Addressing data imbalance in urban informal settlement mapping from earth observation using ensemble learning: A case study in Rio de Janeiro

Thomas Hallopeau, Youssef Fouzai, Laurent Demagistri, Joris Guérin, Vanderlei Pascoal de Matos, Renata Gracie, Helen Gurgel, Christovam Barcellos, Nadine Dessay

Science of Remote Sensing 12 (2025) 100273

Contents lists available at ScienceDirect

Science of Remote Sensing

journal homepage: www.sciencedirect.com/journal/science-of-remote-sensing

Full length article

Addressing data imbalance in urban informal settlement mapping from earth observation using ensemble learning: A case study in Rio de Janeiro

Thomas Hallopeau [a,*], Youssef Fouzai [a], Laurent Demagistri [a,d], Joris Guérin [a], Vanderlei Pascoal de Matos [c,d], Renata Gracie [c,d], Helen Gurgel [b,d], Christovam Barcellos [c,d], Nadine Dessay [a,d]

[a] Espace-Dev, French National Research Institute for Sustainable Development (IRD), University of Montpellier, Montpellier, France
[b] Department of Geography, University of Brasília (UnB), Brasília, Brazil
[c] ICICT, Fundação Oswaldo Cruz (Fiocruz), Rio de Janeiro, Brazil
[d] International Joint Laboratory Sentinela (IRD, UnB, Fiocruz), Brazil/France

ARTICLE INFO

Keywords:
Informal settlement mapping
Urban remote sensing
Data imbalance
Ensemble learning

ABSTRACT

Informal settlements pose major challenges for public health, infrastructure, and urban planning due to their high density and unregulated growth. Remote sensing has emerged as a key tool for mapping these areas, but strong class imbalance – where informal settlements represent a small fraction of urban land – remains a critical barrier. Existing methods often rely on simple undersampling, discarding valuable training data from formal residential zones. We propose BALISE, a novel ensemble learning approach that leverages the full extent of available data to improve informal settlement detection from remote sensing. Our framework combines Sentinel-2 multispectral imagery and the Copernicus Digital Elevation Model with auxiliary features derived from OpenStreetMap. We evaluate our method on a use case in Rio de Janeiro, using a spatial cross-validation strategy that rigorously tests generalization across five different urban zones. BALISE improves both F1-score and Kappa coefficient by approximately 2 points over standard undersampling, offering a robust and transferable tool for remote sensing-based urban analysis in fragmented and socially complex environments.

1. Introduction

Informal settlements are often characterized by unregulated growth, insecure tenure, and a lack of basic infrastructure (United Nations High Commissioner for Refugees (UNHCR), 2024). They pose serious challenges to public health, urban resilience, and social equity, particularly in rapidly urbanizing regions (Corburn and Sverdlik, 2019; Kaiser et al., 2025). According to the United Nations, over 1 billion people live in informal settlements worldwide, a number projected to increase significantly (United Nations Statistics Division (UNSD), 2025). Mapping and monitoring these areas is critical for achieving the Sustainable Development Goals (SDGs), particularly Goal 1 (no poverty) and Goal 11.1.1, which calls for ensuring access for all to adequate, safe, and affordable housing and basic services (United Nations Department of Economic and Social Affairs, 2015).

Remote sensing has emerged as a valuable tool for the detection and monitoring of informal settlements (Kuffer et al., 2016b; Raj et al., 2024). Early approaches were based on very high-resolution imagery and object-based image analysis with ontologies (Hofmann et al., 2008; Kohli et al., 2012). Later, machine learning approaches emerged, based on texture and morphology or integrating various spectral indices (Kuffer et al., 2016a; Wurm et al., 2017). More recently, deep learning methods have been widely developed, mainly with medium or very high-resolution satellite images (Persello and Stein, 2017; Wurm et al., 2019; Hafner et al., 2022; da Silva et al., 2025; Büttner et al., 2025). These methods offer good results, but lack interpretability (Raj et al., 2024), and fail to include auxiliary data sources to satellite imagery that could be beneficial from urban land use classification (Shaamala et al., 2025).

Despite these advancements, one major challenge persists, the strong class imbalance between formal and informal urban areas (Friesen et al., 2018). Informal settlements typically occupy a much smaller surface area than formal neighborhoods, creating skewed datasets that degrade model performance (He and Garcia, 2009). Several studies have investigated the effect of class imbalance in slum detection tasks (Stark et al., 2019), but none have proposed methods that actively leverage this imbalance rather than merely mitigating it. Existing techniques typically rely on simple undersampling, which discards a

* Corresponding author.
E-mail addresses: thomas.hallopeau@ird.fr (T. Hallopeau), joris.guerin@ird.fr (J. Guérin).

https://doi.org/10.1016/j.srs.2025.100273
Received 30 June 2025; Received in revised form 4 August 2025; Accepted 13 August 2025
Available online 29 August 2025

significant portion of formal residential samples, resulting in substantial information loss. To address this gap, we propose BALISE (data Balancing for Informal Settlement mapping with Ensembles), a novel ensemble method that maximizes residential data usage by effectively leveraging data imbalance to improve informal settlement detection.

To validate the effectiveness of BALISE, we developed an approach to map Rio's favelas through binary classification of geographic cells, inspired by Owen and Wong (2013). Informal settlements exhibit varying characteristics, as evidenced by the semantic uncertainty surrounding their definition (Gevaert et al., 2019). Although Rio de Janeiro is also located in Latin America, its geography differs markedly from that of Guatemala City. Rio is a coastal city, shaped by mountainous terrain and characterized by a tropical wet climate, in contrast to Guatemala City, which sits at 1500 m above sea level and experiences a subtropical highland climate (Beck et al., 2018). These geographic differences likely influence the characteristics of informal settlements and the selection of relevant indicators for analysis. Hence, we built an improved multivariate approach for Rio de Janeiro, specifically selecting indicators suited to our study area.

To evaluate the performance of our mapping approach, and more specifically to study the impact of BALISE, we introduce a novel robust methodology for validating our approach through spatial cross-validation (Roberts et al., 2017) based on a 5-region partition of Rio de Janeiro. Previous work by Owen and Wong (2013) involved randomly splitting all geographic cells into training and test sets, which poses a key limitation: individual informal settlements could be present in both sets. This overlap leads to overly optimistic estimates of generalization capacity, as the model may rely on spatial correlations rather than truly learning to detect informal settlements in unseen areas. In contrast, our evaluation method ensures that no favela is split between training and test sets, allowing for a more rigorous assessment of the model's ability to generalize to new informal settlements. Our experimental results demonstrate that BALISE improves informal settlement detection accuracy, achieving approximately a 2-point improvement in both F1-score and Kappa coefficient compared to traditional undersampling approaches.

2. Materials and methods

2.1. Study area

Rio de Janeiro is located on the southeastern coast of Brazil, bordering the Atlantic Ocean. It is the capital of the state of Rio de Janeiro and the second-largest city in Brazil. The city is characterized by a complex topography, with coastal plains, steep mountain ranges, and a variety of natural features such as the Tijuca Forest and Sugarloaf Mountain (Pão de Açúcar). The surrounding landscape includes lush rainforests, wetlands, and rocky hills, providing a unique contrast to the urban areas. The city's urban structure spans over 1200 square kilometers and is densely populated, with a highly diverse range of neighborhoods, from affluent coastal areas like Copacabana and Ipanema to the informal settlements, favelas. The city's urbanization is influenced by both historical colonial development and rapid modern growth, with a combination of remaining colonial-era buildings and contemporary infrastructure. The juxtaposition of upscale neighborhoods and favelas, which creates a striking social and economic contrast, is one of Rio's most distinctive characteristics.

