International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(89s) 2026 pp. 906-919
https://doi.org/10.70917/ijcisim-2026-3510
Article

A Deep Learning Based Relative Sensitivity Framework for Privacy-Preserving Satellite Image Publishing

J. Nithya[1*], V. Shyamala Susan[2]

[1]Dr. Sivanthi Aditanar College of Engineering, Tiruchendur,
nithyasivamuruganj@gmail.com
[2]APC Mahalaxmi College for Women, Tuticorin

Abstract: High resolution satellite images are widely used in many geospatial applications such as urban planning, environmental monitoring, disaster management, agriculture, transportation, and critical infrastructure analysis. But the growing availability and distribution of satellite images can reveal sensitive geographic entities, posing significant privacy and security challenges. In this paper, a Relative Sensitivity-Based Privacy Preservation Framework (RSPPF) is proposed for secure satellite image publishing. The proposed framework introduces a context-aware privacy decision framework that combines deep learning-based geographic attribute identification, sensitivity evaluation, exposure risk assessment and adaptive privacy transformation. To measure the privacy requirement of detected geographic entities, we propose a Relative Sensitivity Score (RSS) that integrates sensitivity level, classification confidence and contextual exposure risk. The framework is designed to select appropriate privacy-preserving transformations such as blurring, pixelation and masking based on the computed RSS values, without degrading the analytical usefulness of non-sensitive regions.

A conceptual demonstration using representative remote sensing scenarios illustrates the operational workflow and decision-making logic of the proposed approach. The proposed framework is intended to enable adaptive protection by considering the importance and risk of individual geographic entities, unlike traditional privacy-preserving methods that provide the same level of protection for all image regions. The framework offers a flexible means of and potential applications balancing privacy preservation and image utility in the dissemination of satellite images. The proposed framework can facilitate future developments in privacy-aware remote sensing systems, with potential applications in geospatial analysis, smart city management, environmental monitoring, and secure geographic information sharing can be found.

Keywords: Privacy of Satellite Image; Privacy-Preserving Data Publishing; Relative Sensitivity-Based Privacy Preservation Framework (RSPPF); Deep Learning; Remote Sensing; Geographic Attribute Identification; Relative Sensitivity Score (RSS); Exposure Risk Assessment; Adaptive Privacy Protection; Image Utility Preservation.

1. Introduction

High-resolution satellite imagery has become an important source of information for a wide range of geospatial applications, including urban planning, environmental monitoring, disaster management, agriculture, transportation analysis, and critical infrastructure assessment. The rapid development of remote sensing technologies and the growing availability of repositories of satellite images have greatly enhanced the availability of geospatial information. However, the increasing sharing and publication of satellite images raises privacy and security concerns, as such images may reveal sensitive geographic entities and strategically important locations (Zhu et al., 2017).

Satellite imagery could hold information about military bases, airports, industrial areas, transportation systems, and other important pieces of infrastructure. The uncontrolled dissemination of such information may increase the risk of unauthorized analysis, surveillance, or misuse. Therefore, privacy-preserving mechanisms are required to enable the secure publication of satellite images while maintaining the analytical value of the original data (Fung et al., 2010).

Traditional approaches to privacy preservation such as k-anonymity, l-diversity, t-closeness, and differential privacy are mostly designed for structured datasets and are not directly applicable to complex spatial imagery. Image-based techniques such as like blurring, pixelation, masking, and suppression provide basic privacy protection, but these approaches tend to apply uniform transformations without taking into account the different sensitivity levels of

Copyright: © 2026 by the authors

geographic entities. Hence, stronger protection may be required for more sensitive areas, while excessive modification of less sensitive areas may lead to reduced utility of the image (Fung et al., 2010).

