

A Hybrid Remote Sensing Pipeline for Landslide and Soil Scar Mapping

Gabriela Alves Marinho, Vinicius Nunes Viana da Costa and Eronides F. da Silva Neto

Abstract—This paper presents a hybrid pipeline for monitoring terrain degradation in satellite imagery, focusing on landslide and soil-scar mapping. A deep segmentation network (FCDU-Net) identifies candidate scar regions, followed by a supervised patch-level classifier that validates them using spectral and textural features. Both stages use RGB, an RGB-derived vegetation proxy (pseudo-NDVI), and GLCM texture descriptors selected via ReliefF. Across 30 runs on 7,092 patches, the approach improves mean F1-score (0.864 to 0.905) and area under the curve (0.955 to 0.983) over segmentation alone, with statistically significant gains (Welch t-test, $p < 0.001$, Cohen’s $d \approx 1.35$), supporting explicit stage separation.

Keywords— Remote sensing, landslide mapping, terrain degradation, deep learning, semantic segmentation, patch classification.

I. INTRODUCTION

Soil erosion and rainfall-induced landslides cause tens of billions of tons of soil loss per year globally [1], damaging infrastructure, agricultural productivity, and human lives, particularly in countries with mountainous relief and dense urban occupation. Early detection is difficult because degraded areas often appear mixed with exposed soil, unpaved roads, recent agricultural disturbance, rocky outcrops, and heterogeneous vegetation terrain elements that share visual characteristics with true scar regions in optical imagery.

Recent progress in semantic segmentation has improved pixel-level delineation of landslide scars. However, local masks remain vulnerable to spectrally similar terrain elements, generating false alarms over bare soil and illumination artifacts. Conversely, purely patch-level classifiers produce binary decisions without spatial evidence. In practical hazard mapping, both are needed: localization and confirmation.

This paper proposes a hybrid remote sensing pipeline that explicitly separates these two tasks. A deep segmentation network first identifies candidate scar regions at the pixel level; a supervised patch-level classifier then evaluates broader contextual evidence before confirming the event. The method combines Red–Green–Blue (RGB) imagery with an RGB-derived vegetation proxy, referred to as pseudo-NDVI, and Gray-Level Co-occurrence Matrix (GLCM) texture descriptors. Conventional NDVI estimates vegetation condition by comparing red and near-infrared (NIR) reflectance. Because the available images contain only visible RGB channels and no true NIR band, this study uses a calibrated green–red

approximation. GLCM descriptors summarize local texture patterns, while ReliefF selects the most discriminative auxiliary features.

The approach is guided by three assumptions. First, degradation mapping benefits from spatial localization, because a useful model should indicate the region supporting its prediction. Second, contextual reasoning is necessary, because local pixel textures alone are often insufficient to confirm a true event. Third, richer descriptors are preferable to RGB alone, because vegetation loss and terrain texture provide complementary evidence directly related to landslide dynamics. These assumptions motivate a pipeline in which segmentation and classification are treated as complementary stages: the first delimiting suspicious structures, the second confirming whether the broader patch context supports a genuine degradation event.

The main contributions of this work are: (i) a modular two-stage architecture that combines segmentation interpretability with the robustness of patch-level global confirmation; and (ii) a statistically rigorous evaluation protocol based on 30 independent runs with synchronized seeds, demonstrating consistent performance gains and reduced inter-run variance over a strong segmentation-only baseline.

This paper is organized as follows. Section II reviews related work. Section III describes the proposed method. Section IV presents the experimental setup. Section V reports results and analysis. Section VI concludes with limitations and future directions.

II. RELATED WORK

U-Net and densely connected encoders remain important references for pixel-level mapping [2], [3], but recent GeoAI research increasingly addresses complex backgrounds, scale variation, uncertain boundaries, and long-range context through attention, Transformers, and multi-scale fusion [4]. Foreground-aware convolutional attention suppresses responses from visually similar structures such as bare soil, roads, shadows, and agricultural areas [5], whereas multi-branch CNN–Transformer designs integrate local detail and global context within the segmenter [6]. These trends improve internal feature discrimination but generally retain a single segmentation network as the final decision mechanism.

