Railway Sinkhole Detection: An Integrated Approach Combining Geomorphological Segmentation and Descriptor-Based Machine Learning

Maryem Bouali*†, Fakhreddine Ababsa†, Rani El Mehouche*,
Muhammad Ali Sammuneh*, Bahar Salavati‡, Flavien Viguier‡
* Institut de Recherche IR, Ecole Spéciale des Travaux Publics (ESTP),
Cachan, France
Email: mbouali@estp.fr, relmeouche@estp.fr, msammuneh@estp.fr
† PIMM, Arts et Métiers ParisTech, CNRS, CNAM, HESAM University,
Paris, France
Email: fakhreddine.ababsa@ensam.eu
‡ Directions Techniques Réseau - DGII DTR IP3M DM Matrice, SNCF Réseau,
Saint Denis, France
Email: bahar.salavati@reseau.sncf.fr, flavien.viguier@reseau.sncf.fr

***Abstract*—Railway sinkholes represent a critical geotechnical hazard for rail infrastructure safety and require reliable detection methods based on high-resolution topographic data. This paper proposes an integrated methodology combining geomorphological analysis and machine learning for the automated detection of railway sinkholes from Light Detection and Ranging (LiDAR)-derived Digital Elevation Models (DEMs). The proposed framework relies on a two-stage approach. First, a geomorphological processing chain is applied to generate high-resolution DEMs and delineate topographic depressions using watershed segmentation. This step enables the systematic extraction of candidate basins potentially corresponding to sinkholes. Second, each extracted basin is characterized using a set of handcrafted geometric and morphometric descriptors capturing properties, such as shape, depth, volume, and spatial structure. These features are then used within a supervised learning framework to discriminate true sinkholes from other terrain irregularities. Several classification algorithms, are implemented and compared to assess their performance and robustness. Experimental results demonstrate that integrating feature-based machine learning with geomorphological segmentation significantly improves detection accuracy and reduces false positives compared to a purely morphological approach. The proposed workflow contributes to spatial terrain analysis and infrastructure monitoring by providing a scalable and data-driven methodology for railway risk assessment based solely on elevation data.**

***Keywords-Railway Sinkholes*; *LiDAR Point Clouds*; *Digital Elevation Models (DEM)*; *Geomorphological Analysis*; *Supervised Learning*.**

I. INTRODUCTION

A railway sinkhole is a localized ground collapse caused by the upward propagation of a void roof failure until rupture reaches the surface. It typically originates from the collapse of shallow underground cavities of natural origin (karst dissolution, gypsum dissolution, internal erosion) or anthropogenic origin (abandoned quarries, underground galleries, buried infrastructures). The process begins with cavity roof failure, followed by the progressive formation of a collapse chimney that migrates upward (Figure 1). When the overburden becomes too thin to sustain arching effects, a sudden surface rupture occurs, forming a circular or funnel-shaped crater. While early stages may develop slowly, the final collapse is often abrupt and difficult to predict. Precursory signs are frequently subtle or absent [1], making sinkholes inherently hazardous.

Railway infrastructures are particularly vulnerable to such instabilities. Due to strict geometric tolerances required for safe operation, even small vertical displacements—on the order of a few centimeters—may lead to speed restrictions or traffic [2]. Several recent incidents in France have resulted in emergency traffic suspensions and costly geotechnical interventions. Beyond safety concerns, these events generate significant operational and economic impacts.

These constraints highlight the need for reliable methods capable of identifying surface depressions potentially associated with sinkholes. Given the morphological variability of sinkhole manifestations and the highly structured nature of railway environments, detection approaches must distinguish true geotechnical instabilities from benign topographic irregularities. Developing robust, scalable detection strategies is therefore essential for improving risk prioritization in railway monitoring systems.

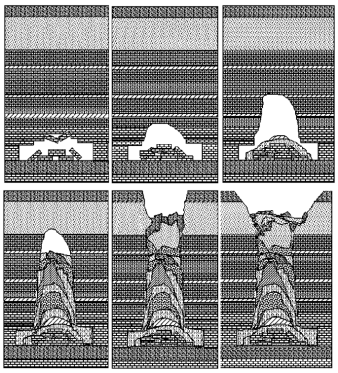

Figure 1. Stages of sinkhole formation [3].

