GEOProcessing 2026

The Eighteenth International Conference on Advanced Geographic Information Systems, Applications, and Services

ISBN: 978-1-68558-384-2

May 24 - 28, 2026

Venice, Italy

GEOProcessing 2026 Editors

Alexey Cheptsov, Information Service Center of the University of Stuttgart (TIK), Germany

GEOProcessing 2026

Forward

The Eighteenth International Conference on Advanced Geographic Information Systems, Applications, and Services (GEOProcessing 2026), held between May 24, 2026, and May 28, 2026, in Venice, Italy, continued a series of events addressing various aspects of managing geographical information and web services.

The goal of the GEOProcessing 2026 conference was to bring together researchers from the academia and practitioners from the industry in order to address fundamentals of advances in geographic information systems and the new applications related to them using Web Services. Such systems can be used for assessment, modeling, and prognosis of emergencies.

We take here the opportunity to warmly thank all the members of the GEOProcessing 2026 technical program committee, as well as all the reviewers. The creation of such a high-quality conference program would not have been possible without their involvement. We also kindly thank the authors who dedicated time and effort to contribute to GEOProcessing 2026. We truly believe that, thanks to all these efforts, the final conference program consisted of top-quality contributions. We also thank the members of the GEOProcessing 2026 organizing committee for their help in handling the logistics of this event.

We hope that GEOProcessing 2026 was a successful international forum for the exchange of ideas and results between academia and industry for the promotion of progress in the field of geographical information systems, applications, and services.

GEOProcessing 2026 Chairs

GEOProcessing 2026 Steering Committee Chair
Claus-Peter Rückemann, Universität Münster / DIMF / Leibniz Universität Hannover, Germany

GEOProcessing 2026 Steering Committee
Thomas Ritz, FH Aachen, Germany
Yerach Doytsher, Technion - Israel Institute of Technology, Haifa, Israel
Roger Tilley, Sandia National Laboratories, USA
Alexey Cheptsov, Information Service Center of the University of Stuttgart (TIK), Germany
Jui-Hsin (Larry) Lai, Ping-An Technology - Research Lab, USA

GEOProcessing 2026 Publicity Chairs
José Miguel Jiménez, Universitat Politecnica de Valencia, Spain
Francisco Javier Díaz Blasco, Universitat Politècnica de València, Spain
Ali Ahmad, Universitat Politècnica de València, Spain
Sandra Viciano Tudela, Universitat Politecnica de Valencia, Spain
Laura Garcia, Universidad Politécnica de Cartagena, Spain

GEOProcessing 2026
Committee

GEOProcessing 2026 Steering Committee Chair

Claus-Peter Rückemann, Universität Münster / DIMF / Leibniz Universität Hannover, Germany

GEOProcessing 2026 Steering Committee

Thomas Ritz, FH Aachen, Germany
Yerach Doytsher, Technion - Israel Institute of Technology, Haifa, Israel
Roger Tilley, Sandia National Laboratories, USA
Alexey Cheptsov, Information Service Center of the University of Stuttgart (TIK), Germany
Jui-Hsin (Larry) Lai, Ping-An Technology - Research Lab, USA

GEOProcessing 2026 Publicity Chairs

José Miguel Jiménez, Universitat Politecnica de Valencia, Spain
Francisco Javier Díaz Blasco, Universitat Politècnica de València, Spain
Ali Ahmad, Universitat Politècnica de València, Spain
Sandra Viciano Tudela, Universitat Politecnica de Valencia, Spain
Laura Garcia, Universidad Politécnica de Cartagena, Spain

GEOProcessing 2026 Technical Program Committee

Mahdi Abdelguerfi, University of New Orleans, USA
Alia I. Abdelmoty, Cardiff University, Wales, UK
Rushan Abeygunawardana, University of Colombo, Sri Lanka
Danial Aghajarian, Georgia State University, USA
Nuhcan Akçit, Middle East Technical University, Turkey
Zaher Al Aghbari, University of Sharjah, UAE
Hoda Allahbakhshi, UZH | Digital Society Initiative, Switzerland
Bharath H. Aithal, IIT Kharagpur, India
Rajendran Shobha Ajin, Resilience Development Initiative (RDI), Bandung, Indonesia
Md Mahbub Alam, Dalhousie University, Canada
Heba Aly, University of Maryland, College Park, USA
Oussama Annad, Higher School of Advanced Technologies (Ecole Nationale Supérieure des Technologies Avancées ENSTA), Algeria
Francisco Javier Ariza López, Escuela Politécnica de Jaén - Universidad de Jaén, Spain
Thierry Badard, Centre de Recherche en Géomatique (CRG) | Université Laval, Canada
Fabian Barbato, UDELAR-School of Engineering, Uruguay
Melih Basaraner, Yildiz Technical University, Turkey
Peter Baumann, rasdaman GmbH Bremen / Jacobs University Bremen, Germany
Mikael Brunila, McGill University, Canada
Lorenzo Carnevale, University of Messina, Italy
Mete Celik, Erciyes University, Turkey