A second direction combines optical imagery with physically richer observations and evaluates transfer across regions. Hybrid CNN–Transformer mapping has incorporated Sentinel-2 spectral bands, indices, DEM-derived topography, and deep transfer learning [7]; recent benchmarks span multiple regions, sensors, and backbones [8]. Sen12Landslides further

Gabriela Alves Marinho Universidade de Pernambuco, Recife-PE, Brazil. E-mail: gabriela.amarinho@upe.br; Vinicius Nunes Viana da Costa, Universidade de Pernambuco, Recife-PE, Brazil., E-mail: vinicius.nvcosta@upe.br; Eronides F. da Silva Neto, Department of Electronics and Systems, Federal University of Pernambuco (UFPE), Recife-PE, Brazil. E-mail: eronides.silvaneto@ufpe.br

provides multi-temporal Sentinel-1 SAR, Sentinel-2, elevation, and event metadata for geographically diverse evaluation [9]. These sources contain information unavailable from RGB alone. Accordingly, the present study is not framed as multi-spectral or multimodal: it targets RGB-only settings using an RGB-derived vegetation proxy and GLCM texture descriptors [10], selected with ReliefF [11].

Joint classification–segmentation designs are less common than single-network segmenters. The conceptually closest work is ICSSN [12], which shares a feature extractor, uses object- and sub-object-level contrastive learning, and iteratively couples classification and segmentation during training. In contrast, the present pipeline uses independent branches trained separately and combines them only at inference through explicit decision-level fusion: segmentation preserves pixel-level localization, while the classifier provides patch-level semantic confirmation. Patch classification is also explored with multi-encoder Vision Transformers and ensemble learning over Sentinel-1 and Sentinel-2 data [13], but such global decisions do not by themselves delineate landslide boundaries.

Thus, the contribution is not a new segmentation backbone. Most recent methods seek to resolve ambiguity by increasing internal segmenter complexity or adding attention, multi-scale context, transfer learning, and complementary sensing modalities. This work instead investigates a modular decomposition in which independently observable spatial and semantic evidence are combined by a transparent rule. Because operational efficiency also depends on parameter count, preprocessing, latency, throughput, and memory—quantities explicitly assessed in recent lightweight segmentation studies [14]—no computational advantage is claimed without corresponding measurements.

III. PROPOSED METHOD

Fig. 1 shows the pipeline as a block diagram. The shared input (RGB + pseudo-NDVI + GLCM, filtered by ReliefF) feeds FCDU-Net on the segmentation branch, producing a binary scar mask, and the patch encoder on the classification branch, producing event probability p . The training phase applies data augmentation before feature extraction; no augmentation is used at inference. A logical AND gate emits a positive label only when both branches agree.

A. Input representation

Each sample is represented as a 128×128 patch. RGB channels provide the visible-color information. In multispectral remote sensing, vegetation condition is commonly estimated using the Normalized Difference Vegetation Index (NDVI), which compares red reflectance with near-infrared (NIR) reflectance. However, the images used in this study contain only visible RGB channels and do not include a true NIR band. Therefore, an RGB-derived vegetation proxy, hereafter referred to as pseudo-NDVI, was computed by replacing the unavailable NIR band with a calibrated linear combination, $\text{pseudo-NIR} = aR + bG + cB$. The coefficients were selected




Fig. 1

BLOCK DIAGRAM OF THE PROPOSED HYBRID PIPELINE. TRAINING PHASE: PATCHES UNDERGO DATA AUGMENTATION (FLIPS, 90 ROTATIONS) BEFORE FEATURE EXTRACTION; FCDU-NET (STAGE 1, DICE LOSS) AND THE PATCH CLASSIFIER (STAGE 2, BCE LOSS) ARE TRAINED INDEPENDENTLY.