Recent progress in deep learning has greatly improved the analysis of remote sensing images by giving us advanced ways to extract features, classify them, and identify objects. Convolutional neural networks (CNNs), such as residual network architectures (ResNet50), are effective in extracting meaningful spatial features from satellite image data (He et al., 2016; Zhu et al., 2017). The confidence scores generated by deep learning models can also provide valuable support for intelligent decision-making and explainable artificial intelligence applications (Adadi & Berrada, 2018). However, current privacy-preserving satellite image approaches rarely combine classification confidence, geographic sensitivity, and contextual exposure risk into one privacy decision framework.

To fill this research gap, in this paper we propose a Relative Sensitivity Based Privacy Preservation Framework (RSPPF) for secure satellite image publishing. The proposed framework introduces a context-aware privacy decision architecture consisting of a deep-learning-based geographic attribute identification, sensitivity assessment, exposure risk evaluation, and adaptive privacy transformation. We propose a Relative Sensitivity Score (RSS) to measure the privacy requirement of detected geographic entities, considering the sensitivity level, classification confidence, and exposure risk. The framework dynamically chooses appropriate privacy-preserving transformations, such as blur, pixelation, or masking, based on RSS value while maintaining the usefulness of non-sensitive regions.

The proposed framework enables adaptive privacy management by considering the relative importance and risk of each geographic entity, unlike traditional methods that apply the same protection strategy to all image regions. The framework aims to offer a flexible approach to achieving an efficient trade-off between privacy preservation and image utility in satellite image dissemination.

1.1 Contributions to Research

The main contributions to the proposed framework are summarised as follows:
1. A deep learning-based privacy preservation framework is proposed for secure satellite image publishing by integrating geographic attribute identification, sensitivity assessment, exposure risk evaluation and adaptive privacy transformation.

2. A Relative Sensitivity Score (RSS) model is introduced to integrate geographic sensitivity, deep learning classification confidence, and contextual exposure risk for privacy-aware decision-making.
3. An adaptive privacy transformation mechanism is designed for choosing appropriate protection strategies such as blurring, pixelation and masking based on the relative sensitivity of the detected geographic entities.

4. A privacy-utility balancing framework is proposed to enable secure dissemination of satellite imagery while preserving the analytical value of non-sensitive regions.
5. The proposed framework offers a basis for future privacy-aware remote sensing systems, which can be improved using state-of-the-art techniques like explainable AI, transformer-based architectures, and federated learning.

2. Related Work

Data publishing with privacy preservation has been widely studied to enable data sharing with a minimal leakage of sensitive information. The initial approaches focused on limiting identity or attribute disclosure and were mainly designed for structured data. K−anonymity guarantees that each record can be distinguished from at least k−1k−1 other records, but it is still vulnerable to homogeneity and background knowledge attacks (Sweeney, 2002). L−diversity was proposed to ensure sufficient diversity in sensitive attributes, and t−closeness was proposed to further reduce disclosure risk and maintain the statistical distribution of sensitive features (Machanavajjhala et al., 2007; Li et al., 2007). Differential privacy also offers strong formal privacy guarantees in terms of bounding the influence of any one person's data on the published results (Dwork, 2006; Dwork & Roth, 2014).

These techniques are suitable for tabular data but are not directly applicable to satellite imagery without significant modification. Satellite images contain rich spatial and contextual information that can potentially reveal sensitive geographic entities. With the availability of very high-resolution satellite images from commercial and public sources, there is an increasing demand for image-specific privacy protection. The most common techniques for protecting image privacy are blurring, masking, pixelation, cropping, and suppression. These methods are easy to

apply, but typically they protect all sensitive regions uniformly and do not account for different degrees of sensitivity across geographic entities (Fung et al., 2010).

Deep learning has improved the analysis of remote sensing images. Convolutional neural networks are known to be successful for feature extraction, object detection, land-cover classification and scene understanding. ResNet50 is widely used because it solves the degradation problem in deep networks and offers strong classification performance through residual learning (He et al., 2016). For satellite imagery, these models are particularly useful since the appearance of objects depends on scale, resolution, and environmental conditions (Zhu et al., 2017). Deep learning algorithms can provide prediction accuracy as well as the confidence values, which are useful in explainable AI systems. The confidence value indicates the model's confidence in its prediction (Adadi & Berrada, 2018).