Arun Sharma, University of Minnesota, Twin Cities, USA
Shih-Lung Shaw, University of Tennessee, Knoxville, USA
Yosio E. Shimabukuro, Brazilian Institute for Space Research - INPE, Brazil
Rajat Shinde, Indian Institute of Technology Bombay, India
Dericks Praise Shukla, IIT Mandi, India
Srushti Rashmi Shirish, Woven planet holdings, Tokyo, Japan
Francesco Soldovieri, Istituto per il Rilevamento Elettromagnetico dell'Ambiente - Consiglio Nazionale delle Ricerche (CNR), Italy
Katia Stankov, University of British Columbia, Canada
Behnam Tahmasbi, University of Maryland, College Park, USA
Ergin Tari, Istanbul Technical University, Turkey
Brittany Terese Fasy, Montana State University, USA
Roger Tilley, Sandia National Laboratories, USA
Goce Trajcevski, Iowa State University, USA
Linh Truong-Hong, Delft University of Technology, Netherlands
Taketoshi Ushiama, Kyushu University, Japan
Munazza Usmani, Fondazione Bruno Kessler, Trento, Italy
Marlène Villanova-Oliver, Univ. Grenoble Alpes - Grenoble Informatics Lab, France
Massimo Villari, University of Messina, Italy
Tin Vu, Microsoft Corporation, USA
Hong Wei, University of Maryland, College Park, USA
John P. Wilson, University of Southern California, USA
Jianhong Cecilia Xia, Curtin University, Australia
Ningchuan Xiao, The Ohio State University, USA
Yanan Xin, Institute of Cartography and Geoinformation | ETH Zurich, Switzerland
Daisuke Yamamoto, Nagoya Institute of Technology, Japan
Xiaojun Yang, Florida State University, USA
Zhaoming Yin, Google LLC, USA
Qiangqiang Yuan, School of Geodesy and Geomatics | Wuhan University, China
F. Javier Zarazaga-Soria, University of Zaragoza, Spain
Bingyi Zhang, University of Southern California, USA
Zenghui Zhang, Shanghai Jiao Tong University, China
Shenglin Zhao, Tencent, Shenzhen, China
Qiang Zhu, University of Michigan - Dearborn, USA
Rui Zhu, Institute of High Performance Computing - A*STAR, Singapore
Alexander Zipf, Heidelberg University, Germany

Copyright Information

For your reference, this is the text governing the copyright release for material published by IARIA.

The copyright release is a transfer of publication rights, which allows IARIA and its partners to drive the dissemination of the published material. This allows IARIA to give articles increased visibility via distribution, inclusion in libraries, and arrangements for submission to indexes.

I, the undersigned, declare that the article is original, and that I represent the authors of this article in the copyright release matters. If this work has been done as work-for-hire, I have obtained all necessary clearances to execute a copyright release. I hereby irrevocably transfer exclusive copyright for this material to IARIA. I give IARIA permission or reproduce the work in any media format such as, but not limited to, print, digital, or electronic. I give IARIA permission to distribute the materials without restriction to any institutions or individuals. I give IARIA permission to submit the work for inclusion in article repositories as IARIA sees fit.

I, the undersigned, declare that to the best of my knowledge, the article is does not contain libelous or otherwise unlawful contents or invading the right of privacy or infringing on a proprietary right.

Following the copyright release, any circulated version of the article must bear the copyright notice and any header and footer information that IARIA applies to the published article.

IARIA grants royalty-free permission to the authors to disseminate the work, under the above provisions, for any academic, commercial, or industrial use. IARIA grants royalty-free permission to any individuals or institutions to make the article available electronically, online, or in print.

IARIA acknowledges that rights to any algorithm, process, procedure, apparatus, or articles of manufacture remain with the authors and their employers.

I, the undersigned, understand that IARIA will not be liable, in contract, tort (including, without limitation, negligence), pre-contract or other representations (other than fraudulent misrepresentations) or otherwise in connection with the publication of my work.

Exception to the above is made for work-for-hire performed while employed by the government. In that case, copyright to the material remains with the said government. The rightful owners (authors and government entity) grant unlimited and unrestricted permission to IARIA, IARIA's contractors, and IARIA's partners to further distribute the work.

