Evaluating the Impact of Feature Extraction and Clustering Techniques in Highway Guardrail Classification

Luis Flávio Ferreira Monteiro,*, Gabriely Barbosa*, Julio Cezar Soares Silva*,
Thiago Meirelles Ventura† and Raoni Florentino da Silva Teixeira‡
*Nova Rota do Oeste, Mato Grosso, Brazil
Email: {luismonteiro,gabrielybarbosa,juliosilva}@rotadooeste.com.br
†Instituto de Computação, Universidade Federal de Mato Grosso, Mato Grosso, Brazil
Email: thiago.ventura@ufmt.br
‡Faculdade de Ciência e Tecnologia, Universidade Federal de Mato Grosso, Mato Grosso, Brazil
Email: raoni.teixeira@ufmt.br

***Abstract*—This paper presents an approach for the detection of highway guardrails using camera-based systems and advanced machine learning techniques. The proposed methodology combines feature extraction with Convolutional Neural Networks (CNN), specifically MobileNetV2, ResNet18, and VGG16, and clustering algorithms applied to these features. The effectiveness of the models is evaluated through clustering and classification metrics, with a particular emphasis on using the Gaussian Mixture Model (GMM) for forming more cohesive and well-separated clusters compared to K-means. The results indicate that the combination of ResNet18 with GMM provides high accuracy in distinguishing between concrete and metal guardrails, outperforming other tested combinations. This study contributes to the advancement of automatic guardrail detection on highways, providing insights for applications in road asset management.**

***Index Terms*—Clustering, Representation Learning, Object Detection, Road Management**

I. Introduction

Artificial Intelligence (AI) models are developed to solve complex and/or repetitive tasks, with training approaches tailored to the specific problem at hand, typically through supervised or unsupervised learning. Supervised training is particularly critical for object detection, as it requires a well-curated dataset of labeled examples that reflect diverse conditions, including variations in angles, lighting, and other relevant factors [1]. This comprehensive dataset ensures the model can generalize effectively across different scenarios.

Challenges like concept drift can degrade classifier performance over time, and training often requires large volumes of manually labeled data [2]. In contrast, unsupervised learning provides a valuable alternative, as it doesn't rely on labeled data, and is especially useful when labels are hard to obtain. Moreover, unsupervised algorithms like Gaussian Mixture Models (GMM) are capable of addressing concept drift [3].

Guardrails are essential for road safety, acting as protective barriers that prevent vehicles from leaving the roadway and reduce the severity of accidents, particularly on highways, bridges, and areas with steep drop-offs or sharp curves. They come in various forms, such as concrete barriers, which offer high durability and strength for high-speed traffic, and metal guardrails, which provide flexibility to absorb impact and minimize the force transferred to vehicles. The proper classification and monitoring of guardrails are vital for infrastructure management, ensuring that the appropriate type is used and maintained in optimal condition [4].

Advanced Machine Learning (ML) methods for guardrail classification can enhance the accuracy and reliability of these safety measures, leading to better-maintained roads and a reduction in accident-related fatalities and injuries. Clustering features learned by an encoder helps in organizing and understanding complex data representations. This approach not only facilitates the discovery of intrinsic structures and subclasses within the dataset but also enhances the performance of subsequent tasks such as classification or object detection. Effective clustering reduces the need for extensive manual labeling by enabling automated grouping of similar objects [5], [6].

This study proposes a method to evaluate encoders that learn representations of concrete and metal guardrails on highways using a diverse database. We consider a baseline approach using K-means and an alternative approach based on GMM. The GMM approach is particularly effective in detecting more complex patterns under varying environmental conditions, such as different lighting and overlapping objects, thereby ensuring reliable performance in real-world scenarios. This evaluation aims to enhance the robustness and accuracy of guardrail material classification, contributing to improved infrastructure monitoring and safety.

The paper is organized as follows: Section 2 reviews related work. Section 3 details the materials and methods for guardrail detection, including data, encoders, clustering algorithms, and evaluation metrics. Section 4 presents the results and Section 5 concludes with suggestions for future research.

II. Related Work

Recent works combine deep representation learning with clustering methods for image clustering. The methods [7], [8] first learn deep features from input data, and then perform clustering on the learned features in the feature space, in separate two steps. The literature highlights dissimilarities in object embeddings, with Lang et al. [9] exploring hyperbolic embeddings for object detection using keypoint-based and transformer-based methods to enhance classification and localization accuracy.