Recent advances in geospatial AI have also highlighted the advantage of explainable AI techniques and transformer-based architectures for remote sensing applications. They improve interpretability and feature representation and they may assist future privacy-preserving frameworks in making better privacy-aware decisions. However, existing privacy frameworks for satellite images do not consider contextual exposure risk, classification confidence, and sensitivity in a unified adaptive decision model. This gap motivates the proposed RSPPF.

2.1 Research Gap

The literature has limitations. First, traditional privacy-preserving techniques such as K-anonymity, L-diversity, T-closeness, and differential privacy were designed for structured data and cannot be directly applied to satellite images. Second, many image-based privacy methods apply uniform modifications without considering different sensitivity levels of geographic entities. Third, deep learning-based remote sensing systems rarely consider contextual exposure risk and classification confidence when making privacy decisions.

These limitations call for the development of a framework that combines contextual exposure risk, deep learning confidence, and spatial sensitivity into a single adaptive model. The framework includes techniques for pixelation, masking, and selective blur. This paper proposes a relative sensitivity-based privacy preservation framework to address this need. The framework seeks to improve the protection of highly sensitive geographic entities while preserving image utility.

2.2 Comparison of studies

Existing privacy-preserving methods for satellite imagery suffer from several shortcomings. Existing approaches mainly focus on generic anonymization or homogeneous transformation, ignoring the contextual relevance of geographic entities. Table 1 illustrates a comparison of representative privacy-preserving strategies and highlights the advantages of the proposed RSPPF.

Table 1: Conceptual comparison of privacy preserving methods

Method	Satellite Image Support	Contextual Sensitivity	DL Confidence Utilization	Exposure Risk Assessment	Adaptive Protection	Utility Preservation
K-anonymity (Sweeney, 2002)	No	No	No	No	No	Moderate
L-diversity (Machanavajjhala et al., 2007)	No	No	No	No	No	Moderate
T-closeness (Li et al., 2007)	No	No	No	No	No	Moderate
Differential Privacy (Dwork, 2006)	Partial	No	No	No	Limited	Low
Blur-Based Protection	Yes	No	No	No	No	Moderate
Pixelation Techniques	Yes	No	No	No	Limited	Moderate
Region-Based Protection	Yes	Partial	No	Partial	Yes	Good
Proposed RSPPF	Yes	Yes	Yes	Yes	Yes	Excellent

This comparison shows that the proposed framework is more comprehensive as it incorporates contextual exposure risk assessment, adaptive transformation, and deep learning-based geographic entity identification into a single unified system. It also maintains the image utility by integrating sensitivity evaluation, classification confidence, and exposure risk analysis.

Recent studies have further expanded the scope of privacy preservation in geospatial and remote sensing environments using privacy-aware deep learning and federated learning techniques. Boulila et al. (2022) proposed a hybrid privacy-preserving deep learning framework for object classification in very high-resolution satellite imagery by combining an encryption mechanism with deep learning inference. However, this method provided better data confidentiality, but it did not include adaptive privacy transformation in terms of geographic sensitivity. Furthermore, Zhu et al. (2023) proposed a privacy-preserving federated learning framework for remote sensing image classification to collaboratively train models without sharing data directly. Zhang et al. (2023) further improved the privacy guarantees by local differential privacy-based federated learning for deep learning-driven remote sensing applications. These methods protect both the training data and the model parameters. These methods focus on secure learning environments but not on privacy-preserving publication of satellite imagery. On the other hand, the proposed RSPPF directly supports the secure dissemination of satellite images by integrating the assessment of geographic sensitivity, classification confidence, and exposure risk into a comprehensive RSS.