The remainder of this manuscript is structured as follows. Section II reviews existing sinkhole detection approaches in railway and geotechnical contexts. Section III presents the

proposed two-stage framework combining geomorphological segmentation and supervised classification. Section IV reports experimental results on both training and independent test datasets. Finally, Section V discusses limitations and perspectives for large-scale deployment.

II. Related Work

The detection of railway sinkholes is challenging due to their localized nature and their potentially rapid impact on track stability. Detection-oriented approaches aim to identify already formed sinkholes or advanced degradation signs from surface or near-surface observations. Existing methods can be broadly categorized into four families.

Track geometry monitoring relies on measurement trains that record longitudinal level, alignment, cross-level, and gauge. Although primarily developed for maintenance purposes, these indicators have been investigated as indirect proxies for substructure degradation [4]. Time-series analysis can reveal localized settlements possibly linked to sinkhole formation, but such approaches remain indirect, provide no explicit information about cavity geometry, and may lack sensitivity when structural stiffness compensates small-scale instabilities [5].

Surface morphological analysis has expanded with the availability of airborne Light Detection and Ranging (LiDAR) and Unmanned Aerial Vehicle (UAV) photogrammetry, enabling automated sinkhole detection from high-resolution DEMs. Typical methods identify closed depressions using morphometric indices, contour-tree structures, or statistical susceptibility mapping [6], [7]. While effective for surface-visible sinkholes, DEM-based approaches are sensitive to anthropogenic structures and topographic artifacts. In railway environments, engineered platform geometries complicate the discrimination between genuine sinkholes and benign features, limiting purely morphometric thresholding strategies.

Ground Penetrating Radar (GPR) detects near-surface anomalies by analyzing reflected electromagnetic waves [8]. Railway-mounted systems provide decimetric-resolution profiles of the track substructure, and machine learning techniques have been proposed for automated anomaly detection in radargrams [9]. However, interpretation depends strongly on soil conditions and acquisition parameters, and GPR requires dedicated inspection campaigns, which limits continuous large-scale monitoring.

Satellite Radar Interferometry (InSAR) estimates ground deformation from multi-temporal radar acquisitions. Persistent Scatterer and SBAS approaches enable millimetric subsidence monitoring along railway corridors [10], and have identified progressive deformation associated with sinkhole activity in karst or mining regions [11]. Nevertheless, InSAR measures surface deformation rather than cavities directly and is less sensitive to small-scale or abrupt collapses without prior gradual settlement.

In this context, our work focuses on surface morphological analysis from high-resolution LiDAR-derived DEMs. Instead of relying solely on morphometric thresholding, we propose a two-stage framework combining watershed-based geomorphological segmentation with supervised classification, aiming to preserve morphological sensitivity while improving statistical discrimination in structured railway environments.

III. Methodology

The proposed methodology addresses the limitations of purely geomorphological segmentation when applied to engineered railway DEMs. In previous work, a marker-controlled watershed framework was introduced to delineate depressions from LiDAR-derived DEMs, successfully extracting sinkhole-like basins [12]. However, in railway environments, this approach generates numerous false positives due to ballast roughness, drainage features, and other anthropogenic micro-relief patterns. To overcome this limitation, a two-stage framework is adopted: (i) watershed segmentation is retained as a high-recall candidate generator, and (ii) the extracted basins are reformulated as a supervised classification problem using a compact set of handcrafted geometric and morphometric descriptors.

Stage I: Geomorphological Watershed Segmentation

The first stage applies a marker-controlled watershed segmentation to LiDAR-derived DEMs to extract candidate topographic depressions (Figure 2). The complete framework was introduced in our previous work [12]; only the main principles are summarized here.