As Gumaelius cites [4], guardrail detection is becoming increasingly requested, focusing on user navigation safety on highways. Therefore, there are various methods for detecting and mapping these objects using different approaches, such as cameras, radar, and LiDAR. In the study by Zhu et al. [10], a system was developed for guardrail detection using images captured by cameras on intelligent vehicles. The method employed involved modeling the guardrails by dividing them into models for the right and left sides, considering transverse distance, the distance to the right and left guardrails, as well as the slope and curvature. Subsequently, directional filters were applied, and finally, object detection was performed. This system was able to accurately detect guardrails on both sides in curves.

In the study by Kim and Song [11], a hybrid method was used that employed both radar and LiDAR to identify the drivable area through the detection of guardrail and curb reflections. For this, radar sensors and a monocular camera were used, dividing the process into four stages: region of interest selection, where road barriers act as delimiters; radar-based estimation; and the remaining stages of clustering and selection were performed with a series of calculations to obtain these estimates.

As mentioned by Gumaelius [4], data collected with LiDAR are frequently used for object detection and tracking. LiDAR-based approaches are advantageous in as it allows for the acquisition of both the distance of each point and measurement uncertainties. In [12], the authors performed real-time guardrail detection and tracking by creating a database with LiDAR, processing these data with specific regions of interest (ROIs). To detect guardrails, relevant features were extracted from the dataset, followed by clustering, and then temporal similarity and dissimilarity scoring methods were applied to verify whether it was a guardrail and whether it remained the same over time or was replaced by another.

Although LiDAR-based systems excel in obstacle avoidance for autonomous cars, they may not be well-suited for maintenance applications that demand precise recognition of assets composed of various material types, requiring differentiation based on reflectivity and texture. Furthermore, these systems may have difficulty detecting small or partially obscured objects, which are frequently encountered in intricate maintenance environments [13].

The proposed approach aims to use RGB camera images to detect guardrails and extract features to determine their material, which can be concrete or metal. This approach is designed to leverage image processing techniques and ML algorithms to accurately classify the material type. By analyzing visual characteristics such as texture, color, and shape, the system can distinguish between concrete and metal guardrails, enabling more precise assessments for infrastructure monitoring and maintenance.

III. Materials and Methods

A. Proposed approach

Figure 1 shows the system workflow. The images of guardrails were captured by a Ladybug5+ 360 camera mounted on a car, using only images from the lateral views of the highway. The dataset is constituted of 3,110 images. The data labeling was performed manually, ensuring that each image was properly tagged with the corresponding class: metal guardrail (Figure 2), concrete guardrail (Figure 3), or background.

Within this dataset, 92 images of metal guardrails and 55 images of concrete guardrails were clustered ($\approx 5\%$ of the dataset). After clustering guardrails into two different groups, a YOLOV8n [14] model was trained on this new dataset, with 80% of the remaining dataset used for training and 20% for testing.

This study only considered pre-trained encoders to perform feature extraction. With the encoded images, a clustering algorithm was employed to allocate the images into groups corresponding to the number of defined classes. Images were moved to a folder corresponding to their cluster, with labels changed from 0 to concrete guardrails and 1 to metal guardrails. This process facilitated the organization of test images according to the classification performed by the model. A supervised analysis was performed to verify the effectiveness of the generated clusters in labeling the true labels of the images. The experiments were conducted using an Intel Xeon W7-2495X processor, an NVIDIA GeForce RTX A5500 GPU, and 128GB of RAM.

B. Image Encoding Process

Let $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$ represent an input image, where H is the height, W is the width, and C is the number of channels. The encoder is a function f_θ, parameterized by θ, that maps an input image $\mathbf{I}_i$ to a feature vector $\mathbf{z}_i$ in a high-dimensional feature space:

$$\mathbf{z}_i = f_\theta(\mathbf{I}_i), \quad \text{for } i = 1, 2, \ldots, n \tag{1}$$

The set of all encoded features can be represented as:

$$\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_n\} \tag{2}$$

Here, $\mathbf{Z}$ is the collection of feature vectors, each $\mathbf{z}_i \in \mathbb{R}^d$, where d is the dimension of the feature space. The dimensionality d was reduced to 50 using PCA.