Table of Contents

Physics-Informed Signed Distance Fields for Flood Arrival Time Prediction and Evacuation Routing

Yameng Guo
yameng.guo@ugent.be
dept. Business Informatics and Operations Management
Ghent University
Gent, Belgium

Seppe vanden Broucke
seppe.vandenbroucke@ugent.be
dept. Business Informatics and Operations Management
Ghent University
Gent, Belgium
Research Centre for Information Systems Engineering
KU Leuven
Leuven, Belgium

***Abstract*—Accurate flood prediction with explicit arrival times is critical for effective evacuation planning. Traditional physics-based models provide reliable predictions but are computationally expensive, while existing data-driven approaches typically predict binary flood extent or water depth without explicit timing information. We propose a novel framework that represents flood evolution as a temporal Signed Distance Field (SDF), enabling efficient prediction of both flood boundaries and arrival times. Our approach combines spatial-temporal deep learning with physics-informed constraints in a variational optimization framework. The predicted SDF naturally supports time-dependent evacuation routing by providing continuous distance-to-flood information and explicit arrival time fields. The proposed idea will be compared with grid-based deep learning approaches as well as physical constraints approaches regarding real-time evacuation planning and arrival time accuracy.**

Keywords-Signed distance function; Flood prediction; Evacuation routing.

I. Introduction

Flooding is one of the most devastating natural disasters, affecting millions worldwide annually. Effective evacuation requires knowledge of not only *where* flooding will occur, but critically *when* each location will be inundated. Traditional approaches face a fundamental trade-off between computation time and accuracy performance. Physics-based models (e.g., Hydrologic Engineering Center's River Analysis System (HEC-RAS) [1], LISFLOOD-FP [2]) provide reliable predictions by solving shallow water equations but require hours of computation, while data-driven methods [3] offer speed but typically predict only binary flood extent or depth fields without explicit timing.

To address this, we propose representing flood evolution as a Signed Distance Field (SDF) [4][5], where $\varphi(x, y, t)$ encodes the signed distance to the flood boundary at each location and time. Compared to traditional ML grid-based methods and physical modeling, SDF provides advantages in both geometric encoding and physics constraint formulation. The direct and continuous encoding of flood extent enables natural extraction of arrival times from level set evolution and helps the model learn continuous boundaries rather than discrete water depths. Additionally, SDF uses information more efficiently for sparse flooding patterns. For instance, for zero values, depth grids will waste the capacity on zeros, while the distance field utilizes zero as meaningful input indicating the flood boundary. Moreover, SDF enables simpler, geometrically cleaner physical constraints that allow constraining boundary motion without modeling all fluid dynamics.

In summary, this work aims to introduce the following key contributions:

- Apply temporal SDF evolution to flood prediction with explicit arrival time modeling.
- Use a physics-informed variational optimization framework to ensure realistic propagation.
- Enable end-to-end system from prediction to evacuation routing with safety guarantees.
- Implement comprehensive evaluation against both physics-based and data-driven baselines.

To the best of our knowledge, this is the first work applying temporal SDF evolution to flood prediction.

The paper is organized as follows: In Section III, we discuss the methodology of using SDF representations, the model architecture, and the design of the physics-aware loss function. In Section IV, we describe the experimental setup, including the dataset, baselines, and selected evaluation metrics. The initial results are presented in Section V, and we conclude in Section VI.

II. Related Work

SDFs have been widely used in computer vision [6], 3D graphics [7], and robotic systems [8], where they play a critical role in geometry processing, mapping, and planning.

Beyond these domains, SDFs have also been applied in structural engineering tasks such as topology optimization [5], and have been extended to incorporate physical awareness in scenarios such as fluid simulation [9].

Despite their success in various physics-aware applications, SDFs remain largely unexplored in the context of flood modeling. Current flood modeling approaches are predominantly divided into two categories: physics-based models and machine learning or deep learning models [10][11], typically leveraging multisource data such as rainfall, Digital Elevation Models (DEMs), and land surface characteristics. Although the methodological spectrum continues to expand, the underlying data

 ISBN: 978-1-68558-384-2

representation and prediction targets remain largely unchanged. That is, the study area is typically discretized into grid cells, and the flood state is predicted for each cell.

In this context, SDF-based representations offer a promising alternative by encoding geometric and spatial information in a continuous and physically meaningful manner. Unlike discrete grid-based representations, SDFs naturally capture boundary information and spatial relationships, which are critical for accurately modeling flood dynamics such as water propagation and interaction with terrain. Furthermore, SDFs can be seamlessly integrated with machine learning and deep learning models, providing rich geometric priors that may improve generalization and robustness.