The DEM is first denoised using an adaptive Wiener filter to reduce local noise while preserving terrain morphology [13]. The watershed transform interprets the DEM as a topographic surface and partitions it into catchment basins associated with local minima [14]–[16]. In our implementation, segmentation is applied to the gradient magnitude of the filtered DEM, so that low-gradient areas form basin interiors while high-gradient zones define separating ridges.

To limit over-segmentation, a marker-based formulation is adopted. Morphologically significant minima are selected using an h-minima operator, while border-connected components are discarded. The parameter h is chosen conservatively to preserve all potential sinkholes, even at the cost of generating additional non-relevant basins. The resulting labeled depressions constitute high-recall candidates that are subsequently refined through supervised classification.

Stage II: Basin-Level Feature Extraction and Supervised Classification

Each watershed basin is represented by a compact set of geometric and morphometric descriptors computed at the region level. These features (Table I) characterize basin size, shape, depth, and internal topographic variability.

Because these descriptors span different numerical ranges, feature scaling is applied when required. Standard normalization is used for scale-sensitive classifiers (e.g., logistic regression and Support Vector Machine (SVM)), while tree-based models are trained directly on raw features.

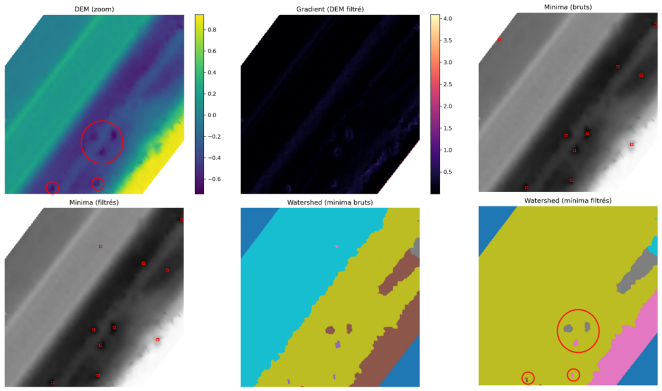

Figure 2. Overview of the non-learning watershed segmentation pipeline. From DEM and gradient computation to raw and filtered minima, and final watershed segmentation. Highlighted regions represent candidate basins used for feature extraction and supervised classification.

TABLE I. BASIN-LEVEL DESCRIPTORS USED FOR SUPERVISED CLASSIFICATION

Feature	**Definition / Meaning**
Area	Number of pixels in the basin (surface extent)
Perimeter	Length of basin boundary
Circularity	$C = \frac{4\pi \cdot \text{area}}{\text{perimeter}^2}$
Elongation	$E = \frac{\text{major axis}}{\text{minor axis}}$
Local depth	$d_{\text{loc}} = \text{Gauss}(DEM) - DEM$
Maximum depth	$\max(d_{\text{loc}})$ over basin
Mean depth	Mean of d_{loc}
Mean slope	Mean terrain slope over basin
Mean gradient	Mean gradient magnitude
Mean roughness	Mean local DEM standard deviation (5×5 window)

Stage III: Supervised Basin Classification

Following feature extraction, each candidate basin is represented by a descriptor vector $\mathbf{f}_k$. The objective of Stage III is to learn a binary decision function $\mathcal{C} : \mathbf{f}_k \rightarrow \{0, 1\}$ that discriminates true sinkholes from non-relevant terrain depressions. Given the moderate size of the dataset and the structured nature of the engineered features, we evaluate four complementary classification models, ranging from linear baselines to ensemble-based non-linear methods.

a) Logistic Regression: Logistic regression serves as a linear probabilistic baseline [17]. It models the posterior probability $P(y = 1 \mid \mathbf{f}) = \sigma(w^\top \mathbf{f} + b)$, yielding a linear decision boundary and directly interpretable feature weights.

b) Support Vector Machine: Support Vector Machines (SVM) aim to maximize the margin between classes [18]. By employing a Radial Basis Function (RBF) kernel, the model captures non-linear decision boundaries while maintaining strong generalization capabilities, particularly suited to small-to-medium datasets.

c) Random Forest: Random Forest is an ensemble learning method composed of decorrelated decision trees trained via bootstrap aggregation and randomized feature selection [19]. The final prediction is obtained through majority voting, allowing the model to capture complex non-linear feature interactions while reducing overfitting.

d) LightGBM: LightGBM is a Light Gradient Boosting Machine (LightGBM) framework based on sequentially constructed decision trees [20]. Each tree is optimized to minimize the residual errors of previous iterations, enabling efficient learning of high-order feature interactions in structured tabular data.