Rio de Janeiro's first favelas appeared in the city center in the 1890s, following the abolition of slavery in Brazil. Throughout the 20th century, additional favelas developed in lowland areas along the railway line to the north, while others formed in the mountainous coastal regions of the southern part of the city. Between the 1950s and 1970s, rapid rural-to-urban migration led to increased density, vertical expansion, and the spread of favelas across all available land, replacing green spaces and creating narrow alleys for circulation (Abreu, 1987). Unlike other types of slums and shanty towns — often temporary and built from recycled materials such as wood, scrap metal, and cardboard (see the UrbanGlossary) - favelas have become more permanent. Most houses are now constructed with exposed bricks and feature ceramic or concrete roofs.

Rio de Janeiro is a frequently used case study for informal settlement mapping (Hofmann et al., 2008; Cunha et al., 2024; Hallopeau et al., 2025; da Silva et al., 2025). Previous studies by Wurm and Taubenböck (2018) have highlighted the internal diversity of favela types in Rio de Janeiro, distinguishing at least four morphological and socio-spatial categories which differ in terms of consolidation, regularity, location, density, and building materials. Some communities are compact and long-established, with densely packed masonry buildings, while others are more recent, scattered, and located on the urban periphery. In addition, certain favelas share visual and structural characteristics with nearby formal neighborhoods. These internal complexity and heterogeneity amplifies the methodological challenges for automated mapping models in this city, and makes Rio de Janeiro a valuable and demanding testbed for evaluating the generalizability and robustness of classification models across diverse informal settlement types.

To support our experimental design (see Section 3.2) and to assess the ability of our approach to generalize across different favela types, we divided the study area into 5 distinct zones. We developed a custom zoning scheme that considers the historical territorialization of favelas and the city's geographical features (e.g., forests and hills), while ensuring an even distribution of favelas among the zones (see Fig. 1). The custom zones are described as follows:

1. Historical Core and Bayfront: This coastal zone, bordered by the Tijuca massif, encompasses the downtown district (Centro), the commercial and administrative core of the city, as well as the port area along Guanabara Bay. It includes some of Rio de Janeiro's most iconic neighborhoods as well as some of the city's oldest favelas. Notable favelas in this zone include Jacarezinho, Complexo do Alemão, and Complexo da Maré.
2. North East and Governor's Island: This primarily industrial an residential zone also hosts some of Rio de Janeiro's oldest and largest favelas. Key favelas in this zone include Vila Joaniza and Morro do Dendê on Governor's Island (Ilha do Governador), as well as the communities of Penha and Grotão, Acari and Pedreira.
3. Between the Massifs: Located between the Pedra Branca and Tijuca forest massifs, this zone features a mix of residential and industrial areas. The favelas here are primarily situated on mountain slopes. Notable examples include Rio das Pedras, Rocinha and Vidigal.
4. Northern Expanses: This zone lies to the north of the Pedra Branca massif and includes both well-established favelas, such as Fazenda Coqueiro and Vila do Vintém, and newer communities, such as Nova Cidade.
5. Emerging Western Periphery: Historically a rural area, this zone is still undergoing urbanization and contains some of Rio's most recently established and smallest favelas, such as Cajueiro.

Our ground truth data showed in Fig. 1 is based on official favela boundaries provided by the Brazilian Institute of Geography and Statistics, which defines favelas, also referred to as *subnormal agglomerations*, as densely populated urban areas with irregular land occupation and inadequate public services (Instituto Brasileiro de Geografia e Estatística, 2022). The heterogeneity of the five zones is further illustrated in the Discussion (see Fig. 5), where we compare the distribution of key morphological and environmental features.

2.2. Dataset construction

2.2.1. Grid cell definition

The study area was divided into cells according to an orthogonal grid with a 150 m resolution, following the approach of Owen and

Fig. 1. Study area in Rio de Janeiro, divided into five distinct urban zones for spatial cross-validation.

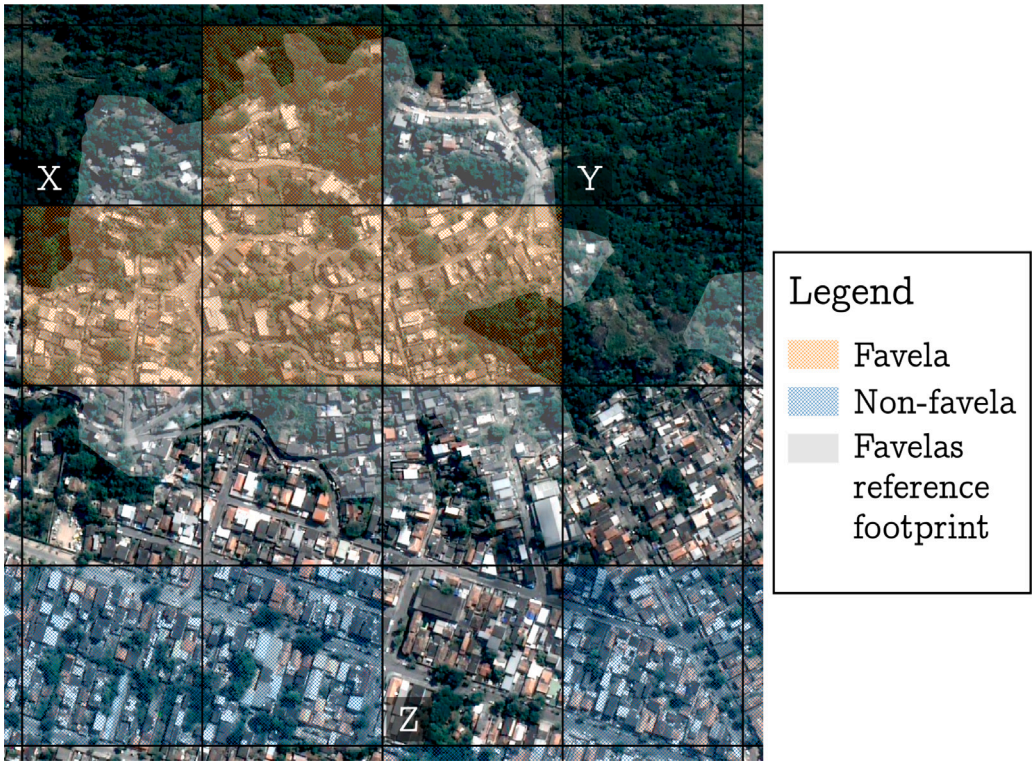

Fig. 2. Example of favela (F) and non-favela (R) cells. Non-retained cells are not shaded.

Wong (2013), as illustrated in Fig. 2. For each geographical cell, the following variables were calculated in order to create a labeled dataset:

- The proportion of surface covered by favelas: computed from the favela outlines provided as a multi-polygon by the Brazilian Institute of Geography and Statistics,
- The proportion of surface covered by vegetation: based on the Normalized Difference Vegetation Index (NDVI), computed from a Sentinel 2 image and thresholded at 0.6 (each pixel with an NDVI greater than or equal to 0.6 was associated with vegetation; this threshold has been adjusted through photo-interpretation),
- The proportion of surface covered by buildings: computed from the Global Human Settlement Layer (GHSL) data of the European Commission.

Based on OpenStreetMap data, the presence or absence of an industrial zone within each cell was also recorded. The primary set of cells

Table 1
Summary of the grid cell features.

Data source	Feature name	Description
Satellite imagery	Vegetation proportion	Proportion of surface covered by vegetation, calculated using NDVI threshold ($\geq$0.6).
	Entropy	Global entropy computed as the mean entropy across the 12 Sentinel-2 bands, normalized between 0 and 1.
Digital elevation model	Slope	Average slope of all pixels in each cell, derived from elevation gradients.
	Profile convexity	Average profile convexity of pixels, representing terrain curvature.
OpenStreetMap	Number of street nodes	Total count of street intersection nodes within each cell.
	Total street length	Total length of streets within each cell.
	Node connectivity statistics	Mean, minimum, and maximum number of streets connected to each node within the cell.

covering the entire city was filtered by removing cells meeting one of the following conditions:

- Proportion of surface covered by buildings strictly less than 50%,
- Proportion of surface covered by vegetation strictly greater than 95%,
- Presence of an industrial or commercial area within the cell. These areas differ significantly from other components of the urban fabric, and therefore, we exclude them from our classification work.