III. Methodology

A. SDF Representation

We define the flood state at time t as:

$$\varphi(x,y,t) = \begin{cases} d((x,y),\partial\Omega_t) & \text{if } (x,y) \in \Omega_t \\ -d((x,y),\partial\Omega_t) & \text{otherwise} \end{cases} \quad (1)$$

where Ω_t is the flooded region, $\partial\Omega_t$ its boundary, and $d(\cdot,\cdot)$ is Euclidean distance. The zero level set $\{\varphi = 0\}$ represents the flood boundary.

Arrival time at location (x, y) is extracted as:

$$t_{\text{arr}}(x,y) = \min\{t : \varphi(x,y,t) \geq 0\} \quad (2)$$

The representation shown in Figure 1 illustrates the difference between a grid-based field and a signed-distance field. In the grid-based field, each discrete cell has its own target water-depth value, whereas in the SDF representation, the target is a continuous boundary of the water area.

B. Temporal SDF Evolution Model

Our model predicts the SDF trajectory $\{\varphi_t\}_{t=0}^{T}$ given initial conditions φ_0, terrain elevation z, and rainfall forcing $r(t)$.

The proposed hybrid spatial-temporal architecture consists of three components, including a spatial encoder-decoder to extract multi-scale spatial features from $[\varphi_t, z, r_t]$, a temporal model to capture long-range temporal dependencies efficiently and a physics layer to enforce constraints during decoding.

C. Physical-constraints Loss

We formulate training as a constrained optimization problem with a physical-constraints loss applying penalty on both data fitting error and physical feasibility.

$$\mathcal{L} = \underbrace{\|\varphi_{\text{pred}} - \varphi_{\text{obs}}\|^2}_{\text{data fit}} + \lambda_1 \underbrace{\|\partial_t\varphi + v\cdot\nabla\varphi\|^2}_{\text{level set PDE}} + \lambda_2 \underbrace{\|\text{ReLU}(\nabla\varphi\cdot\nabla z)\|^2}_{\text{downhill flow}} + \lambda_3 \underbrace{\|\text{volume}_t - \text{rain}_t\|}_{\text{mass balance}} \quad (3)$$

The level-set term enforces $\partial\varphi/\partial t + v\cdot\nabla\varphi = 0$ where v is flood velocity (estimated from terrain/rainfall). This ensures the SDF evolves consistently with flood propagation physics. The downhill flow term is to enforce that flood propagation does not move uphill, and the last mass term is to ensure flooded water volume match rainfall input over time.

D. Evacuation Routing

Given the predicted SDF trajectory and arrival time field, we solve time-dependent evacuation routing as follows ("safe zone computation"): at current time t_c, location (x, y) is safe if:

$$t_{\text{arr}}(x,y) > t_c + \tau_{\text{buffer}} \quad (4)$$

IV. Experimental Design

A. Datasets

We plan to evaluate our approach on both real and synthetic datasets. The real-life dataset includes historical flood events with satellite-derived inundation maps, Digital Elevation Models (DEMs), and rainfall records. In our initial experiments, we use a UK flood event dataset from FloodCastBench [12].

However, existing real-world datasets are significantly limited by the coarse granularity of rainfall data, the relatively small number of recorded flood events, and the frequent absence of a long-term temporal dimension. To address these limitations, we plan to incorporate synthetic data in future work, particularly for rainfall generation and flood extent simulation. Specifically, we will leverage the open-source package Landlab [13], which supports numerical modeling of Earth surface dynamics, including rainfall generation and water-driven erosion processes.

B. Baselines

Two baseline categories will be included, i.e., physics-based models and grid-based deep learning approaches. Physics-based baselines, such as the industry-standard HEC-RAS 2D and LISFLOOD-FP for large-scale floods, can be included into comparison experiment. For deep learning baselines, we use convolutional models with Long Short-Term Memory (LSTM) kernels to predict water depth at each cell over time.

C. Evaluation Metrics

Performance will be evaluated across three dimensions: flood prediction accuracy, evacuation performance, and computational cost.

For flood prediction, we will use Intersection over Union (IoU) [14][15] to measure the overlap between predicted and ground-truth flood masks. We will also evaluate arrival time accuracy using Root Mean Square Error (RMSE) between predicted and actual inundation times.

For evacuation performance, we measure success rate, i.e., the percentage of evacuations reaching safety before flood arrival. Finally, we record time and memory usage during training and inference to assess computational cost.