Experimental Setup: The framework is evaluated on 20 LiDAR-derived DEMs containing confirmed railway sinkholes. For each DEM, marker-controlled watershed segmentation (Stage I) extracts candidate basins, followed by post-segmentation filtering to retain only morphologically valid regions suitable for descriptor computation. This yields 60 labeled basins: 20 true sinkholes and 40 non-sinkhole regions, defined based on spatial overlap with ground-truth masks.

Supervised learning is performed at the basin level using stratified cross-validation to preserve class proportions. Logistic Regression, SVM (RBF kernel), Random Forest, and LightGBM are trained on the training folds and evaluated on the corresponding validation folds. Performance is primarily assessed using Precision–Recall (PR) metrics to account for class imbalance.

IV. Results and Discussion

This section evaluates the proposed watershed-based basin classification framework through a two-stage analysis. First, classifier performance is assessed during training using 5-fold stratified cross-validation on the watershed basins extracted from the training DEMs. Second, the trained models are evaluated on an independent dataset composed exclusively of real railway DEMs in order to assess their generalization capability and robustness under realistic operational conditions characterized by class imbalance, geomorphological variability, and segmentation artifacts.

A. Comparative Analysis of Supervised Classifiers on Training Data

This section evaluates the performance of the supervised classifiers trained on a dataset composed of $N = 60$ watershed basins, including 20 sinkhole basins (positive class) and 40 non-sinkhole basins (negative class). The prevalence of the positive class is therefore $\pi = 20/60 = 0.333$.

In this imbalanced setting, accuracy alone may conceal unfavorable classification behavior, particularly over-prediction of the majority class. Performance is therefore analyzed using Precision–Recall (PR) curves, which are more informative when the primary objective is to detect a relatively rare class. The Area Under the Precision–Recall Curve (AUPRC) is used as the main ranking metric. It is important to note that, in PR space, the expected performance of a random classifier equals

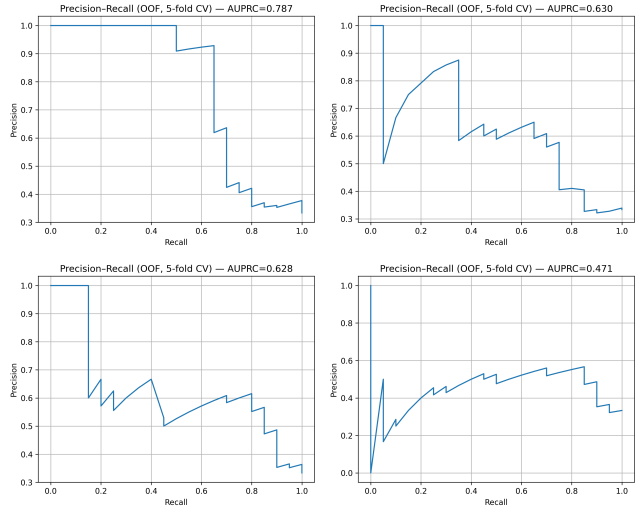

Figure 3. Precision–recall curves (5-fold cross-validation). (a) RF: 0.79, (b) LGBM: 0.63, (c) SVM: 0.63, (d) LR: 0.47.

the class prevalence; therefore, a non-discriminative method would yield an AUPRC close to 0.333.

The PR curves shown in Figure 3 indicate that all models outperform the random baseline (AUPRC $>$ 0.333), confirming the existence of a discriminative signal within the morphometric basin descriptors. However, performance differences are substantial.