For example, cell Y in Fig. 2 was not retained because it is only composed of vegetation. Cell Z was also not retained because it is associated with the presence of an industrial zone in OpenStreetMap (hangar with a dark gray roof, in the top left corner). While the conditions on vegetation and building coverage may overlap, they are designed to ensure that no non-urban cells are included in the dataset.

2.2.2. Grid cell labeling

A class label was assigned to the remaining cells meeting the following secondary criteria:

- *Favela (F) cells*: At least 90% of the surface area is covered by the favela reference polygon. This threshold was determined based on a sensitivity analysis (see Appendix A).
- *Non-favela (R) cells*, most often corresponding to formal residential zones: No favela coverage (0%).

Some cells meet the initial selection criteria but were not associated with any class and were therefore also eliminated, like cell X in Fig. 2. For this cell, the proportion of surface covered by favelas falls within $]0\%, 90\%[$. This cell was not associated with class R because it is covered by a favela, but it is not sufficiently covered to be associated with class F.

2.2.3. Grid cell features

Inspired by Owen and Wong (2013) and Kohli et al. (2019), and based on our observations, we selected nine indicators tailored to our study area, which were calculated for each cell in the dataset. Some indicators are derived from Sentinel-2 multispectral imagery, others from the Copernicus Digital Elevation Model (DEM), and some from the OpenStreetMap road network (see Table 1).

Features derived from Sentinel-2 imagery. We used a cloud-free Sentinel-2 image from April 2019, downloaded from the Copernicus Data Hub. We first applied a 10 m bilinear resampling to harmonize the 12 bands to the finest spatial resolution. Let p_C denote the total number of pixels in the cell C.

- Proportion of surface covered by vegetation. Favelas are typically very dense with little vegetation compared to planned residential areas. A low vegetation proportion may be therefore be an indicator of favelas. We computed the NDVI from the Sentinel-2 image. Pixels with an NDVI value greater than or equal to 0.6 were classified as vegetation. This threshold was determined through photo-interpretation.

$$\mathrm{V}(C) = \frac{|\{\mathrm{p} \in C \mid \mathrm{NDVI(p)} \geq 0.6\}|}{\mathrm{p_C}} \tag{1}$$

- Entropy. High entropy indicates a high spectral variability within a cell, which can be characteristic of favelas where construction materials are highly diverse (metal sheets, bricks, concrete of different colors, etc.), unlike more uniform neighborhoods. The global entropy for each cell was computed as the mean entropy across the 12 Sentinel-2 normalized bands. The entropy $\mathrm{H}_b(C)$ of a band b for a cell C was calculated using Shannon's entropy formula. First, the band was normalized between 0 and 1. Then the histogram of b pixel values within C was used to compute the probability distribution $P_i^{(b)}$, where $i \in \{0, 1, \ldots, 255\}$ (256 quantization levels). $P_i^{(b)}$ is therefore the probability of pixel values falling in the ith intensity level (IL) for band b, with 256 bins spanning the range $[0, 1]$. The global entropy for the cell C is expressed as:

$$\mathrm{H}(C) = \frac{1}{12}\sum_{b=1}^{12}\mathrm{H}_b(C) = \frac{1}{12}\sum_{b=1}^{12}\left[-\sum_{i=0}^{255}\mathrm{P}_i^{(b)}\log(\mathrm{P}_i^{(b)} + \epsilon)\right] \tag{2}$$

$$\text{with } \mathrm{P}_i^{(b)} = \frac{|\{\mathrm{p} \in C \mid \mathrm{IL(p)} = i\}|}{\mathrm{p_C}} \tag{3}$$

To prevent numerical instability due to logarithm operations on zero, a small constant $\epsilon = 10^{-10}$ is added to the probability values.

Features derived from the copernicus DEM. We used the Copernicus DEM, which provides a global digital surface model at a 30 m resolution. After bilinear resampling of the raster data to 10 m resolution, the following features were computed:

- Slope. Favelas often develop on steep terrains where traditional urbanization is more difficult. A high slope value could therefore be an indicator. For each cell, a slope value was calculated as the average slope across all pixels:

$$\mathrm{Slope}(C) = \frac{1}{\mathrm{p_C}}\sum_{p=1}^{\mathrm{p_C}}\mathrm{S_p} = \frac{1}{\mathrm{p_C}}\sum_{p=1}^{\mathrm{p_C}}\arctan\left[\sqrt{\left(\frac{\partial \mathrm{Z}_p}{\partial \mathrm{X}_p}\right)^2 + \left(\frac{\partial \mathrm{Z}_p}{\partial \mathrm{Y}_p}\right)^2}\right] \tag{4}$$

- Profile convexity. High convexity is associated with hilly and rugged terrains, which are common in favelas built on mountain slopes. For each cell, a profile convexity value was calculated as the average convexity across all pixels. It represents the curvature of the surface along the steepest downhill direction.

Table 2
Distribution of favelas (F), non-favela (R) and non-retained cells across the 5 zones of the dataset.

Zone	1	2	3	4	5	Total
F cells	22%	27%	17%	21%	14%	561 cells
R cells	18%	17%	26%	23%	17%	17 743 cells
Non-retained cells	10%	7%	31%	20%	32%	35 113 cells

$$\mathrm{Conv}(C) = \frac{1}{\mathrm{pc}} \sum_{i=1}^{\mathrm{pc}} \mathrm{PC}_i = \frac{1}{\mathrm{pc}} \sum_{i=1}^{\mathrm{pc}} \left[\frac{\partial^2 \mathrm{Z}_i}{\partial \mathrm{X}_i^2} \cdot \cos^2(\theta_i) + \frac{\partial^2 \mathrm{Z}_i}{\partial \mathrm{Y}_i^2} \cdot \sin^2(\theta_i) \right] \tag{5}$$

$$\text{with } \theta_i = \mathrm{atan2}\left(\frac{\partial \mathrm{Z}_i}{\partial \mathrm{Y}_i}, \frac{\partial \mathrm{Z}_i}{\partial \mathrm{X}_i} \right) \tag{6}$$

Features derived from the OpenStreetMap road network. We extracted Rio de Janeiro's road network from OpenStreetMap using the OSMnx Python library (Boeing, 2024). The road network is divided into two distinct vector files, the street lines and the intersection points (referred to as nodes). We computed the following features from these data:

- Number of street nodes. We computed the total number of street nodes within each cell.
- Total length of streets. We computed the total length of streets within each cell.
- Mean, minimum and maximum connections per street node. Each node has a connectivity value corresponding to the number of streets forming the intersection it represents. We computed the average connectivity value of nodes within each cell as well as the minimum and maximum connectivity value among the nodes within each cell.

Favelas tend to have a more chaotic street network with narrow, unstructured alleyways. A high street length and a high number of nodes can indicate a dense and complex network made of irregular branching structures and sinuous streets. A low average node connectivity, along with a wide range between minimum and maximum connectivity values, may reflect the fact that favela streets have fewer interconnections overall but can still feature highly connected hubs at key junctions. This contrasts with the structured grid pattern of planned neighborhoods, where streets tend to have more uniform connectivity.

2.2.4. Dataset statistics

Table 2 highlights our efforts to balance the number of favela (F) cells across the five zones. Similarly, the number of non-favela (R) cells is relatively consistent across the zones. The significant number of non-retained cells in zones 3, 4, and 5 can largely be attributed to the presence of extensive forested areas. Notably, the dataset contains over 30 times more F cells than R cells, underscoring a stark class imbalance.

2.3. Classification model

We used a Random Forest Classifier (RFC, Breiman (2001)) with the Scikit-learn Python library (Pedregosa et al., 2011). This model performs well on heterogeneous data and helps mitigate the risk of overfitting (Fernández-Delgado et al., 2014), which is particularly suitable given the complex urban sprawl of Rio de Janeiro and the varying characteristics of its favelas (Wurm and Taubenböck, 2018). Moreover, RFC is a non-linear model, making it well-suited for our intra-urban classification task, where some favelas may exhibit characteristics that are sometimes very similar to those of formal neighborhoods.