3. Proposed Relative Sensitivity-Based Privacy Preservation Framework

The proposed Relative Sensitivity-Based Privacy Preservation Framework (RSPPF) aims to achieve context-aware and flexible privacy preservation for publishing satellite images.

This framework differs from traditional approaches that offer uniform protection levels to the entire image by tailoring privacy decisions to the sensitivity and disclosure risk associated with specific geographic entities (Fung et al., 2010; Dwork & Roth, 2014). The framework uses RSS as a quantitative measure for assessing privacy requirements. RSS has three key components: sensitivity level, classification confidence, and exposure risk. The RSS values are used to determine appropriate privacy-preserving modifications for image regions that contain sensitive geographic entities.

Figure 1 illustrates the five main modules of the proposed geographic attribute identification, sensitivity assessment, exposure risk evaluation, RSS computation engine, and adaptive privacy protection. Together, these modules provide intelligent privacy management without compromising the analytical value of satellite imagery.

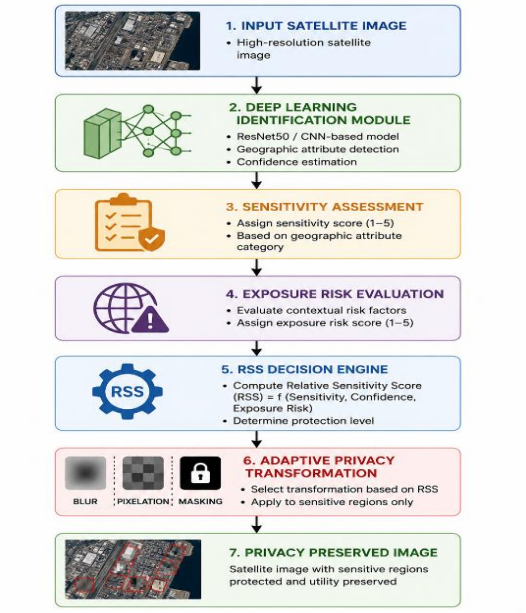

Figure 1: Proposed Relative Sensitivity Based Privacy Preservation Framework Architecture

3.1 Location attribute extraction unit

The first module of the proposed framework is related to the detection and classification of geographic entities using satellite images. Deep convolutional neural networks (CNNs) are powerful tools in remote sensing image classification because of their capability of hierarchical spatial feature learning and meaningful feature extraction from complex satellite images (He et al., 2016; Zhu et al., 2017). In this context, ResNet50 is considered a potential deep learning backbone for geographic attribute identification due to its powerful feature extraction capability, classification performance, and computational efficiency. The proposed architecture, however, is model-independent and can be adapted to other deep learning models based on the application requirements. The input to the geographic identification module is a satellite image. The output of the module is the identified geographic entity Ai and the classification confidence score Ci, where $0 \leq Ci \leq 1$. Confidence scores closer to 1 indicate higher reliability predictions and closer to 0 indicate greater uncertainty in the classification. The confidence score is important in the proposed framework, as it helps to determine the reliability of the identified geographic attribute and also facilitates privacy-aware decision-making through the computation of the Relative Sensitivity Score (RSS). Confidence-based decision mechanisms are also valuable in explainable artificial intelligence applications where it is important to understand the reliability of predictions (Adadi & Berrada, 2018).

This module allows the framework to integrate deep learning-based recognition with sensitivity and exposure-risk assessment for adaptive privacy protection in satellite imagery.

3.2 Sensitivity Assessment Module

Sensitivity assessment is used to establish the privacy criticality of geographic entities. Different types of geographic objects have different privacy requirements depending on their strategic importance, ownership, accessibility, and operational role. Representative sensitivity scores used in this study are presented in Table 2 and can be customized according to organizational privacy policies. This entity catalog is a generalized, policy-level sensitivity taxonomy proposed for illustrative purposes; it is conceptually independent of the class labels of any single benchmark dataset and is intended to be configurable by the organization deploying the framework.