The evaluated encoders include VGG16 [15], MobileNetV2 [16], and ResNet18 [17]. For these architectures, the implementations provided by PyTorch [18] were utilized.

C. Clustering the Encoded Images

The clustering algorithm is then applied to $\mathbf{Z}$ to partition the feature space into k clusters. Let $\mathbf{C} = \{C_1, C_2, \ldots, C_k\}$ represent the clusters, where each cluster C_j contains a subset of the feature vectors:

$$C_j = \{\mathbf{z}_i : i \in I_j\}, \quad \text{for } j = 1, 2, \ldots, k \tag{3}$$

Here, I_j is the index set of the data points that belong to cluster C_j. For K-means clustering, the goal is to minimize the within-cluster variance, defined as:

$$\min_{\mathbf{C}} \sum_{j=1}^{k} \sum_{\mathbf{z}_i \in C_j} \|\mathbf{z}_i - \boldsymbol{\mu}_j\|^2 \tag{4}$$

where $\boldsymbol{\mu}_j$ is the centroid of cluster C_j, calculated as:

$$\boldsymbol{\mu}_j = \frac{1}{|C_j|} \sum_{\mathbf{z}_i \in C_j} \mathbf{z}_i \tag{5}$$

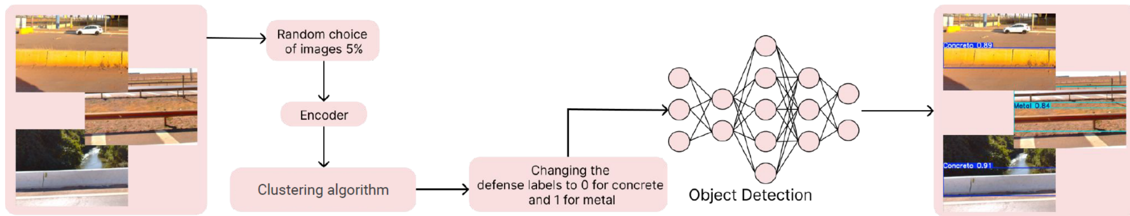

Fig. 1: Overview of the system workflow

Fig. 2: Example of a metal guardrail.

Fig. 3: Example of a concrete guardrail.

For Gaussian Mixture Model (GMM) clustering, the goal is to maximize the likelihood of the data under the model, where each cluster is modeled as a Gaussian distribution:

$$p(\mathbf{z}_i) = \sum_{j=1}^{k} \pi_j \mathcal{N}(\mathbf{z}_i \mid \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) \tag{6}$$

Here, π_j is the mixing coefficient for the j-th Gaussian component, $\boldsymbol{\mu}_j$ is the mean vector, and $\boldsymbol{\Sigma}_j$ is the covariance matrix of the j-th component.

D. *Performance metrics*

The goal of the proposed approach is to assess how the quality of clustering in the encoded guardrail dataset affects the classifier's performance. Thus, three metrics that are widely used in practice to evaluate clustering performance were adopted [19], [20]:

1) Clustering:

- The **Silhouette Index** for a point i is defined as (the higher, the better.):

$$S(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}} \tag{7}$$

  where:
  - $a(i)$ is the average distance from i to points within the same cluster.
  - $b(i)$ is the minimum average distance from i to points in the nearest other cluster.

  The Silhouette Coefficient for a set of samples is given as the mean of the Silhouette Coefficient for each sample.
- The **Calinski-Harabasz Index** is calculated as (the higher the better):

$$CH = \frac{\text{trace}(B_k)}{\text{trace}(W_k)} \times \frac{N - k}{k - 1} \tag{8}$$

  where:
  - B_k is the between-cluster dispersion matrix.
  - W_k is the within-cluster dispersion matrix.
  - N is the total number of points.
  - k is the number of clusters.
- The **Davies-Bouldin Index** is given by (the lower the better):

$$DB = \frac{1}{k} \sum_{i=1}^{k} \max_{j \neq i} \left(\frac{\sigma_i + \sigma_j}{d(\mu_i, \mu_j)} \right) \tag{9}$$

  where:
  - σ_i is the dispersion within cluster i.
  - $d(\mu_i, \mu_j)$ is the distance between the centroids of clusters i and j.