2.4. BALISE: Data Balancing for Informal Settlement mapping with Ensembles

When designing machine learning models, balancing datasets is essential to prevent the error calculation from being skewed towards the majority class (Carvalho et al., 2025). In urban environments, cells corresponding to informal settlements are frequently under-represented compared to residential areas. For example, in our Rio de Janeiro's dataset, after applying the previously mentioned filters, there are approximately 30 times more class R cells than class F cells in the global dataset (see Section 2.2.4).

2.4.1. Baseline approach

To mitigate the class imbalance, the most straightforward approach is to balance the dataset through random undersampling (He and Garcia, 2009). It consists of constructing independent balanced training and test sets by randomly selecting an equal number of class R and class F cells. All class F cells are used, and an equal number of class R cells are randomly selected. Then, the RFC can be trained on the training set and evaluated on the test set to provide an unbiased estimate of the model's performance. This approach serves as our baseline. With this baseline approach, all the available F cells are divided into the training and test sets, as they represent the minority class. However, only a small fraction of R cells are exploited by the model. In the next section, we present a novel ensemble approach to leverage R cells data to their full potential.

2.4.2. Proposed approach (BALISE)

Ensemble methods in machine learning improve classification accuracy and robustness by combining multiple classifiers (Dietterich, 2000). One well-known approach, bagging (bootstrap aggregating), reduces variability by training models on different random subsets of data and aggregating their predictions. This technique is particularly useful when dealing with high variance, as it stabilizes predictions and mitigates the impact of randomness in data sampling. In the baseline approach, the undersampling process involves a lot of variability in classification performance, depending on the randomly sampled non-favela (R) cells. Our proposed BALISE approach (presented in Fig. 3) leverages this variability and improve prediction stability by eliminating non-systematic errors. We randomly undersample the raw training dataset multiple times to create K balanced subsets, each used to train an individual classifier. During inference, each trained classifier generates predictions. A list of K predictions is thus obtained for each cell in the balanced test set. A prediction of 0 corresponds to the non-favela (R) class, and a prediction of 1 corresponds to the favela (F) class. By considering for each cell the mean prediction value, denoted M, the final predicted class is obtained as follows:

$$\text{Predicted class} = \begin{cases} \mathrm{F} & \text{if } M \geq 0,5 \\ \mathrm{R} & \text{else} \end{cases} \tag{7}$$

3. Experiments

3.1. Evaluation metrics

We selected the following four evaluation metrics to assess the detection quality achieved by our approach: Recall, Precision, F1-score, and Kappa coefficient. Precision measures the model's ability to avoid making errors when predicting positive samples, while Recall measures the model's ability to correctly predict all positive samples. In other words, Precision reflects the proportion of predicted favela (F) cells that really are favela cells, while Recall reflects the proportion of actual favela cells that are classified as such. The F1-score reflects the overall quality of the prediction by balancing Precision and Recall. Additionally, the Kappa score evaluates the agreement between the model's predictions and the ground truth, taking into consideration the proportion of agreement expected by random chance. These four evaluation metrics, commonly used in binary classification tasks, provide a comprehensive view of the model performance.

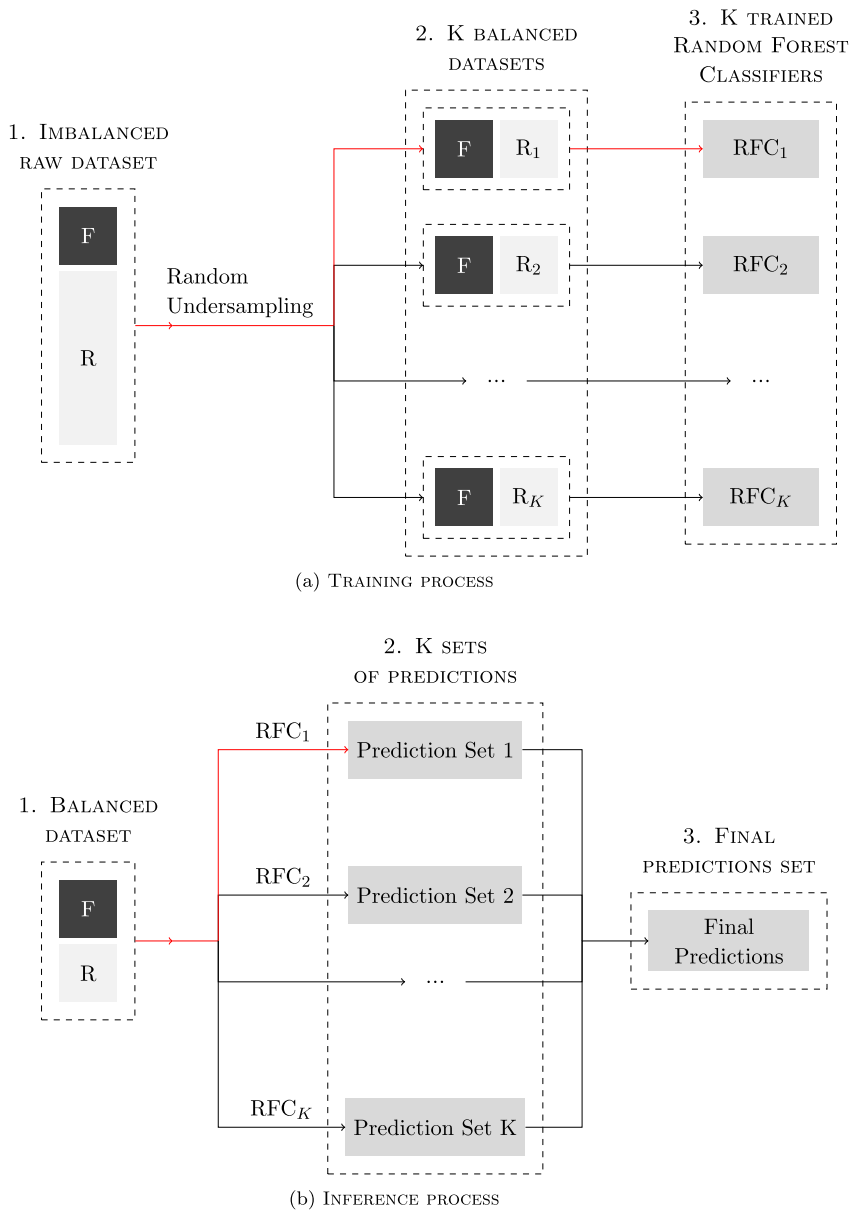

Fig. 3. BALISE (data Balancing for Informal Settlement mapping with Ensembles) is a novel ensemble method that maximizes residential data usage by effectively leveraging data imbalance to improve informal settlement detection. The standard undersampling baseline approach is represented in red.

3.2. Spatial cross-validation

Given the considerable diversity of favelas within the city of Rio de Janeiro, both in terms of morphology and spatial context (Wurm and Taubenböck, 2018), the city provides a challenging and representative testbed for evaluating informal settlement mapping methods. To rigorously assess the generalization ability of our approach across distinct urban conditions, we employed spatial cross-validation (Roberts et al., 2017). The data was partitioned into five spatial subsets corresponding to the five geographical zones introduced in Section 2.1 (see Fig. 1). Successively, four subsets were used for training, and testing was performed on the fifth subset. Each iteration started with a randomly initialized model. This evaluation strategy ensures that the favelas associated to the 5 geographical zones are not split between different training and test sets, thus providing a robust estimate of the model's generalization ability (Beigaitė et al., 2022).