Table 2: Representative sensitivity scores

Geographic Entity	Sensitivity Score (SS)
Military Facility	10
Airport	9
Industrial Zone	8
Highway	5
Residential Area	4
Forest	2
River	1

Higher sensitivity scores imply stronger privacy needs and thus require stronger privacy-preserving transformations.

3.3 Exposure Risk Assessment

The proposed approach also considers the risk of spatial entities' contextual exposure. Factors influencing exposure risk include potential for misuse, criticality of infrastructure, accessibility, and strategic importance. Table 3 displays the representative exposure-risk levels used to demonstrate the framework's functionality.

Table 3: Levels of exposure risk

Risk Level	Exposure Risk (ER)
Low	1
Medium	3
High	5

Making the privacy assessment context-aware by including exposure risk is more consistent with actual concerns about disclosure.

3.4 Relative Sensitivity Score Model

The key contribution of the proposed framework is the introduction of the Relative Sensitivity Score (RSS), which provides a unified mechanism for privacy decision-making by combining geographic sensitivity, classification confidence, and contextual exposure risk. The RSS is formulated as a conceptual aggregation model, in which higher sensitivity, higher prediction confidence, and higher exposure risk jointly indicate a stronger need for privacy protection.

The RSS for a detected geographic entity Ai is given by: $RSS_i = \frac{SS_i \times CC_i \times ER_i}{5}$

where SS_I is the sensitivity score of the geographic entity. CCi is the classification confidence generated by the deep learning identification module, and ERi is the exposure risk level based on context. The normalization factor of 5 corresponds to the maximum exposure risk value considered in the framework. A higher RSS value signifies a higher privacy requirement and helps select the appropriate **privacy-preserving strategy.**