V. Initial Results

We conducted an initial proof-of-concept study on the UK event from the FloodCastBench data set [12] to assess whether an SDF target already provides practical value before introducing the full physics-informed training setup. Each sample consists of six past flood states together with the DEM height map and a scalar mean-rainfall summary. The model predicts the next six flood states in a direct sequence-to-sequence manner. To keep the comparison controlled, we used

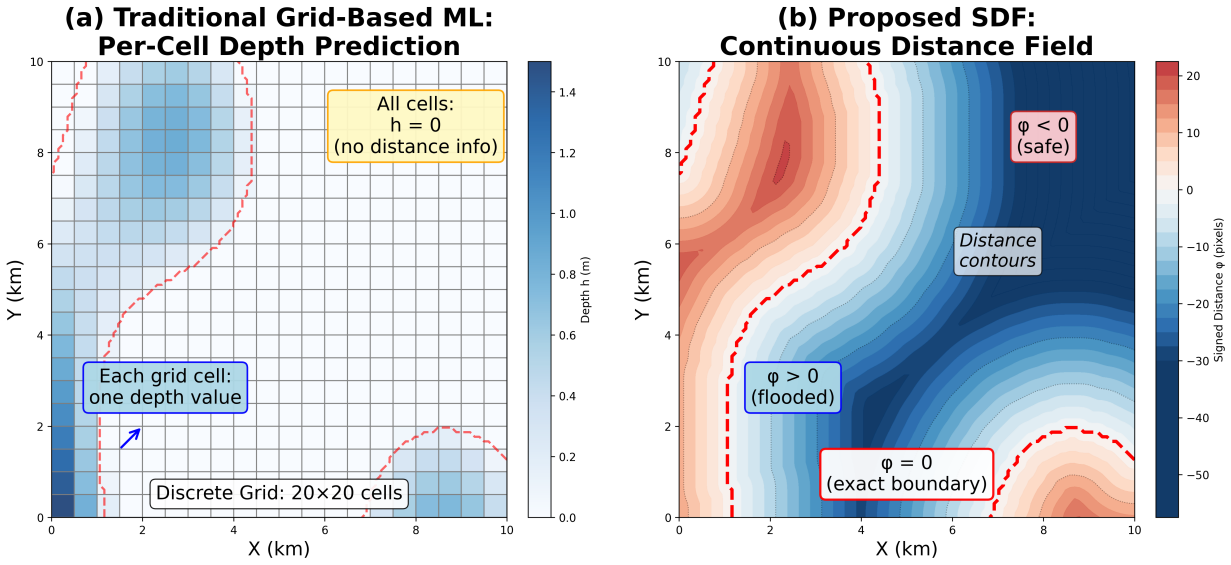

Figure 1. Comparison of flood area representations using a grid-based method and a signed-distance field.

the same small U-Net style backbone across all runs and only varied the prediction target. The resulting baselines compare a signed-distance field with a classic coarse cell-based mask representation.

The targets are organized as follows. For the mask approach, each grid cell is assigned a binary label indicating whether the corresponding area is flooded. For the SDF representation, the target takes continuous values between -1 and 1, where -1 denotes non-flooded areas, 0 represents the SDF boundary, and 1 indicates flooded areas.

The UK event targets were cleaned up and evaluated using a chronological train/validation/test split. We incorporated four complementary evaluation metrics. Flood-extent overlap is measured using Intersection over Union (IoU),

$$\mathrm{IoU} = \frac{|P \cap G|}{|P \cup G|}, \tag{5}$$

and Dice,

$$\mathrm{Dice} = \frac{2|P \cap G|}{|P| + |G|}, \tag{6}$$

where P and G denote the predicted and ground-truth flooded regions. We additionally consider a boundary F1 score to assess whether predicted and reference flood boundaries match within a small spatial tolerance, and an arrival-time mean absolute error

$$\mathrm{F1}_\partial = \frac{2\,\mathrm{Prec}_\partial\,\mathrm{Rec}_\partial}{\mathrm{Prec}_\partial + \mathrm{Rec}_\partial}, \tag{7}$$

where Prec_∂ and Rec_∂ denote boundary precision and recall after allowing a small tolerance around the reference contour, and an arrival-time mean absolute error

$$\mathrm{MAE}_{\mathrm{arr}} = \frac{1}{N}\sum_{i=1}^{N} \left|\hat{t}_{\mathrm{arr},i} - t_{\mathrm{arr},i}\right| \tag{8}$$

to summarize timing differences between predicted and observed first inundation.

Under this setup, the SDF approach achieved slightly improved flood-extent performance compared to the coarse-cell mask baseline, as shown in Table I. This suggests that, at least for the current direct-prediction backbone, the SDF target is viable and already competitive as a representation, even though it does not yet yield a decisive aggregate-metric advantage over the mask baseline. This is mainly due to the fact that this event is generally stable in terms of flood-extent and that after an initial on-fall of rain, the flood depths remain relatively stable over time.