To account for the variability in predictions while assessing our approach, we repeated our evaluation protocol 10 times. This resulted in 50 values for each performance metric: 10 associated with each zone. We then computed the average metric values for each zone and the standard deviation associated to these averages. We applied this evaluation framework for the baseline approach and for the ensemble approach with $K = 100$ (see Section 2.4.2). Finally, we computed the average metrics values over the 5 zones, and the standard deviations, to account for variability of metrics from one area to another.

3.3. Results on balanced test sets

The results are presented in Table 3. Both the baseline and BALISE approaches show significant variation in classification performance across the zones. Zones 1 and 3 (Historical Core and Bayfront, Between the Massifs) are well detected, whereas the remaining three zones

Table 3
Comparison of baseline and ensemble approaches by spatial cross-validation.

Precision		Test zone	Recall	
Undersampling (baseline)	BALISE (ours)		Undersampling (baseline)	BALISE (ours)
0.85 ± 0.02	0.87 ± 0.02	1	0.92 ± 0.01	0.94 ± 0.01
0.83 ± 0.03	0.82 ± 0.03	2	0.79 ± 0.02	0.80 ± 0.01
0.89 ± 0.03	0.89 ± 0.03	3	0.85 ± 0.03	0.85 ± 0.01
0.75 ± 0.02	0.75 ± 0.03	4	0.83 ± 0.04	0.86 ± 0.00
0.77 ± 0.04	0.79 ± 0.04	5	0.69 ± 0.03	0.71 ± 0.00
0.82 ± 0.05	0.82 ± 0.05	Global	0.82 ± 0.08	0.83 ± 0.07

F1-score		Test zone	Kappa	
Undersampling (baseline)	BALISE (ours)		Undersampling (baseline)	BALISE (ours)
0.89 ± 0.01	0.90 ± 0.01	1	0.76 ± 0.02	0.80 ± 0.03
0.81 ± 0.02	0.81 ± 0.01	2	0.62 ± 0.04	0.62 ± 0.03
0.87 ± 0.02	0.87 ± 0.01	3	0.74 ± 0.04	0.74 ± 0.03
0.79 ± 0.02	0.80 ± 0.02	4	0.55 ± 0.04	0.57 ± 0.05
0.73 ± 0.02	0.75 ± 0.02	5	0.48 ± 0.05	0.52 ± 0.04
0.81 ± 0.06	0.83 ± 0.05	Global	0.63 ± 0.11	0.65 ± 0.10

Table 4
Evaluation on imbalanced test set: comparison of baseline and BALISE approaches.

Precision		Test zone	Recall	
Undersampling (baseline)	BALISE (ours)		Undersampling (baseline)	BALISE (ours)
0.15 ± 0.01	0.18 ± 0.00	1	0.89 ± 0.02	0.94 ± 0.00
0.20 ± 0.01	0.19 ± 0.00	2	0.72 ± 0.02	0.80 ± 0.01
0.10 ± 0.00	0.16 ± 0.00	3	0.88 ± 0.02	0.85 ± 0.01
0.07 ± 0.00	0.08 ± 0.00	4	0.72 ± 0.03	0.86 ± 0.00
0.06 ± 0.00	0.09 ± 0.00	5	0.73 ± 0.03	0.72 ± 0.01
0.11 ± 0.05	0.14 ± 0.04	Global	0.79 ± 0.08	0.83 ± 0.07

F1-score		Test zone	Kappa	
Undersampling (baseline)	BALISE (ours)		Undersampling (baseline)	BALISE (ours)
0.26 ± 0.01	0.30 ± 0.00	1	0.20 ± 0.01	0.26 ± 0.00
0.31 ± 0.01	0.30 ± 0.00	2	0.25 ± 0.02	0.24 ± 0.00
0.18 ± 0.01	0.27 ± 0.00	3	0.15 ± 0.01	0.24 ± 0.00
0.12 ± 0.00	0.15 ± 0.00	4	0.08 ± 0.00	0.11 ± 0.00
0.11 ± 0.01	0.16 ± 0.00	5	0.07 ± 0.01	0.12 ± 0.00
0.19 ± 0.08	0.24 ± 0.07	Global	0.15 ± 0.07	0.19 ± 0.07

exhibit a more moderate agreement between predictions and the target, with Zone 5 (Emerging Western Periphery) showing particularly low agreement. The average detection performance across all zones remains good. BALISE had little to no impact on reducing the variability of performance metrics across zones, except for Recall, which showed a very low standard deviation across all zones. However, BALISE improved detection quality across all zones, with the average Kappa coefficient and F1-score increasing by 2 points. It should be noted that in Table 3, the standard deviation values of the Global line are computed over the 5 metrics values above, to account for the variability of metrics from one zone to another (see 3.2). A feature importance analysis is conducted in Appendix B.

3.4. Results on imbalanced test sets

To assess the robustness of BALISE under real-world deployment conditions, we conducted a complementary experiment using imbalanced test sets that preserves the natural class distribution, approximately 30 times more R-class (non-favela) cells than F-class (favela) cells. Table 4 reports the results of this evaluation. Both approaches suffer a sharp drop in precision due to the prevalence of false positives in the highly imbalanced setting. However, BALISE consistently outperforms the baseline across all metrics, particularly in terms of F1-score and Kappa coefficient. Notably, in Zones 1 and 3, BALISE achieves substantial gains in precision and Kappa, while maintaining high recall. These improvements suggest that BALISE provides better stability and generalization, even in the presence of severe class imbalance.

4. Discussion

4.1. Spatial generalization and zone-specific performance

The good results from both the baseline and BALISE demonstrate the relevance of the selected features for our study area. However, spatial cross-validation highlights the difficulty of generalizing from one geographic region to another, as well as from older to newer informal settlements. Our methodology appears more effective in older, denser favelas, characterized by narrow streets and steep terrain. In contrast, newer favelas, still in early stages of urbanization, feature improvised materials, sparse street layouts, and squat houses, making them harder to distinguish from non-favela areas. Additionally, the presence of empty plots and vegetation patches within these recent settlements further complicates classification. This temporal dynamic may explain why classification performance is higher in the older areas of the city (Zones 1, 2, and 3), which have been densely occupied since the 20th century, with streets and avenues shaped by urban regulations. To better illustrate these differences, Fig. 4 includes aerial views of two types of favelas. Santa Maria is an older, consolidated favela located in Zone 1, characterized by dense built-up patterns, narrow alleys, and a steep topography. Jardim da Vitoria is a more recent favela in Zone 5, exhibiting sparse construction, open plots, and more vegetation patches. These visual differences reflect the temporal and morphological diversity of informal settlements across the city, and highlight the challenges of building generalizable classification models. Improving classification performance may require incorporating new features that better capture material heterogeneity or construction stages. Moreover, some of the current features, particularly those derived from the DEM, may be overly specific to the geography of certain zones, limiting the transferability of the learned patterns across space and time. To further quantify the heterogeneity across the five spatial zones used for validation, Fig. 5 presents violin plots (Hintze and Nelson, 1998) for 3 variables, representative of our 3 data sources: vegetation index, derived from Sentinel-2 imagery, average slope, derived from the DEM, and total street length per grid cell, extracted from OSM data. The distributions across zones clearly demonstrate the diversity of environmental and morphological conditions captured in our study area, reinforcing the relevance of spatial cross-validation as a generalization benchmark.

4.2. Impact of class imbalance in realistic test scenarios

Our BALISE approach leverages class imbalance by exploiting the variability in predictions caused by data balancing. Iterating predictions and applying a threshold to their mean eliminates non-systematic errors and improves prediction stability. However, the issue of class imbalance persists during testing. If our approach is applied to an imbalanced dataset, the number of false positives can significantly distort favela detection. Our experiments confirm that while class imbalance at inference time remains a major limitation our approach marks a step forward. The ensemble-based aggregation in BALISE reduces non-systematic errors and helps maintain classification quality under these challenging conditions. A potential improvement could involve rebalancing datasets by first identifying newly urbanized areas, prior to applying our approach. This would correspond to a form of hierarchical classification, in which an initial model detects recent

Fig. 4. Visual comparison of two favela areas in Rio de Janeiro.