INFERENCE PHASE: NO AUGMENTATION; SEGMENTATION BRANCH PRODUCES $\text{SEG}(x)$ IN $\{0,1\}$; CLASSIFIER BRANCH PRODUCES $p(x)$ IN $[0,1]$; AND GATE EMITS $\hat{y} = 1$ IFF $\text{SEG}(x) = 1$ AND $p(x) \geq \tau$.

to maximize the separation between labeled erosion and non-erosion pixels in the calibration data, yielding $a = 0$, $b = 1$, and $c = 0$. Thus, the pseudo-NIR corresponds to the green channel, and the proxy used in this study is

$$\text{pseudo-NDVI} = \frac{G - R}{G + R}, \quad (1)$$

where R and G denote the red and green channels, respectively. This quantity should be interpreted as a normalized green–red vegetation proxy rather than conventional multi-spectral NDVI, which requires a true near-infrared band. The proxy is intended to emphasize variations associated with vegetation cover and exposed soil.

Texture information is extracted through gray-level co-occurrence matrices computed over 11×11 sliding windows, offset distance 1 pixel, four directions (0° , 45° , 90° , 135°), averaged over directions. Eight descriptors are retained: contrast, dissimilarity, homogeneity, angular second moment, energy, correlation, entropy, and variance.

Because not every auxiliary channel contributes equally, ReliefF-based feature selection is applied before training. In the reported experiments, ReliefF retained four attributes—pseudo-NDVI, energy, entropy, and correlation—which were concatenated with the RGB channels to form the final seven-channel input tensor.

B. Segmentation stage

The first stage is FCDU-Net [15], a dense U-Net variant combining skip connections and multi-scale feature reuse. Its role is to produce preliminary masks indicating the most likely scar pixels inside each patch. The architecture follows the encoder-decoder principle with four levels of downsampling

and upsampling, using 3×3 convolutional blocks with Batch Normalization and ReLU activations. Training minimizes a weighted Dice loss to address pixel-level class imbalance. This stage addresses the localization question: where is the suspicious structure in the patch?

C. Patch-level classification stage

The second stage is a supervised patch classifier trained independently from the segmentation network. The classifier contains four convolutional blocks (32, 64, 128, and 256 filters, respectively), each with 3×3 convolutions, ReLU activation, Batch Normalization, and max pooling, followed by global average pooling, a fully connected layer with 128 neurons (ReLU, L2 regularization), dropout of 0.35, and a final sigmoid unit.

The classifier is optimized with binary cross-entropy,

$$L_{\text{BCE}} = -[y \log(p) + (1 - y) \log(1 - p)] \quad (2)$$

where p is the predicted probability that the patch contains a landslide or erosion scar. To compensate for class imbalance, class weights are used:

$$w_{\text{pos}} = \frac{N}{2N_{\text{pos}}}, \quad w_{\text{neg}} = \frac{N}{2N_{\text{neg}}} \quad (3)$$

where N is the total number of samples and N_{pos} , N_{neg} are the class counts. All layers are trained from scratch; an encoder pre-trained for dense segmentation reduced AUC and increased inter-run variance when repurposed for global binary classification.

D. Hybrid decision rule

The final decision is made only when both stages agree:

$$\hat{y} = 1 \iff \text{Seg}(x) = 1 \wedge p(x) \geq \tau. \quad (4)$$

Threshold τ is selected by a grid search over $\{0.50, 0.55, \dots, 0.80\}$ on the validation subset of each run independently, maximizing F1. The test split is never used for threshold selection, avoiding data leakage. This rule is intentionally transparent. Instead of hiding the final alarm inside a single opaque module, it exposes two complementary forms of evidence: local delineation and global semantic confirmation. Bare soil or shadows may activate the segmentation branch, but the classifier can reject these cases when the patch context does not support a true scar. Conversely, fragmented or thin scars can still be recovered when the global patch evidence remains consistent.

IV. EXPERIMENTAL SETUP

The experimental samples were derived from the public Bijie Landslide Dataset introduced by Ji et al. [16]. The original dataset covers Bijie, Guizhou Province, China, and contains optical TripleSat imagery acquired between May and August 2018, together with DEM data and manually delineated landslide boundaries subjected to inspection in the

original study. Instance boundaries are distributed as polygon coordinates in `.txt` files and raster masks. After the preprocessing and patch-extraction procedure adopted in this study, the experimental corpus comprised 7,092 image patches, including 3,084 positive patches with landslide scars and 4,008 negative patches. Positive patches were associated with binary masks, and all samples were standardized to 128×128 pixels. Data augmentation includes horizontal and vertical flips and 90° -multiple rotations, which preserve plausible geometric properties for slope structures while enriching the training distribution.