Random Forest achieves the highest AUPRC (0.787), more than twice the random reference. This result indicates a strong ability to rank basins according to their probability of being sinkholes. In particular, the top-ranked basins exhibit high precision, which is critical in a detection pipeline aiming to reduce false positives prior to inspection or post-processing. The curve shape suggests a high-precision / moderate-recall regime, enabling the adoption of a strict decision threshold to retain only highly reliable detections.

LightGBM and SVM obtain comparable AUPRC values (0.630 and 0.628, respectively), substantially above the random baseline (absolute gain of approximately +0.30 over 0.333), yet clearly inferior to Random Forest. Their curves appear more irregular, reflecting increased sensitivity to the limited dataset size (only 20 positive samples) and to fold variability. These models remain capable of prioritizing sinkholes, but precision decreases more rapidly as recall increases, leading to higher false-positive rates when targeting high recall.

Logistic regression reaches an AUPRC of 0.471, above the random reference (0.333) but significantly lower than the non-linear models. This suggests that linear separation in feature space is insufficient to capture the diversity of basin geometries produced by watershed segmentation. This result is consistent with the problem structure: discriminative morphological signatures likely depend on non-linear interactions between descriptors (e.g., area, circularity, perimeter, and depth).

Aggregated 5-fold cross-validation metrics are reported in Table II. Results are presented as mean $\pm$ standard deviation to quantify variability induced by data partitioning, which is particularly relevant given the limited number of positive samples.

TABLE II. MEAN PERFORMANCE (± STD) FROM 5-FOLD STRATIFIED CROSS-VALIDATION

Model	Acc.	Prec.	Rec.	F1	AUPRC
RF	0.83±0.12	0.83±0.21	0.65±0.20	0.72±0.18	0.79±0.12
LGBM	0.78±0.10	0.60±0.15	0.62±0.18	0.61±0.14	0.63±0.13
SVM	0.76±0.11	0.58±0.17	0.60±0.20	0.59±0.16	0.63±0.15
LR	0.70±0.10	0.54±0.15	0.65±0.26	0.58±0.17	0.47±0.15

The values reported in Table II confirm and complement the PR analysis. Random Forest provides the strongest overall performance, with a mean F1-score of 0.72 ± 0.18 and high precision (0.83 ± 0.21), indicating a favorable balance between recall and false-positive control under moderate class prevalence ($\pi = 0.333$).

LightGBM and SVM exhibit intermediate performance, with F1-scores close to 0.60 and increased variability across folds, reflecting greater sensitivity to the small number of positive samples. Logistic regression, although outperforming the random baseline, remains consistently weaker across metrics, confirming the limitations of a purely linear decision boundary for this task.

Based on these training results, Random Forest is selected as the primary model for subsequent evaluation on an independent test set composed exclusively of real railway DEMs.

B. Comparative Analysis of Supervised Classifiers on the Test Phase

The evaluation phase is conducted on an independent test dataset composed of real DEMs not used during training. Twelve DEMs, each containing one confirmed sinkhole, are processed using the same marker-controlled watershed segmentation pipeline described previously. The segmentation of these 12 DEMs produces a total of 72 valid basins, including 12 true sinkhole basins and 60 basins corresponding to non-relevant morphological structures.

The number of basins generated by the watershed is not constant and strongly depends on the local morphology of each DEM, particularly surface roughness, the presence of topographic irregularities, and the level of over-segmentation induced by the method. Consequently, some DEMs generate a larger number of candidate basins than others, independently of the actual number of sinkholes present. This pronounced imbalance reflects a realistic operational scenario in which sinkholes represent a rare class among numerous candidate depressions. The performance reported in this section therefore evaluates the ability of the models to generalize and discriminate sinkholes under practical conditions, without any post-hoc adjustment of the watershed outputs.

The corresponding confusion matrices are presented in Table III, and the performance metrics associated with the *sinkhole* class (positive class) are summarized in Table IV.