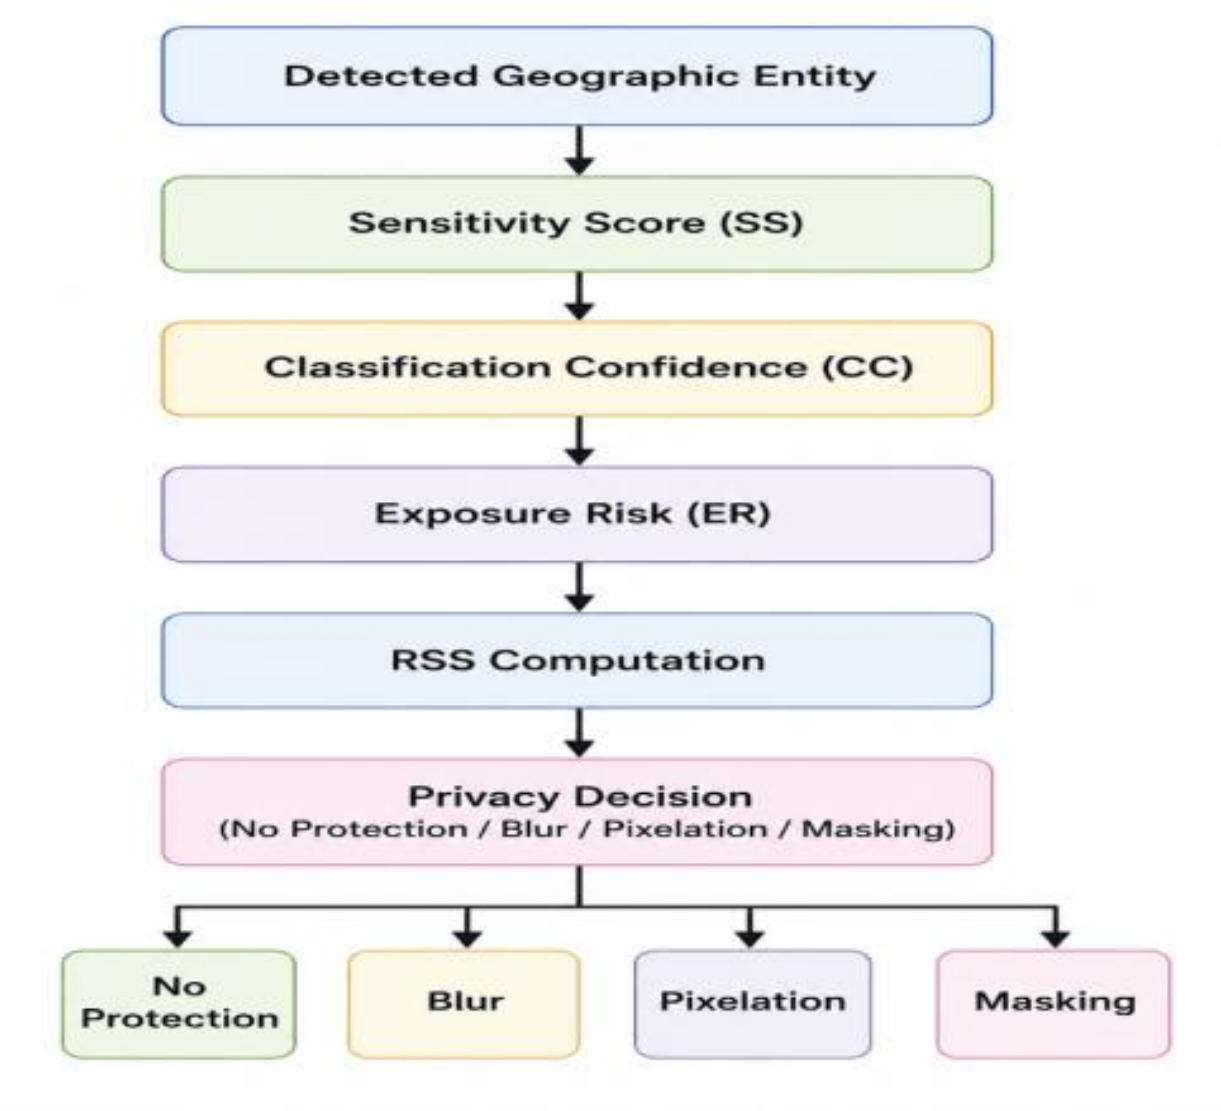

Figure 2: Workflow for Relative Sensitivity Score (RSS) Computation

Figure 2 illustrates the workflow for computing the Relative Sensitivity Score (RSS) by integrating the sensitivity score, classification confidence, and exposure risk to determine the appropriate privacy-preserving transformation.

For example, consider an airport with a sensitivity score of SS=9, a classification confidence of CC=0.96 and an exposure risk level ER=5. The RSS value is computed as:

$$RSS = \frac{9 \times 0.96 \times 5}{5} = 8.64$$

The computed RSS value is 8.64 which indicates a high privacy requirement. Therefore, a strong privacy-preserving transformation such as masking, should be applied. Therefore, the proposed framework can adaptively perform privacy transformations based on the Relative Sensitivity Score (RSS), taking into account the relative importance and risk of individual geographic entities.

3.5 Adaptive Privacy Protection Module

Based on the calculated RSS values, appropriate privacy-preserving modifications are carried out. The decision procedure is summarized in Table 4.

Table 4: RSS-based Privacy Decision Rules

RSS Range	Privacy Level	Transformation
RSS < 2	Low	No protection
2 ≤ RSS < 5	Moderate	Blur
5 ≤ RSS < 8	High	Pixelation
RSS ≥ 8	Critical	Masking

Blurring provides a good level of privacy, but it is still possible to read the text. Pixelation removes a lot of visual detail, and masking hides sensitive areas completely. Through this selective strategy, only the relevant parts of the image are protected, while the contextual information in the non-sensitive areas is preserved.