Each execution uses a stratified train-validation-test split. The study employs 30 independent runs with different random seeds, synchronized across Python, NumPy, and TensorFlow, controlling data splits, weight initialization, and batch ordering. This reduces the influence of stochastic variability typical of deep models and enables a more reliable estimate of mean performance and dispersion.

Both stages are trained with Adam ($\text{lr} = 3 \times 10^{-4}$), batch size 16, early stopping, and ReduceLROnPlateau. For the segmentation baseline, the best operating threshold and minimum area are selected by sweep on the validation set. For the hybrid system, the same features are used in both branches so that the comparison isolates the contribution of the second decision stage. Table I summarizes the main settings.

TABLE I
MAIN EXPERIMENTAL SETTINGS.

Setting	Value
Patch size	128×128
Positive patches	3,084
Negative patches	4,008
Optimizer	Adam
Learning rate	3×10^{-4}
Batch size	16
Runs	30
Augmentation	flips + 90° rotations
τ selection	grid $\{0.50, \dots, 0.80\}$, val., per run

Reported patch-level metrics are F1, precision, accuracy, and AUC, expressed as mean $\pm$ standard deviation over 30 runs. The segmentation baseline additionally reports IoU at threshold 0.5. Statistical significance is assessed with a Welch t-test on the F1 distributions; practical significance with Cohen's d .

V. RESULTS AND DISCUSSION

Table II presents the comparison between the segmentation baseline and the proposed hybrid method over 30 runs. The hybrid design improved every metric retained in the comparative analysis. Mean F1 increased from 0.864 ± 0.038 to 0.905 ± 0.020 , precision from 0.872 ± 0.052 to 0.911 ± 0.025 , accuracy from 0.919 ± 0.041 to 0.947 ± 0.011 , and AUC from 0.955 ± 0.018 to 0.983 ± 0.006 .

Because the final decision requires agreement between both branches, operational deployment should include a threshold-sensitivity analysis aligned with the acceptable miss rate for the target terrain.

TABLE II

COMPARISON BETWEEN THE SEGMENTATION BASELINE AND THE PROPOSED HYBRID SYSTEM OVER 30 RUNS.

Metric	Baseline	Proposed
F1	0.864 ± 0.038	0.905 ± 0.020
Precision	0.872 ± 0.052	0.911 ± 0.025
Accuracy	0.919 ± 0.041	0.947 ± 0.011
AUC	0.955 ± 0.018	0.983 ± 0.006
IoU (seg.)	0.760	— *

* The proposed system produces a patch-level binary decision rather than a pixel mask; segmentation IoU is therefore not directly comparable for the final stage.

Pipeline: FCDU-Net Segmentation → Patch-level Classification




Fig. 2

PIPELINE: FCDU-NET SEGMENTATION FOLLOWED BY PATCH-LEVEL CLASSIFICATION. LEFT TO RIGHT: ORIGINAL IMAGE, PROBABILITY HEATMAP, BINARY MASK AFTER POST-PROCESSING ($\tau=0.50$, $\text{MIN_AREA}=40$ PX), AND IMAGE-LEVEL CLASSIFICATION RESULT.

Fig. 2 presents qualitative examples of the operation of the proposed pipeline. In the positive example, FCDU-Net produces a high-probability region spatially consistent with the visible landslide scar, while the classifier confirms the patch-level class. In the negative case, the probability map remains close to zero, and the sample is correctly rejected by the classifier.

These gains are meaningful beyond average performance. The lower standard deviation of the proposed method indicates more stable behavior across runs, which is desirable in real monitoring systems. A Welch t-test applied to the F1 distributions yielded $p < 0.001$, supporting the claim that the improvement is statistically significant rather than an artifact of split selection. The observed effect size is large (Cohen's $d \approx 1.35$), reinforcing that the confirmation stage changes the effective modeling capacity of the system rather than merely refining a decision threshold.