TABLE I. INITIAL HELD-OUT RESULTS ON THE CLEANED UK FLOOD EVENT.

Model	IoU	Dice	$\mathrm{F1}_\partial$	$\mathrm{MAE}_{\mathrm{arr}}$
Coarse mask	0.9975	0.9987	0.9998	0.0074
SDF	0.9976	0.9988	0.9988	0.0059

At the same time, the qualitative behavior of the SDF target is encouraging. As can be observed in the comparison of Figure 2, Figure 3 and Figure 4, the SDF representation yields a continuous geometric field that is easier to overlay on top of terrain and more naturally highlights evolving flood structure than a purely binary occupancy map. On the particular frame shown, the SDF approach presents much better evaluation metrics than the averaged results in Table I. The initial experiments therefore support the core motivation of the paper: SDF is a promising geometric representation for flood prediction, even though larger gains will likely depend on stronger temporal models and the full physics-informed formulation proposed in this work.

These initial results refine the direction of the full study. Rather than assuming that SDF will dominate grid-based targets out of the box, the experiments indicate that the most promising benefit of SDF lies in boundary-aware structure, interpretability, and its compatibility with future physics-informed constraints.

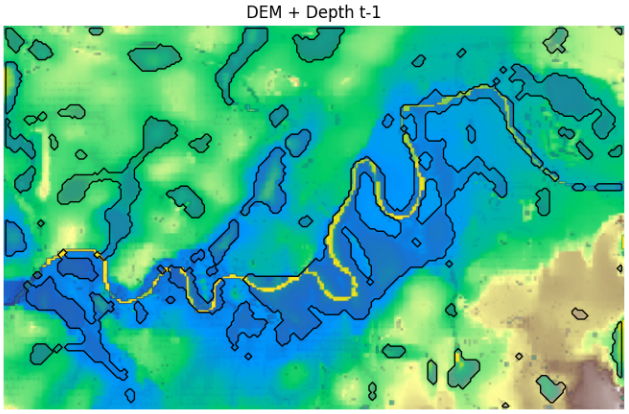

Figure 2. Initial DEM of the cleaned UK flood event.

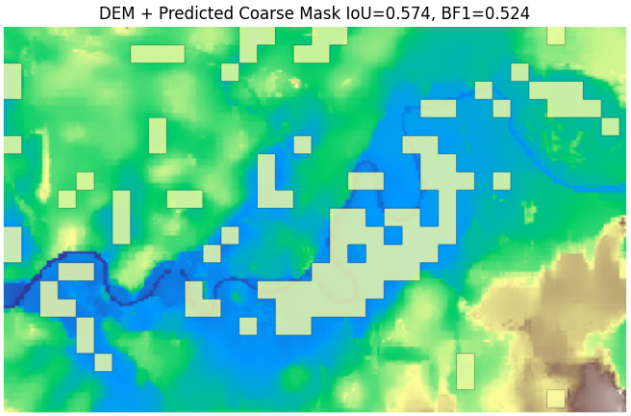

Figure 3. The results of coarse cell-based mask prediction (without explicit flood boundary delineation).

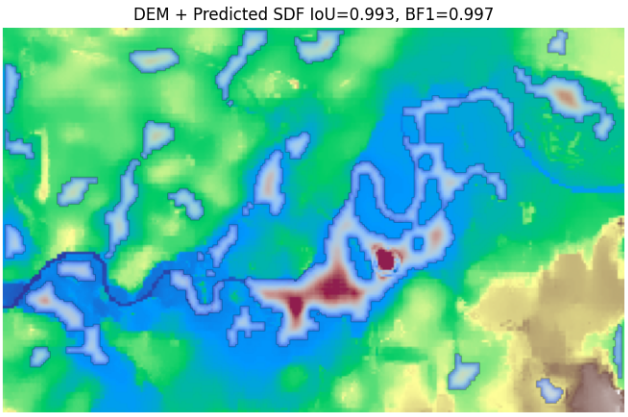

Figure 4. The continuous signed-distance field prediction (with explicit flood boundary delineation).

This makes SDF a strong candidate representation for the next phase of work.

VI. Conclusion and Future Work

We present a novel physics-informed SDF framework for flood prediction that combines geometric representation, temporal evolution modeling, and optimization-based constraints. By explicitly targeting arrival times alongside flood extent, the framework is intended to bridge the gap between fast data-driven methods and reliable physics-based models. Our initial UK case study already indicates that SDF is a viable and visually compelling representation, with performance comparable to a grid-based mask baseline under a matched small U-Net backbone and with clearer geometric structure for qualitative analysis.