The analysis of the confusion matrices (Table III) reveals several important trends. Logistic regression, although correctly

TABLE III. CONFUSION MATRICES FOR THE FOUR MODELS (TEST SET: 72 BASINS)

Logistic Regression

	Pred. 0	Pred. 1
True 0	37	23
True 1	6	6

SVM (RBF)

	Pred. 0	Pred. 1
True 0	42	18
True 1	3	9

Random Forest

	Pred. 0	Pred. 1
True 0	55	5
True 1	3	9

LightGBM

	Pred. 0	Pred. 1
True 0	51	9
True 1	4	8

Table IV
CLASSIFICATION SCORES FOR THE *sinkhole* CLASS

Model	Accuracy	Precision	Recall	F1-score
Logistic Regression	0.597	0.207	0.500	0.293
SVM (RBF)	0.708	0.333	0.750	0.462
LightGBM	0.819	0.471	0.667	0.552
Random Forest	0.889	0.643	0.750	0.692

identifying half of the sinkholes (6/12), produces a very large number of false positives (23), confirming insufficient linear separability between basin classes. The RBF SVM substantially reduces false positives (18) while detecting the majority of sinkholes (9/12), benefiting from its ability to model non-linear decision boundaries.

Tree-based methods (Random Forest and LightGBM) exhibit the most balanced confusion matrices, with few false negatives (3–4) and a limited number of false positives (5–9), indicating improved robustness to local DEM variations. Random Forest achieves the clearest separation, with only 5 false positives and 3 false negatives.

The metrics reported in Table IV confirm these observations. Logistic regression yields moderate recall (0.500) but extremely low precision (0.207), resulting in a limited F1-score (0.293). The SVM significantly improves performance, particularly in recall (0.750), but still suffers from numerous false alarms. LightGBM provides intermediate performance, balancing sensitivity and error reduction. Random Forest emerges as the most reliable and balanced model, achieving the highest F1-score (0.692) and the best trade-off between sinkhole detection and false-positive control. Its stability can be attributed to the aggregation of weakly correlated trees, which smooths geomorphological noise while effectively capturing non-linear relationships between basin descriptors.

Comparing cross-validation results with independent test performance reveals overall consistency. For Random Forest, LightGBM, and SVM, test scores fall within the variability range defined by the mean ± standard deviation observed during cross-validation, indicating satisfactory generalization capacity. Logistic regression constitutes a notable exception, with a marked performance drop on the test set, highlighting the intrinsic limitations of a linear model in a strongly imbalanced and non-linear problem setting.

Figure 4 illustrates classification results on a representative test DEM, comparing the initial watershed segmentation with predictions produced by the four evaluated models.

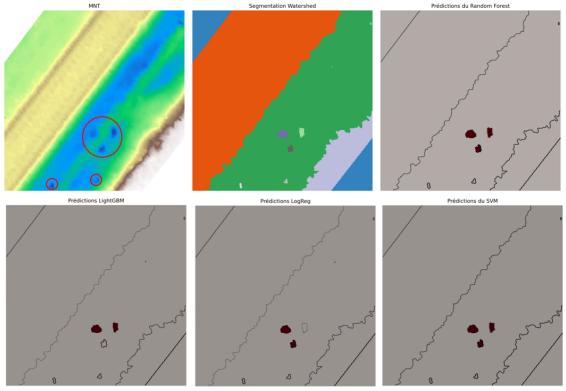

Figure 4. Comparative visualization of final predictions for the four models: original DEM (top left) with true sinkholes outlined in red, watershed segmentation (top center), Random Forest predictions (top right), LightGBM (bottom left), Logistic Regression (bottom center), and SVM (bottom right). Basins predicted as sinkholes are shown in brown.

Performance Analysis of the Random Forest Model: Figure 5(a) presents feature importance computed using the Gini impurity criterion within the Random Forest. This metric reflects how frequently each descriptor contributes to node splitting across the ensemble and therefore highlights structurally dominant variables in the learned decision process. The descriptors *circularity* and *area* emerge as the most influential, indicating that global shape properties define the primary decision boundaries in the basin feature space. Additional variables such as *depth_mean*, *elongation*, *perimeter*, and *slope_mean* also contribute meaningfully to tree construction.