A. Comparison with Literature

The proposed strategy was also compared against literature-based implementations of ResM-FusionNet [17], CTHNet [18], and ICSSN [12]. Because official code was unavailable, the architectures were reconstructed from the published descriptions; consequently, implementation details not fully specified in the original papers may not have been reproduced. All three models were trained on the same 7,092-patch dataset under identical splits, number of epochs, and augmentation settings.

TABLE III

COMPARISON WITH LITERATURE-BASED REIMPLEMENTATIONS.

Model	Impl. F1	Ref. F1	Impl. IoU	Ref. IoU
ResM-FusionNet	0.75	≈ 0.857	0.60	≈ 0.75
CTHNet	0.77	≈ 0.728	0.62	≈ 0.572
ICSSN	0.63	≈ 0.54	0.46	≈ 0.66
Proposed (hybrid)	0.905 ± 0.020	—	—	—

† The reimplemented CTHNet (0.77) slightly exceeds its reference (≈ 0.728), while ResM-FusionNet (0.75) falls below its reference (≈ 0.857). These differences should be interpreted cautiously because the original results were obtained on other datasets and the reported patch-level classification and pixel-level segmentation metrics are not automatically equivalent. The most controlled comparison is therefore Table II, in which the FCDU-Net baseline and hybrid pipeline use the same data, splits, features, and settings.

Under the adopted patch-level evaluation protocol, the proposed hybrid pipeline achieved the highest F1-score among the evaluated implementations. This result does not establish universal superiority over the compared segmentation architectures, which differ in output granularity, architectural objectives, computational requirements, and original evaluation settings. Within this study, the result supports the specific functional decomposition between pixel-level localization and patch-level confirmation.

B. Generalization Experiment

To conduct a preliminary external transfer check beyond the benchmark distribution, an auxiliary evaluation was performed without fine-tuning, threshold retraining, or dataset-specific adaptation. The purpose was to examine the model behavior under a terrain appearance, vegetation pattern, and erosive dynamic not represented in the training data.

The check used an external set of 23 patches (16 positive and 7 negative) drawn from Garanhuns, PE, Brazil, a region with distinct mountainous relief not represented in the training benchmark. The pseudo-NDVI and GLCM features were recomputed using the same procedure adopted for the main experiment, and the trained model was applied with fixed $\tau = 0.50$. Under this setting, the model achieved $F1 = 0.93$. Although this result should be treated as exploratory—the 23-patch set is too small to support reliable F1 confidence intervals—it provides a preliminary indication that the adopted RGB-derived vegetation and texture representation may retain

useful degradation cues outside the main benchmark distribution. Broader cross-region validation is required before drawing conclusions about geographic robustness.

VI. CONCLUSION AND LIMITATIONS

In operational use, the proposed design supports a more cautious alarm policy. A pure segmentation network may return plausible masks over spectrally confusing regions such as exposed ground or recently disturbed agricultural parcels. The second-stage classifier acts as a confirmation mechanism, improving prioritization when many candidate regions must be screened quickly. In a centralized workflow, segmentation masks support manual cartographic review, while the binary confirmation score ranks patches by urgency. This does not eliminate the need for expert supervision, but reduces the false-alarm rate in routine monitoring.

Limitations include: (i) the generalization experiment covers a single external subset and is not a controlled cross-domain evaluation; (ii) the benchmark inherits its own annotation and sampling assumptions, requiring future validation with new sensors, seasons, and terrain types; (iii) practical deployment still depends on data availability, annotation quality, and threshold calibration. A complete operational assessment should additionally quantify parameter count, preprocessing cost, inference latency, throughput, and peak memory consumption.

This paper presented a hybrid pipeline for landslide and soil-scar mapping from satellite imagery. By explicitly separating scar delineation from event confirmation, the system improved F1 from 0.864 to 0.905, reduced inter-run variance, and reached AUC = 0.983 over 30 independent runs. Statistical testing confirms that the gain is significant ($p < 0.001$) and large (Cohen's $d \approx 1.35$). The exploratory evaluation on 23 patches from Garanhuns, PE, provides preliminary evidence motivating broader cross-region validation.

Future work should validate the system on additional geographic regions and sensors, incorporate topographic priors from digital elevation models when available, investigate probability calibration for risk-oriented alert triage, and explore temporal consistency over multi-date image stacks for early-warning applications.