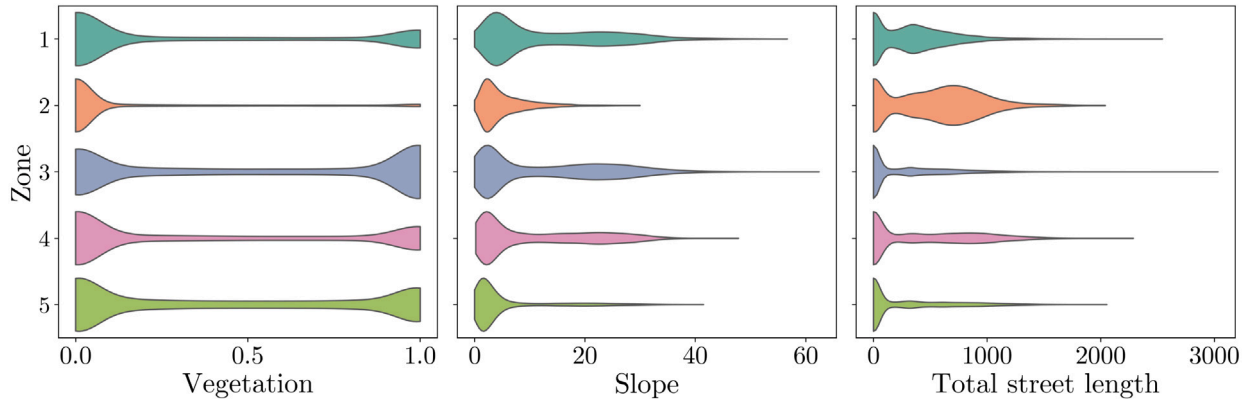

Fig. 5. Distribution of 3 representative features across the 5 zones.

urban expansion zones, where informal settlements are more likely to emerge, followed by the application of BALISE restricted to these areas. Such a strategy would bring the class distribution closer to balance, thereby enhancing detection performance. Hierarchical classification approaches have proven effective in complex intra-urban classification tasks (Shackelford and Davis, 2003).

4.3. Limitations of binary labeling and toward fuzzy classification

Finally, this cell-based classification method introduces challenges related to the use of coverage thresholds for label construction. Only cells with at least 90% of their area covered by reference favela polygons are labeled as class F. This choice helps reduce label noise and improves training performance, as shown in Appendix A. However, it also excludes cells with partial coverage, which may still contain relevant information, especially in transition zones between formal and informal urban fabric. These partially covered cells are excluded from both class F and class R, creating a structural bias in the training data. While a high threshold like 90% leads to clearer labels and better model performance during training, it may result in suboptimal predictions at real inference time, where no such filtering is applied. Urban informal settlements sometimes present gradual transitions rather than sharply defined boundaries (Gevaert et al., 2019), and our current strategy does not fully capture this continuum. Addressing this limitation will likely

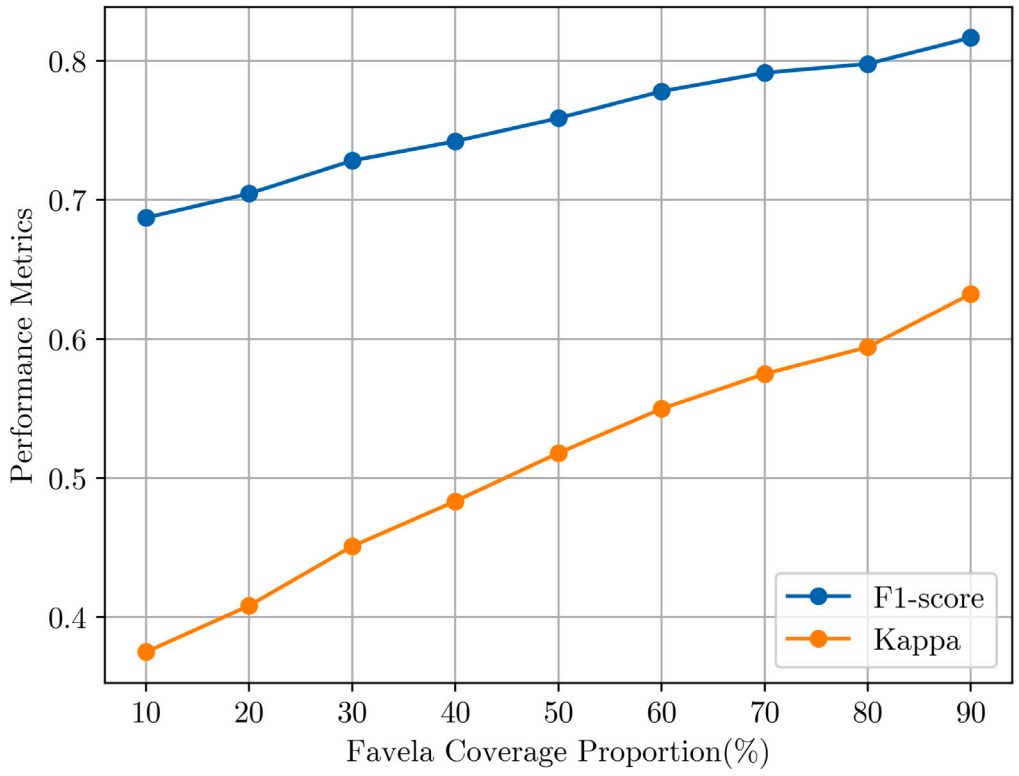

Fig. A.6. Sensibility analysis to the favela coverage proportion.

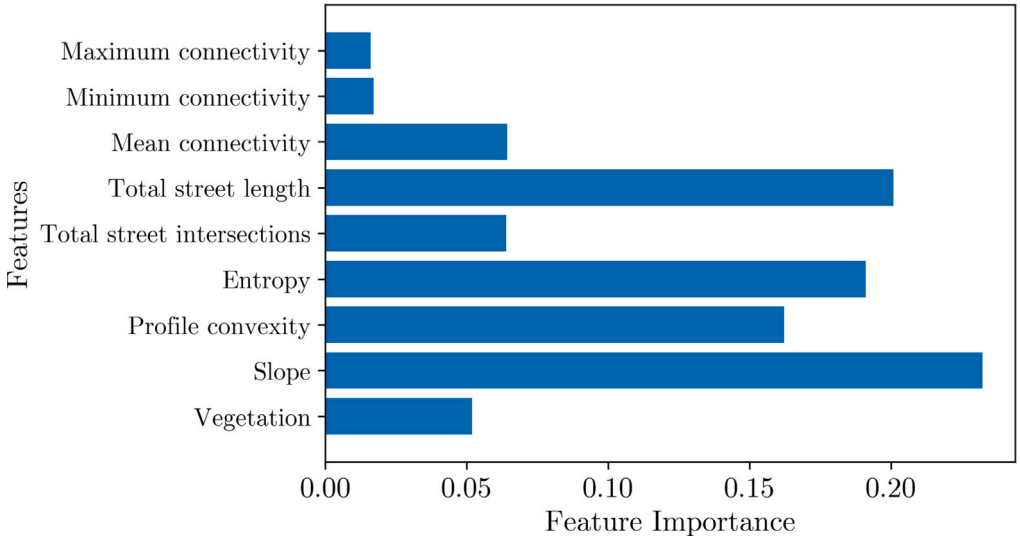

Fig. B.7. Feature importances for the 9 selected features.

require more nuanced label assignment strategies in future work. One promising direction is using probabilistic labeling, where each grid cell is assigned a soft or fuzzy membership score for the favela class based on partial overlaps. Recent research demonstrates that distributing uncertainty across labels improves calibration and generalization in classification tasks using label distributions rather than one-hot encoding (Koller et al., 2024). Moreover, reviews of classification under label noise show how soft label assignments such as probabilistic mixtures or fuzzy clustering can yield more robust models by explicitly modeling uncertainty in annotation (Frenay and Verleysen, 2014).

