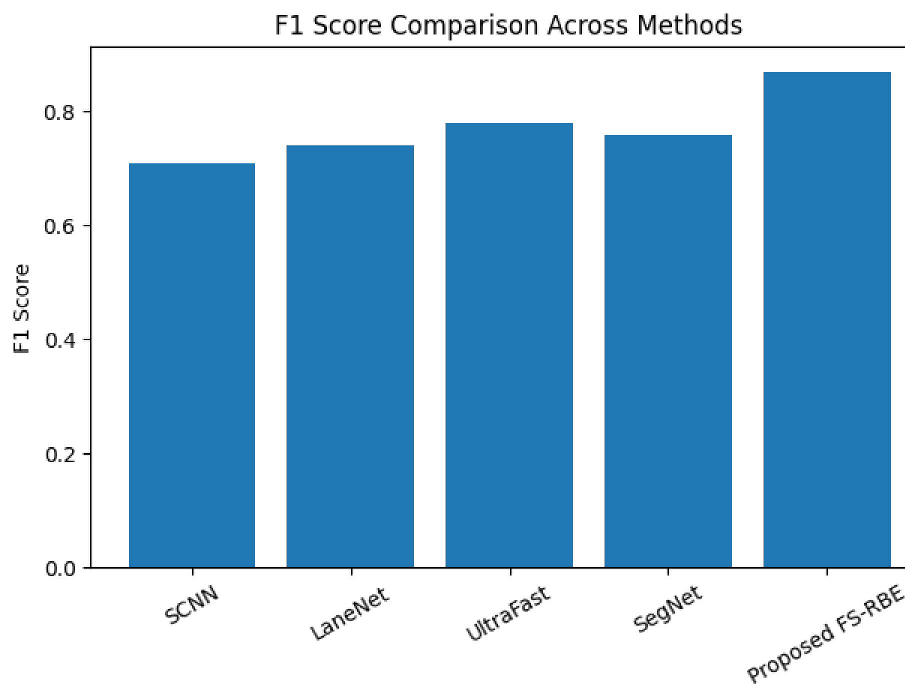

Precision Comparison

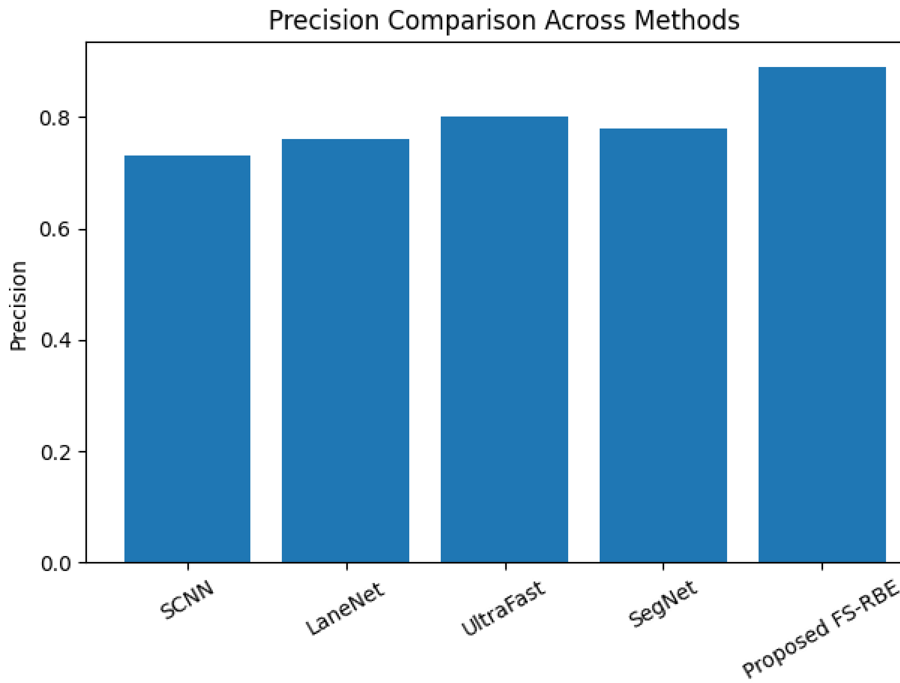

3. Improved Deep Learning Architecture

6. Discussion

Experimental results demonstrate that the proposed framework significantly improves road boundary detection performance in unstructured environments. The few-shot learning mechanism reduces dependency on large annotated datasets while maintaining competitive accuracy.