REFERENCES

- [1] United Nations Office for Disaster Risk Reduction (UNDRR) and International Science Council (ISC), "UNDRR-ISC Hazard Information Profiles – 2025 Update: GH0403 Soil Erosion," tech. rep., UNDRR, 2025.
- [2] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, Springer, 2015.
- [3] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2261–2269, 2017.
- [4] X. Chen, W. Li, C.-Y. Hsu, S. T. Arundel, and B. Higman, "Harnessing geospatial artificial intelligence and deep learning for landslide inventory mapping: Advances, challenges, and emerging directions," *Remote Sensing*, vol. 17, no. 11, p. 1856, 2025.
- [5] C.-C. Yu, Y.-D. Chen, H.-Y. Cheng, and C.-L. Jiang, "Semantic segmentation of satellite images for landslide detection using foreground-aware and multi-scale convolutional attention mechanism," *Sensors*, vol. 24, no. 20, p. 6539, 2024.
- [6] K. Wang, D. He, Q. Sun, L. Yi, X. Yuan, and Y. Wang, "A novel network for semantic segmentation of landslide areas in remote sensing images with multi-branch and multi-scale fusion," *Applied Soft Computing*, vol. 158, p. 111542, 2024.
- [7] L. Wu, R. Liu, N. Ju, A. Zhang, J. Gou, G. He, and Y. Lei, "Landslide mapping based on a hybrid CNN-Transformer network and deep transfer learning using remote sensing images with topographic and spectral features," *International Journal of Applied Earth Observation and Geoinformation*, vol. 126, p. 103612, 2024.
- [8] N. Chandra, H. Vaidya, K. Abhinav, and S. R. Meena, "Global landslide mapping using U-Net architecture with diverse backbones across multi-regional and multi-sensor remote sensing datasets," *Frontiers in Earth Science*, vol. 13, p. 1710586, 2026.
- [9] P. Höhn, K. Heidler, R. Behling, and X. X. Zhu, "A spatio-temporal dataset for satellite-based landslide detection," *Scientific Data*, vol. 12, p. 1772, 2025.
- [10] R. M. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-3, no. 6, pp. 610–621, 1973.
- [11] I. Kononenko, E. Šimec, and M. Robnik-Šikonja, "Overcoming the myopia of inductive learning algorithms with ReliefF," *Applied Intelligence*, vol. 7, no. 1, pp. 39–55, 1997.
- [12] Z. Lu, Y. Peng, W. Li, J. Yu, D. Ge, L. Han, and W. Xiang, "An iterative classification and semantic segmentation network for old landslide detection using high-resolution remote sensing images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–13, 2023.
- [13] I. Nasios, "Multi-modal landslide detection from Sentinel-1 SAR and Sentinel-2 optical imagery using multi-encoder vision transformers and ensemble learning," *Remote Sensing Applications: Society and Environment*, vol. 42, p. 102037, 2026.
- [14] Y. Li, "The research on landslide detection in remote sensing images based on improved DeepLabv3+ method," *Scientific Reports*, vol. 15, p. 7957, 2025.
- [15] G. Xu, Y. Wang, L. Wang, L. P. Soares, and C. H. Grohmann, "Feature-based constraint deep CNN method for mapping rainfall-induced landslides in remote regions with mountainous terrain: An application to Brazil," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 2644–2659, 2022.
- [16] S. Ji, D. Yu, C. Shen, W. Li, and Q. Xu, "Landslide detection from an open satellite imagery and digital elevation model dataset using attention boosted convolutional neural networks," *Landslides*, vol. 17, no. 6, pp. 1337–1352, 2020.
- [17] X. Ren, X. Wu, D. Zhai, X. Wang, N. He, and M. Tarif, "ResM-FusionNet for efficient landslide detection algorithm with a hybrid architecture," *Scientific Reports*, vol. 15, no. 1, p. 13080, 2025.
- [18] J. Li, J. Zhang, and Y. Fu, "CTHNet: A CNN-Transformer hybrid network for landslide identification in loess plateau regions using high-resolution remote sensing images," *Sensors*, vol. 25, no. 1, p. 273, 2025.