A more predictive evaluation is provided by permutation importance (Figure 5(b)), which measures performance degradation when each variable is independently shuffled. Unlike the Gini-based ranking, this analysis reveals that *area* is by far the most decisive descriptor for effective classification, followed by *circularity* and *perimeter*. Secondary geometric variables (*rough_mean*, *grad_mean*, *slope_mean*, *elongation*) exhibit minimal influence on predictive performance, suggesting redundancy or weak discriminative contribution. The negligible role of *depth_max* confirms its limited relevance for the detection task.

The overall discriminative capacity of the model is illustrated by the Receiver Operating Characteristic (ROC) curve in Figure 5(c). The Area Under the Curve (AUC = 0.799) indicates good separation between sinkhole basins and non-relevant depressions. The curve lies clearly above the random baseline, demonstrating that the model captures meaningful structure despite the limited dataset size and inherent variability. Decision threshold optimization yields an optimal threshold of 0.763, maximizing the F1-score at 0.700. The distribution of predicted probabilities (Figure 5(d)) further supports this

interpretation. Most basins receive low probabilities (< 0.4), reflecting a conservative prediction regime, while true sinkholes predominantly cluster in the high-probability region (> 0.75), naturally justifying the selected threshold. The intermediate interval (0.4–0.7) corresponds to geometrically ambiguous basins where morphological cues are less distinctive. This structured distribution explains why the optimal operating point lies well above the default threshold of 0.5.

Overall, the Random Forest effectively extracts the morphological patterns associated with railway sinkholes while maintaining robustness to geometric heterogeneity and terrain variability.

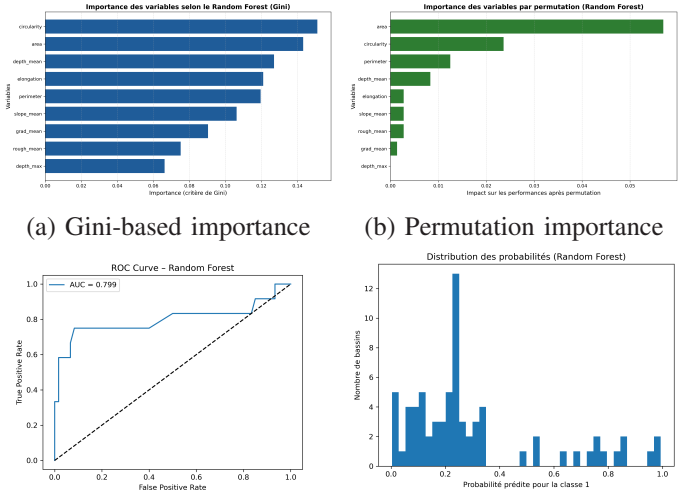

(a) Gini-based importance (b) Permutation importance

(c) ROC curve (AUC = 0.799) (d) Probability distribution

Figure 5. Random Forest performance analysis.

V. Conclusion And Future Work

This study proposed and evaluated a structured two-stage framework for the detection of railway sinkholes from LiDAR-derived Digital Elevation Models. The methodology combines (i) a geomorphological watershed segmentation to extract candidate depressions and (ii) supervised classification of basin-level morphometric descriptors to distinguish true sinkholes from non-relevant surface irregularities.

Experimental results demonstrate that the watershed segmentation provides a consistent and physically grounded partition of the terrain, enabling systematic candidate extraction while preserving all relevant depressions. When coupled with supervised learning, this representation proves effective for discriminating sinkholes within a challenging and imbalanced setting. Among the evaluated classifiers, tree-based methods — particularly Random Forest — achieve the best overall performance, reaching an F1-score of 0.692 on an independent test set composed exclusively of real railway DEMs. Feature importance analyses consistently emphasize the dominant role of geometric descriptors such as area, perimeter, and circularity, confirming that basin shape constitutes a robust and interpretable signature of railway sinkholes in the considered dataset.

Overall, the proposed framework establishes a reproducible and operational baseline for automated railway sinkhole detection based on geomorphological reasoning and supervised learning. While the approach relies on predefined morphometric descriptors and watershed-based candidate generation, it provides a transparent and explainable decision process that is well adapted to structured terrain analysis.