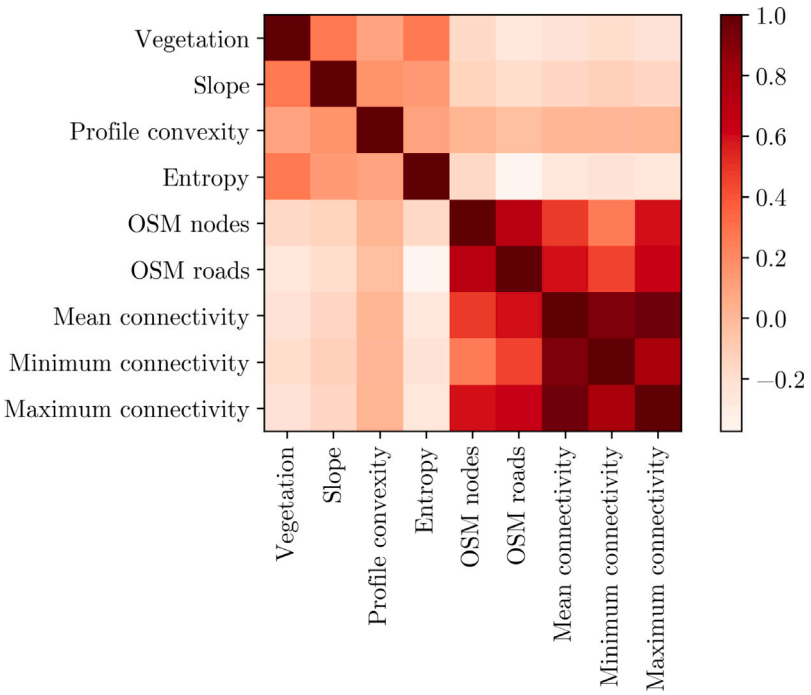

Fig. B.8. Correlation matrix of the 9 selected features.

5. Conclusion

This study highlights the potential of multisource satellite remote sensing combined with spatial machine learning for mapping informal settlements in complex urban landscapes such as Rio de Janeiro. By combining features derived from Sentinel-2 multispectral imagery, Copernicus DEM, and OSM road networks, we developed a classification approach tailored to the unique characteristics of our study area. Our robust evaluation framework, grounded in spatial cross-validation, offers a realistic assessment of the generalization capabilities of our approach and reveals significant challenges in generalizing predictions across geographic regions.

Our ensemble-based approach, called BALISE, proved advantageous to tackle data balancing more efficiently, improving prediction stability, and enhancing overall classification metrics. However, persistent challenges such as class imbalance at inference and the ambiguity introduced by partially covered cells in label assignment remain critical limitations. Addressing these issues through improved feature selection and advanced labeling strategies will be essential to make the methodology more robust and generalizable. Future work should explore additional data sources and features to increase generalization capabilities.

All code associated with this work is available at ensemble-learning-rio. This work was supported by the French Space Agency (CNES) through the CNES-TOSCA committee (tele-epidemiology thematic section).

CRediT authorship contribution statement

Thomas Hallopeau: Writing – original draft, Software, Investigation, Data curation, Conceptualization. **Youssef Fouzai:** Software, Investigation, Data curation. **Laurent Demagistri:** Writing – review & editing, Supervision, Investigation. **Joris Guérin:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Vanderlei Pascoal de Matos:** Writing – review & editing, Validation. **Renata Gracie:** Writing – review & editing, Validation. **Helen Gurgel:** Writing – review & editing, Validation. **Christovam Barcellos:** Writing – review & editing, Validation. **Nadine Dessay:** Writing – review & editing, Validation, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This work was supported by the French Space Agency (CNES) through the CNES-TOSCA committee (tele-epidemiology thematic section).

Appendix A. Sensibility analysis to the minimal favela coverage proportion

While defining our initial labeled dataset, we applied a cell filtering scheme to remove favela (F) cells with a coverage proportion below a fixed threshold (90% in practice, see Section 2.2.2). Here, we analyze how different values of this threshold influence the performance metrics for the baseline approach. For each threshold value, we performed 10 spatial cross-validations and then compared the mean performance metrics, as shown in Fig. A.6. Although fewer favela (F) cells are available for training when the threshold is higher, performance improves.

Appendix B. Feature analysis

We evaluated the importance of our 9 features in the classification task. Fig. B.7 shows the impurity based (or Gini) features importances given by the Scikit-learn Python library (Pedregosa et al., 2011). Features derived from the Copernicus DEM (profile convexity and slope) are significant, as well as entropy (derived from Sentinel-2 imagery) and road length (derived from OSM). The OSM-derived auxiliary features are complementary to the multi-source satellite features, each providing relevant information for the classification task.

Fig. B.8 presents the correlation matrix of our features, constructed using Pearson product-moment correlation coefficients, which were computed with the Numpy Python library (Harris et al., 2020). It can

be observed that the OSM-based features exhibit significant correlations among themselves.

Data availability

Link to code and data available in the manuscript.

References

Abreu, M.D.A., 1987. A Evolução Urbana do Rio de Janeiro. IPLANRIO/Zahar, Rio de Janeiro, p. 147.

Beck, H.E., Zimmermann, N.E., McVicar, T.R., Vergopolan, N., Berg, A., Wood, E.F., 2018. Present and future Köppen–Geiger climate classification maps at 1-Km resolution. Sci. Data 5 (1), 180214. http://dx.doi.org/10.1038/sdata.2018.214.

Beigaitė, R., Mechenich, M., Žliobaitė, I., 2022. Spatial cross-validation for globally distributed data. In: Pascal, P., Ienco, D. (Eds.), Discovery Science. Springer Nature Switzerland, Cham, pp. 127–140. http://dx.doi.org/10.1007/978-3-031-18840-4_10.

Boeing, G., 2024. Modeling and analyzing urban networks and amenities with OSMnx. In: Working paper. URL: https://geoffboeing.com/publications/osmnx-paper/.

Breiman, L., 2001. Random forests. Mach. Learn. 45 (1), 5–32. http://dx.doi.org/10.1023/A:1010933404324.

Büttner, N., Stalder, S., Volpi, M., Suel, E., Harttgen, K., 2025. Large-scale slum mapping in sub-Saharan Africa's major cities: Remote sensing and deep learning reveal strong slum growth in the urban periphery between 2016 and 2022. Habitat Int. 161, 103403. http://dx.doi.org/10.1016/j.habitatint.2025.103403.

Carvalho, M., Pinho, A.J., Brás, S., 2025. Resampling approaches to handle class imbalance: A review from a data perspective. J. Big Data 12 (1), 71. http://dx.doi.org/10.1186/s40537-025-01119-4.

Corburn, J., Sverdlik, A., 2019. Informal settlements and human health. In: Nieuwenhuijsen, M., Khreis, H. (Eds.), Integrating Human Health Into Urban and Transport Planning: A Framework. Springer International Publishing, Cham, pp. 155–171. http://dx.doi.org/10.1007/978-3-319-74983-9_9.

Cunha, H., Diniz da Silva, A., Martins, B., Guedes, B., Nunes, I., Maranhão, M., Conforto, M., 2024. Detection of slums in Rio de Janeiro through satellite images. Dataset Rep. 3, 107–113. http://dx.doi.org/10.58951/dataset.2024.019.

da Silva, J.P., Rodrigues-Jr, J.F., de Albuquerque, J.P., 2025. On the power of CNNs to detect slums in Brazil. Comput. Environ. Urban Syst. 121, 102306. http://dx.doi.org/10.1016/j.compenvurbsys.2025.102306.

Dietterich, T.G., 2000. Ensemble methods in machine learning. In: Multiple Classifier Systems. Springer, Berlin, Heidelberg, pp. 1–15. http://dx.doi.org/10.1007/3-540-45014-9_1.