At the same time, several parts of the full idea remain to be implemented and tested in the next phase. First, the current results are based on simple direct-prediction baselines rather than the full physics-informed temporal architecture. Second, the physical-constraints loss still needs to be integrated and evaluated systematically, including the level-set evolution term, downhill-flow consistency, and mass-balance constraints. Third, the arrival-time and evacuation-routing components have been formulated conceptually but still need end-to-end empirical validation. Finally, broader experiments are required to test whether the representation advantage of SDF becomes more pronounced under stronger temporal backbones (e.g. autoregressive backbone models), harder forecasting regimes, and more realistic operational settings.

Overall, the present study provides encouraging early evidence for the representation itself and clarifies the most promising next steps for turning the proposed SDF framework into a complete physics-informed flood forecasting and evacuation-support system.

References

[1] G. W. Brunner, “Hec-ras river analysis system. hydraulic reference manual. version 1.0.”, *User manual. US Army Corps of Engineers, Hydrologie Engineering Center*, 1995.

[2] P. Bates, M. Trigg, J. Neal, and A. Dabrowa, “Lisflood-fp”, *User manual. School of Geographical Sciences, University of Bristol. Bristol, UK*, 2013.

[3] A. Mosavi, P. Öztürk, and K. Chau, “Flood prediction using machine learning models: Literature review”, *CoRR*, vol. abs/1908.02781, 2019. arXiv: 1908.02781.

[4] S. Osher and R. Fedkiw, “Signed distance functions”, in *Level set methods and dynamic implicit surfaces*, Springer, 2003, pp. 17–22.

[5] O. Jezek et al., “Smooth geometry extraction from SIMP topology optimization: Signed distance function approach with volume preservation”, *Adv. Eng. Softw.*, vol. 212, p. 104 071, 2026. DOI: 10.1016/J.ADVENGSOFT.2025.104071.

[6] D. Vicini, S. Speierer, and W. Jakob, “Differentiable signed distance function rendering”, *ACM Trans. Graph.*, vol. 41, no. 4, 125:1–125:18, 2022. DOI: 10.1145/3528223.3530139.

[7] J. Shim, C. Kang, and K. Joo, “Diffusion-based signed distance fields for 3d shape generation”, in *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, IEEE, 2023, pp. 20 887–20 897. DOI: 10.1109/CVPR52729.2023.02001.

[8] B. Liu, G. Jiang, F. Zhao, and X. Mei, “Collision-free motion generation based on stochastic optimization and composite signed distance field networks of articulated robot”, *IEEE Robotics Autom. Lett.*, vol. 8, no. 11, pp. 7082–7089, 2023. DOI: 10.1109/LRA.2023.3311357.

[9] C. Zhang, “Sdfibm: A signed distance field based discrete forcing immersed boundary method in openfoam”, *Comput. Phys. Commun.*, vol. 255, p. 107 370, 2020.

[10] D. K. Hakim, R. Gernowo, and A. W. Nirwansyah, "Flood prediction with time series data mining: Systematic review", *Natural Hazards Research*, vol. 4, no. 2, pp. 194–220, 2024.

[11] F. Yang et al., "Rapid urban flood inundation forecasting using a physics-informed deep learning approach", *Journal of Hydrology*, vol. 643, p. 131 998, 2024.

[12] Q. Xu, Y. Shi, J. Zhao, and X. X. Zhu, "Floodcastbench: A large-scale dataset and foundation models for flood modeling and forecasting", *Scientific Data*, vol. 12, no. 1, p. 431, 2025.

[13] K. R. Barnhart et al., "Landlab v2. 0: A software package for earth surface dynamics", *Earth Surface Dynamics*, vol. 8, no. 2, pp. 379–397, 2020. DOI: 10.5194/esurf-8-379-2020.

[14] H. Rezatofighi et al., "Generalized intersection over union: A metric and a loss for bounding box regression", in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, Computer Vision Foundation / IEEE, 2019, pp. 658–666. DOI: 10.1109/CVPR.2019.00075.

[15] Y. Wang et al., "A data-driven approach for flood prediction using grid-based meteorological data", *Hydrological Processes*, vol. 37, no. 3, e14837, 2023.

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.

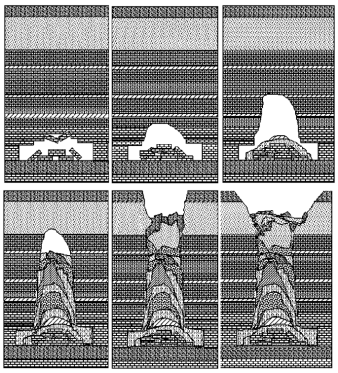

Figure 1. Stages of sinkhole formation [3].

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.