2) Classification: The following metrics for imbalanced datasets was adopted:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{10}$$

where:

- TP (True Positives) are the correctly predicted positive cases.
- FP (False Positives) are the incorrectly predicted positive cases.

$$\text{Recall} = \frac{TP}{TP + FN} \tag{11}$$

where:

- FN (False Negatives) are the actual positive cases that were incorrectly predicted as negative.

By plotting precision against recall at various threshold levels, the Precision-Recall Curve (PRC) provides insight into the trade-offs between these two metrics, allowing for a more nuanced assessment of the classifier's ability to identify positive instances.

IV. Results

A. Clustering Results

The clustering metrics for K-means and GMM across the three architectures are summarized in Tables I.

TABLE I: Clustering Metrics Comparison for K-means and GMM using MobileNetV2, ResNet18, and VGG16

Metric	MobileNetV2	ResNet18	VGG16
K-means			
Silhouette Score	0.0820	0.0846	0.3444
Calinski-Harabasz Score	3.3339	2.8517	0.5157
Davies-Bouldin Score	12.1327	11.1420	2.9319
GMM			
Silhouette Score	0.0568	0.2395	0.3053
Calinski-Harabasz Score	11.9591	47.1855	2.5644
Davies-Bouldin Score	2.6732	1.6534	0.5577

MobileNetV2: The GMM outperforms K-means in Calinski-Harabasz and Davies-Bouldin Scores, indicating better-formed clusters.

ResNet18: GMM consistently surpasses K-means across all metrics, making it the better choice for this architecture.

VGG16: Although K-means performs slightly better in Silhouette Score, GMM shows superior results in Calinski-Harabasz and Davies-Bouldin Scores, suggesting it is preferable for most applications.

B. Classification Results

The classification results for the best and worst algorithm and architecture and clustering combinations are shown in Figure 4. The PRC evaluates classification performance, showing that ResNet18 with GMM provides robust performance, especially for concrete. Conversely, MobileNetV2 with K-means, while performing well for metal, is less effective for concrete. Thus, considering the evaluated dataset the encoder and clustering algorithm combination may influence in the performance of a downstream task, such as classification.

V. Conclusion

This work presented a method for highway guardrail type detection using camera-based systems and advanced ML techniques. The adopted methodology, which involved feature extraction with CNN and clustering these features, proved effective in distinguishing types of guardrails. By utilizing models such as MobileNetV2, ResNet18, and VGG16, it was possible to compare the efficacy of different network architectures in clustering tasks and evaluate the quality of the results obtained.

The results showed that the combination of feature extraction with the clustering algorithm can significantly influence detection accuracy. Among the evaluated models, VGG16 learned features that, when clustered with GMM, achieved the best performance in terms of Silhouette and Davies-Bouldin indices, suggesting better cluster separation and cohesion. Thus, GMM demonstrated superior capability in modeling clusters with more complex structures, which may be particularly useful for future detection improvements.

Future work will explore the integration of additional techniques, such as the use of anchor boxes

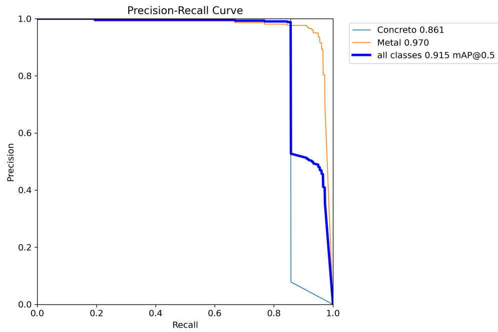

(a) Best Case: ResNet18 with GMM

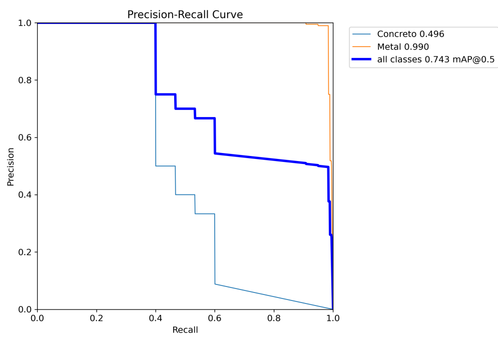

(b) Worst Case: MobileNetV2 with K-means

Fig. 4: PRC for the best and worst architecture and clustering combinations

and semantic segmentation, to further enhance detection accuracy. Moreover, applying image generation models and masked image modeling techniques could offer new advancements, as discussed by Rakowski et al. [21] and Xie et al. [22]. The incorporation of damage detection techniques and the use of more sophisticated models, such as Mask R-CNN [23], also hold promise for improving the ability to identify and classify different types of guardrails with greater precision.