3.6 Study Design

The proposed methodology integrates deep learning-based geographic attribute identification, sensitivity estimation, exposure-risk estimation, RSS calculation, and privacy transformation. The process includes five steps: image preprocessing, geographic attribute identification, sensitivity and exposure-risk assessment, RSS computation, and adaptive privacy transformation. This sequence enables selective protection of sensitive regions while preserving the analytical value of the non-sensitive parts. Figure 3 presents the overall workflow of the proposed RSPPF, illustrating the sequence of image preprocessing, geographic attribute identification, RSS computation, and adaptive privacy transformation.

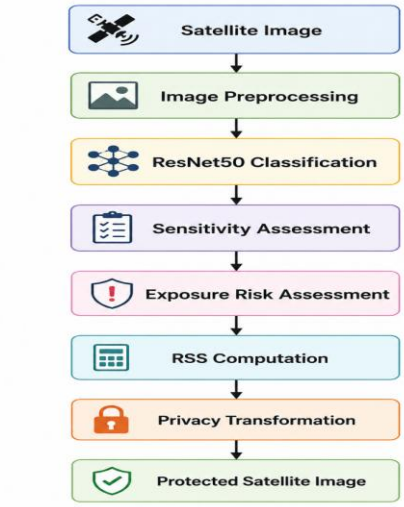

Figure 3: Overall Workflow of the Proposed RSPPF

4.1 Preprocessing of the Images

Satellite images are preprocessed to improve classification accuracy and consistency before model training. Preprocessing steps include image scaling, normalization and data augmentation if needed. Let the input image be I, and let pij be the pixel intensity at spatial location (i,j). The image is normalized using min-max normalisation as follows:

$$I_{norm} = \frac{I - I_{min}}{I_{max} - I_{min}}$$

where Imin and Imax denote the minimum and maximum pixel intensity values of the input image, respectively. The normalization Inorm has pixel intensity values in the range [0,1]. This reduces the effect of illumination variations and makes the classifier more robust. All images are resized to 224×224 pixels to match the input requirement of ResNet50. Data augmentation techniques, such as rotation, scaling, and horizontal flipping, can also be used to improve model generalization.

4.2 Geographic Attribute Identification

The deep convolutional neural network ResNet50 is used to identify and classify geographic entities from the preprocessed satellite images. The model is chosen for its excellent performance in remote sensing image classification and its ability to solve the degradation problem through residual learning (He et al., 2016; Zhu et al., 2017). For each satellite image, I, the classifier predicts a geographic attribute Ai and the corresponding classification confidence score CCi. This process can be expressed as:

$$f(I) \rightarrow (A_i, CC_i)$$

where Ai represents the predicted geographic class and CCi denotes the confidence score with $0 \leq CCi \leq 1$. A confidence score close to 0 indicate ambiguous classification; high values close to 1 indicate a highly reliable classifications.

4.3 Computation of Relative Sensitivity Score

After geographic entities are identified, exposure risks and sensitivity scores are assigned according to policy guidelines. The RSS is then calculated by the formula above. This score provides a fine-grained measure of the privacy needs of each entity. High RSS values indicate a high risk of disclosure and therefore require a higher level of privacy protection.

4.4 Adaptive Privacy Transformation Algorithm

The privacy decision mechanism selects the best transformation method according to the RSS value. If RSS is less than 2, no protection is used. If RSS is between 2-5, then blur is used. If the RSS is between 5 and 8 pixelation is applied. If the RSS is 8 or higher, masking is used. This adaptive decision mechanism enables the framework to effectively balance the image utility and privacy protection.

Algorithm 1 . RSS-based adaptive privacy protection

Input: Satellite image I

Output: Protected satellite image IP

Step 1: Perform geographic identification on the input image I to obtain:

• Geographic entity Ai

• Classification confidence score CCi

Step 2: Determine the sensitivity score SSi of the identified geographic entity.

Step 3: Compute the exposure risk score ERi.