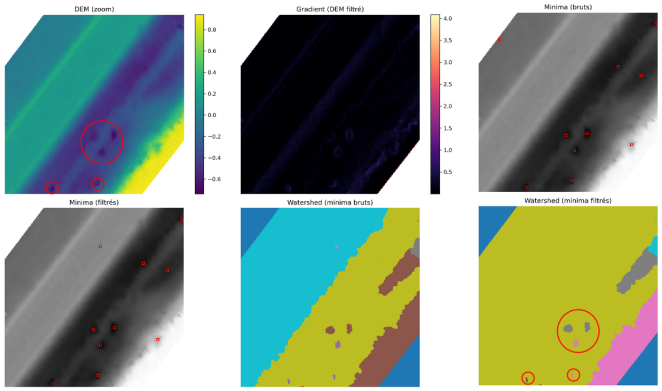

Figure 2. Overview of the non-learning watershed segmentation pipeline. From DEM and gradient computation to raw and filtered minima, and final watershed segmentation. Highlighted regions represent candidate basins used for feature extraction and supervised classification.

Table I
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

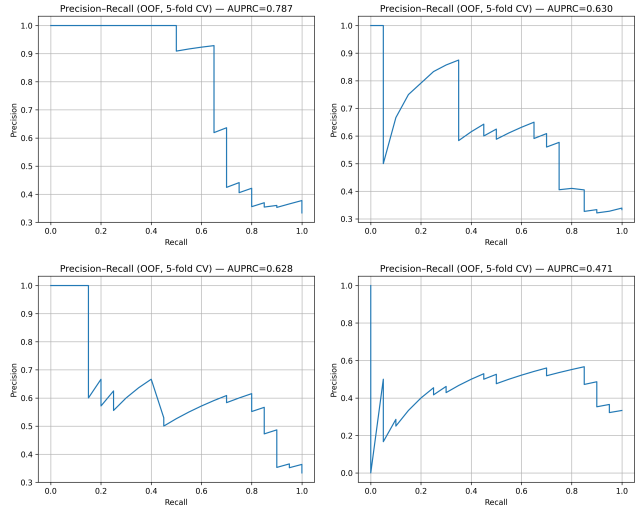

Figure 3. Precision–recall curves (5-fold cross-validation). (a) RF: 0.79, (b) LGBM: 0.63, (c) SVM: 0.63, (d) LR: 0.47.

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

Table II
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

is particularly relevant given the limited number of positive samples.

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
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.

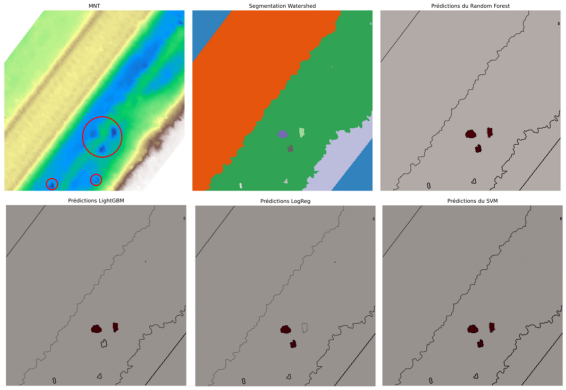

Figure 4. Comparative visualization of final predictions for the four models: original DEM (top left) with true sinkholes outlined in red, watershed segmentation (top center), Random Forest predictions (top right), LightGBM (bottom left), Logistic Regression (bottom center), and SVM (bottom right). Basins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

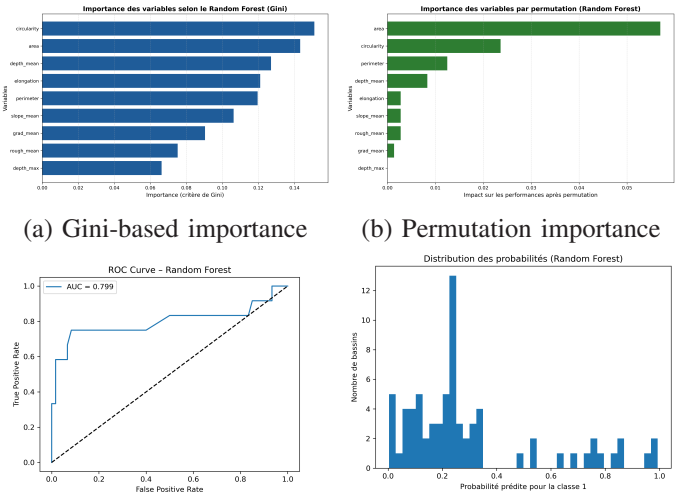

(a) Gini-based importance (b) Permutation importance

(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.

 ISBN: 978-1-68558-384-2