Acknowledgment

The authors would like to thank the National Land Transport Agency (ANTT) for the resources provided under Resolution No. 483, of March 24, 2004, as well as for the infrastructure and coordination provided by Nova Rota do Oeste. The authors would also like to thank the Research Support Foundation of Mato Grosso – FAPEMAT (process FAPEMAT-PRO-2022/01047) for their support during the development of this work.

References

[1] J. Wang, T. Zhang, Y. Cheng, and N. Al-Nabhan, "Deep learning for object detection: A survey," *Computer Systems Science and Engineering*, vol. 38, no. 2, p. 165–182, 2021.

[2] Y. Pu, J. Sun, N. Tang, and Z. Xu, "Deep expectation-maximization network for unsupervised image segmentation and clustering," *Image and Vision Computing*, vol. 135, p. 104717, 2023.

[3] G. H. F. M. Oliveira, L. L. Minku, and A. L. I. Oliveira, "Tackling virtual and real concept drifts: An adaptive gaussian mixture model approach," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 2, pp. 2048–2060, 2023.

[4] N. Gumaelius, "Guardrail detection for landmark-based localization," 2022.

[5] A. Saxena, M. Prasad, A. Gupta, N. Bharill, O. P. Patel, A. Tiwari, M. J. Er, W. Ding, and C.-T. Lin, "A review of clustering techniques and developments," *Neurocomputing*, vol. 267, p. 664–681, Dec. 2017.

[6] H. Huang, C. Wang, X. Wei, and Y. Zhou, "Deep image clustering: A survey," *Neurocomputing*, vol. 599, p. 128101, 2024.

[7] F. Tian, B. Gao, Q. Cui, E. Chen, and T.-Y. Liu, "Learning deep representations for graph clustering," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 28, no. 1, Jun. 2014.

[8] A. Coates and A. Y. Ng, *Learning Feature Representations with K-Means*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 561–580.

[9] C. Lang, A. Braun, L. Schillingmann, and A. Valada, *On Hyperbolic Embeddings in Object Detection*. Springer International Publishing, 2022, p. 462–476.

[10] H. Zhu and B. Guo, "A beam guardrail detection algorithm using lidar for intelligent vehicle," in *2018 IEEE 8th Annual International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*, 2018, pp. 1398–1402.

[11] T. Kim and B. Song, "Detection and tracking of road barrier based on radar and vision sensor fusion," *Journal of Sensors*, vol. 2016, p. 1–8, 2016.

[12] H. Zhu, F. Chen, and B. Guo, "Joint beam guardrail detection and tracking by lidar for real-time applications," in *2019 Chinese Automation Congress (CAC)*. IEEE, Nov. 2019.

[13] A. Elamin and A. El-Rabbany, "Uav-based image and lidar fusion for pavement crack segmentation," *Sensors*, vol. 23, no. 23, 2023.

[14] G. Jocher, A. Chaurasia, J. Qiu, C. Fang *et al.*, "Yolov8n: A light-weight object detection model," *Ultralytics Documentation*, 2023, available at https://github.com/ultralytics/ultralytics.

[15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2015.

[16] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," 2019.

[17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," 2015.

[18] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," 2019.

[19] K. Gonzalez and S. Misra, "Unsupervised learning monitors the carbon-dioxide plume in the subsurface carbon storage reservoir," *Expert Systems with Applications*, vol. 201, p. 117216, 2022.

[20] M. E. Ferrão, P. Prata, and P. Fazendeiro, "Utility-driven assessment of anonymized data via clustering," *Scientific Data*, vol. 9, no. 1, Jul. 2022.

[21] C. C. Rakowski and T. Bourlai, "On enhancing crack semantic segmentation using stylegan and brownian bridge diffusion," *IEEE Access*, vol. 12, pp. 34 769–34 784, 2024.

[22] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu, "Simmim: A simple framework for masked image modeling," 2022.

[23] H. Jin, Z. Lan, and X. He, "On highway guardrail detection algorithm based on mask rcnn in complex environments," in *2021 7th International Conference on Systems and Informatics (ICSAI)*, 2021, pp. 1–6.