Fernández-Delgado, M., Cernadas, E., Barro, S., Amorim, D., 2014. Do we need hundreds of classifiers to solve real world classification problems? J. Mach. Learn. Res. 15 (1), 3133–3181.

Frenay, B., Verleysen, M., 2014. Classification in the presence of label noise: a survey. IEEE Trans. Neural Netw. Learn. Syst. 25 (5), 845–869. http://dx.doi.org/10.1109/TNNLS.2013.2292894.

Friesen, J., Taubenböck, H., Wurm, M., Pelz, P.F., 2018. The similar size of slums. Habitat Int. 73, 79–88. http://dx.doi.org/10.1016/j.habitatint.2018.02.002.

Gevaert, C.M., Kohli, D., Kuffer, M., 2019. Challenges of mapping the missing spaces. In: 2019 Joint Urban Remote Sensing Event. JURSE, pp. 1–4. http://dx.doi.org/10.1109/JURSE.2019.8809004.

Hafner, S., Ban, Y., Nascetti, A., 2022. Unsupervised domain adaptation for global urban extraction using Sentinel-1 SAR and Sentinel-2 MSI data. Remote Sens. Environ. 280, 113192. http://dx.doi.org/10.1016/j.rse.2022.113192.

Hallopeau, T., Guérin, J., Demagistri, L., Fouzai, Y., Gracie, R., de Matos, V.P., Gurgel, H., Dessay, N., 2025. Mapping Rio de Janeiro's favelas: General-purpose vs. satellite-specific neural networks. In: XXI SimpÓsio Brasileiro de Sensoriamento Remoto. Galoá, Salvador, Brazil.

Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M.H., Brett, M., Haldane, A., del Río, J.F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., Oliphant, T.E., 2020. Array programming with NumPy. Nat. 585 (7825), 357–362. http://dx.doi.org/10.1038/s41586-020-2649-2.

He, H., Garcia, E.A., 2009. Learning from imbalanced data. IEEE Trans. Knowl. Data Eng. 21 (9), 1263–1284. http://dx.doi.org/10.1109/TKDE.2008.239.

Hintze, J.L., Nelson, R.D., 1998. Violin plots: A box plot-density trace synergism. Amer. Statist. 52 (2), 181–184. http://dx.doi.org/10.1080/00031305.1998.10480559.

Hofmann, P., Strobl, J., Blaschke, T., Kux, H., 2008. Detecting informal settlements from QuickBird data in Rio de Janeiro using an object based approach. In: Blaschke, T., Lang, S., Hay, G.J. (Eds.), Object-Based Image Analysis: Spatial Concepts for Knowledge-Driven Remote Sensing Applications. Springer, Berlin, Heidelberg, pp. 531–553. http://dx.doi.org/10.1007/978-3-540-77058-9_29.

Instituto Brasileiro de Geografia e Estatística, 2022. Favelas e comunidades urbanas. https://www.ibge.gov.br/geociencias/organizacao-do-territorio/tipologias-do-territorio/15788-favelas-e-comunidades-urbanas.html.

Kaiser, Z.R.M.A., Sakil, A.H., Baikady, R., Deb, A., Hossain, M.T., 2025. Building resilience in urban slums: Exploring urban poverty and policy responses amid crises. Discov. Glob. Soc. 3 (1), 8. http://dx.doi.org/10.1007/s44282-025-00142-3.

Kohli, D., Kuffer, M., Gevaert, C.M., 2019. The generic slum ontology: Can a global slum repository be created? In: 2019 Joint Urban Remote Sensing Event. JURSE, pp. 1–4. http://dx.doi.org/10.1109/JURSE.2019.8809034.

Kohli, D., Sliuzas, R., Kerle, N., Stein, A., 2012. An ontology of slums for image-based classification. Comput. Environ. Urban Syst. 36 (2), 154–163. http://dx.doi.org/10.1016/j.compenvurbsys.2011.11.001.

Koller, C., Kauermann, G., Zhu, X.X., 2024. Going beyond one-hot encoding in classification: Can human uncertainty improve model performance in earth observation? IEEE Trans. Geosci. Remote Sens. 62, 1–11. http://dx.doi.org/10.1109/TGRS.2023.3336357.

Kuffer, M., Pfeffer, K., Sliuzas, R.V., 2016b. Slums from space: 15 years of slum mapping using remote sensing. Remote. Sens. 8 (6), 455. http://dx.doi.org/10.3390/rs8060455.

Kuffer, M., Pfeffer, K., Sliuzas, R., Baud, I., 2016a. Extraction of slum areas from VHR imagery using GLCM variance. IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens. 9 (5), 1830–1840. http://dx.doi.org/10.1109/JSTARS.2016.2538563.

Owen, K.K., Wong, D.W., 2013. An approach to differentiate informal settlements using spectral, texture, geomorphology and road accessibility metrics. Appl. Geogr. 38, 107–118. http://dx.doi.org/10.1016/j.apgeog.2012.11.016.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al., 2011. Scikit-learn: Machine learning in Python. J. Mach. Learn. Res. 12 (Oct), 2825–2830.

Persello, C., Stein, A., 2017. Deep fully convolutional networks for the detection of informal settlements in VHR Images. IEEE Geosci. Remote. Sens. Lett. 14 (12), 2325–2329. http://dx.doi.org/10.1109/LGRS.2017.2763738.

Raj, A., Mitra, A., Sinha, M., 2024. Deep learning for slum mapping in remote sensing images: a meta-analysis and review. https://arxiv.org/abs/2406.08031v1.

Roberts, D.R., Bahn, V., Ciuti, S., Boyce, M.S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J.J., Schroeder, B., Thuiller, W., Warton, D.I., Wintle, B.A., Hartig, F., Dormann, C.F., 2017. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. arXiv:11343/216809.

Shaamala, A., Yigitcanlar, T., Nili, A., Nyandega, D., 2025. Machine learning applications for urban geospatial analysis: A review of urban and environmental studies. Cities 165, 106139. http://dx.doi.org/10.1016/j.cities.2025.106139.

Shackelford, A., Davis, C., 2003. A hierarchical fuzzy classification approach for high-resolution multispectral data over urban areas. IEEE Trans. Geosci. Remote Sens. 41 (9), 1920–1932. http://dx.doi.org/10.1109/TGRS.2003.814627.

Stark, T., Wurm, M., Taubenböck, H., Zhu, X.X., 2019. Slum mapping in imbalanced remote sensing datasets using transfer learned deep features. In: 2019 Joint Urban Remote Sensing Event. JURSE, pp. 1–4. http://dx.doi.org/10.1109/JURSE.2019.8808965.

United Nations Department of Economic and Social Affairs, 2015. Transforming our world: The 2030 agenda for sustainable development. https://sdgs.un.org/2030agenda.

United Nations High Commissioner for Refugees (UNHCR), 2024. Informal settlements. https://emergency.unhcr.org/emergency-assistance/settlement-and-shelter/settlement-shelter-interventions/informal-settlements.

United Nations Statistics Division (UNSD), 2025. The sustainable development goals report. https://unstats.un.org/sdgs/report/2025.

Wurm, M., Stark, T., Zhu, X.X., Weigand, M., Taubenböck, H., 2019. Semantic segmentation of slums in satellite images using transfer learning on fully convolutional neural networks. ISPRS J. Photogramm. Remote Sens. 150, 59–69. http://dx.doi.org/10.1016/j.isprsjprs.2019.02.006.

Wurm, M., Taubenböck, H., 2018. Detecting social groups from space – assessment of remote sensing-based mapped morphological slums using income data. Remote. Sens. Lett. 9 (1), 41–50. http://dx.doi.org/10.1080/2150704X.2017.1384586.

Wurm, M., Weigand, M., Schmitt, A., Geiß, C., Taubenböck, H., 2017. Exploitation of textural and morphological image features in Sentinel-2A data for slum mapping. In: 2017 Joint Urban Remote Sensing Event. JURSE, pp. 1–4. http://dx.doi.org/10.1109/JURSE.2017.7924586.