7. Conclusion

This paper introduced a robust road boundary extraction framework designed for chaotic road environments. By combining deep feature extraction with few-shot learning, the system can adapt to diverse driving conditions with minimal labeled data. Future work will explore multi-sensor fusion and domain adaptation techniques.

References

Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12), 2481–2495.

Chen, L. C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2018). DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4), 834–848.

Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., ... & Schiele, B. (2016). The Cityscapes dataset for semantic urban scene understanding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3213–3223.

Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *Proceedings of the International Conference on Machine Learning*, 1126–1135.

Neven, D., De Brabandere, B., Proesmans, M., & Van Gool, L. (2018). Towards end-to-end lane detection: An instance segmentation approach. *IEEE Intelligent Vehicles Symposium*, 286–291.

Pan, X., Shi, J., Luo, P., Wang, X., & Tang, X. (2018). Spatial as deep: Spatial CNN for traffic scene understanding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7266–7275.

Qin, Z., Wang, H., & Li, X. (2020). Ultra fast structure-aware deep lane detection. *Proceedings of the European Conference on Computer Vision*, 276–291.

Snell, J., Swersky, K., & Zemel, R. (2017). Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems*, 4077–4087.

Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., ... & Darrell, T. (2020). BDD100K: A diverse driving dataset for heterogeneous multitask learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2636–2645.

Bertozzi, M., & Broggi, A. (1998). GOLD: A parallel real-time stereo vision system for generic obstacle and lane detection. *IEEE Transactions on Image Processing*, 7(1), 62–81.

Chen, L. C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2018). DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4), 834–848.

Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., ... & Schiele, B. (2016). The Cityscapes dataset for semantic urban scene understanding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3213–3223.

Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *Proceedings of the International Conference on Machine Learning*, 1126–1135.

Koch, G., Zemel, R., & Salakhutdinov, R. (2015). Siamese neural networks for one-shot image recognition. *Proceedings of the ICML Deep Learning Workshop*.

McCall, J. C., & Trivedi, M. M. (2006). Video-based lane estimation and tracking for driver assistance: Survey, system, and evaluation. *IEEE Transactions on Intelligent Transportation Systems*, 7(1), 20–37.

Neven, D., De Brabandere, B., Proesmans, M., & Van Gool, L. (2018). Towards end-to-end lane detection: An instance segmentation approach. *IEEE Intelligent Vehicles Symposium*, 286–291.

Pan, X., Shi, J., Luo, P., Wang, X., & Tang, X. (2018). Spatial as deep: Spatial CNN for traffic scene understanding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7266–7275.

Qin, Z., Wang, H., & Li, X. (2020). Ultra fast structure-aware deep lane detection. *Proceedings of the European Conference on Computer Vision*, 276–291.

Snell, J., Swersky, K., & Zemel, R. (2017). Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems*, 4077–4087.

Teichmann, M., Weber, M., Zoellner, M., Cipolla, R., & Urtasun, R. (2018). MultiNet: Real-time joint semantic reasoning for autonomous driving. *IEEE Intelligent Vehicles Symposium*, 1013–1020.

Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., ... & Darrell, T. (2020). BDD100K: A diverse driving dataset for heterogeneous multitask learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2636–2645.

How do I cite this article?
Jay Gupta et.al, Robust Road Boundary Extraction in Unstructured Environments: A Few-Shot Adaptation Approach, Excel International Journal of Technology, Engineering and Management, 2026; Volume -13, Issue-2_Page_104-116. DOI Link:
https://doi.org/10.64971/j.cph.eijtem.v13.i2.14.2026