Future work may explore extensions toward richer spatial representations, including learning-based approaches operating directly on three-dimensional data or point clouds. Such developments could further enhance robustness and generalization across heterogeneous railway environments, while building upon the geomorphological insights established in this study.

References

[1] J. N. van der Merwe, “The development of a time-based probabilistic sinkhole prediction method for coal mining in the witbank and highveld coalfields”, *Journal of the Southern African Institute of Mining and Metallurgy*, vol. 120, no. 6, pp. 393–402, 2020.

[2] B. Indraratna, S. Nimbalkar, and D. Christie, “The performance of rail track incorporating the effects of ballast fouling and variable axle loads”, *Journal of Geotechnical and Geoenvironmental Engineering*, vol. 136, no. 10, pp. 1447–1457, 2010.

[3] C. Caudron, “Caractérisation de l’aléa fontis lié aux carrières souterraines abandonnées”, Ph.D. dissertation, Institut National Polytechnique de Lorraine, Nancy, France, 2007.

[4] C. Esveld, *Modern Railway Track*. MRT-Productions, 2001.

[5] C. Vale and S. M. Lurdes, “Stochastic model for the geometrical rail track degradation process in the portuguese railway northern line”, *Reliability Engineering & System Safety*, vol. 116, pp. 91–98, 2013.

[6] Q. Wu, H. Deng, and Y. Liu, “Automated sinkhole extraction from topographic data using contour trees”, *Geomorphology*, vol. 253, pp. 273–286, 2016.

[7] M. Bouali, F. Ababsa, M. A. Sammuneh, R. E. Meouche, B. Salavati, and F. Viguier, “Enhanced railway sinkhole detection and characterization using lidar-based 3d modeling and geometric analysis”, *Pattern Recognition Letters*, vol. 197, pp. 378–384, 2025. DOI: 10.1016/j.patrec.2025.06.015.

[8] A. M. Alani, F. Tosti, and A. Benedetto, “Applications of ground penetrating radar (gpr) in bridge deck monitoring and assessment”, *Journal of Applied Geophysics*, vol. 97, pp. 45–54, 2013.

[9] I. Giannakis, A. Giannopoulos, and C. Warren, “A machine learning approach for buried object classification using ground penetrating radar”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 7, pp. 3943–3953, 2016.

[10] A. Ferretti, C. Prati, and F. Rocca, “Permanent scatterers in sar interferometry”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 1, pp. 8–20, 2001.

[11] J. P. Galve, F. Gutiérrez, J. Guerrero, and G. Bonachea, “Sinkholes in the ebro valley evaporite karst detected from dinsar”, *Geomorphology*, vol. 229, pp. 1–15, 2015. DOI: 10.1016/j.geomorph.2014.02.033.

[12] M. Bouali, M. A. Sammuneh, R. E. Meouche, F. Ababsa, B. Salavati, and F. Viguier, “An automated technique for detecting sinkholes from lidar dems in railway environments”, in *Intelligent Systems and Pattern Recognition*, Springer, 2025, pp. 327–343. DOI: 10.1007/978-3-031-82412-2_22.

[13] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. MIT Press, 1949.

[14] L. Vincent and P. Soille, “Watersheds in digital spaces: An efficient algorithm based on immersion simulations”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 6, pp. 583–598, 1991. DOI: 10.1109/34.87344.

[15] S. Beucher and F. Meyer, “The morphological approach to segmentation: The watershed transformation”, in *Mathematical Morphology in Image Processing*, E. R. Dougherty, Ed., Marcel Dekker, 1993, pp. 433–481.

[16] P. Soille, *Morphological Image Analysis: Principles and Applications*, 2nd ed. Springer, 2003.

[17] D. W. Hosmer, S. Lemeshow, and R. X. Sturdivant, *Applied Logistic Regression*, 3rd ed. Wiley, 2013.

[18] C. Cortes and V. Vapnik, “Support-vector networks”, *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[19] L. Breiman, “Random forests”, *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[20] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, “Lightgbm: A highly efficient gradient boosting decision tree”, in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 3146–3154.