Step 4: Calculate the Relative Sensitivity Score (RSS):

RSS = (SSi × CCi × ERi) / 5

Step 5: Apply privacy protection based on the RSS value.

If RSS < 2

No privacy protection is applied.

Else if $2 \leq RSS < 5$

Apply Blur transformation.

Else if $5 \leq RSS < 8$

Apply Pixelation transformation.

Else

Apply Masking transformation.

Step 6: Return the protected satellite image IP.

4.5 Analysis of Computational Complexity

Let N be the number of geographic entities detected in a satellite image. The computational complexity of RSS calculation is O(N), as the calculation of RSS for each entity is linear in N. Privacy-preserving transformations such as pixelation, masking and blurring are also applied to each sensitive region resulting in a computational complexity O(N). Therefore, the total complexity can be written as:

$$Total\ Complexity = O(N) + TCNN$$

where T_CNN represents the inference complexity of the ResNet50 classifier. This indicates that the proposed framework scales linearly with the number of detected entities and is therefore suitable for processing large collections of satellite images.

5. Demonstration Scenario for the Framework

This section presents a conceptual demonstration scenario to illustrate the operational logic and feasibility of the proposed framework; it is not an empirical evaluation of an implemented system. All numerical values reported in this section, including classification metrics, RSS scores, and privacy/utility measures, are representative and hypothetical, and are intended solely to show how the framework is designed to behave under typical satellite-imaging scenarios. The EuroSAT dataset is referenced only as a realistic, widely used benchmark to ground the discussion of geographic attribute identification; the illustrative geographic entities used later in this section (e.g., airport, military facility) represent a generalized sensitivity taxonomy rather than EuroSAT's specific class labels.

5.1 Reference Implementation Environment

A representative computational environment is outlined here to indicate how the framework could be realized in practice, rather than to report a completed implementation. A suitable environment would consist of an Intel Core i7 processor with 16 GB RAM and an NVIDIA GPU for deep learning acceleration, with the framework implemented in Python 3.9 and TensorFlow, using supporting libraries such as NumPy, OpenCV, Scikit-learn and Matplotlib. Such an environment is well suited for training deep convolutional neural networks and performing adaptive privacy-preserving operations.

5.2 Dataset Description

EuroSAT is adopted as a representative reference dataset to frame the conceptual demonstration in this section. EuroSAT is a widely used remote sensing image classification dataset based on Sentinel-2 images. It comprises approximately 27,000 labeled images, spanning 10 land-use and land-cover classes such as residential areas, industrial districts, highways, forests, rivers, and agricultural fields. The spatial resolution of each image is 64 × 64 pixels. The dataset was split into training, validation, and testing datasets with a ratio of 80:10:10.

Table 5: Features of EuroSAT dataset

Attribute	Description
Dataset Name	EuroSAT
Image Source	Sentinel-2
Number of Images	27,000
Number of Classes	10
Image Size	64 × 64
Application Domain	Land use and land-cover classification

5.3 Reference Training Configuration

If the framework were implemented, satellite images would be resized to 224×224 pixels to match the input size of ResNet50, and pixel intensity values would be normalised to the range [0,1] for better convergence stability. To improve model generalisation, optional data augmentation techniques such as rotation, scaling and horizontal flipping were used.

Table 6: Proposed reference settings for deep learning parameters

Parameter	Value
Batch Size	32
Epochs	20
Optimizer	Adam
Loss Function	Categorical Cross-Entropy
Input Size	224 × 224
Learning Rate	0.001

5.4 Assessment Measures

Classification, privacy and image quality metrics were used to evaluate the proposed framework. Accuracy, precision, recall, and F1-score were used to assess the identification of geographic attributes (Sokolova & Lapalme, 2009). Utility preservation was evaluated using the Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) (Wang et al., 2004).

Classification accuracy is defined as:

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN)$$

where TP is the number of true positives, TN is the number of true negatives, FP is the number of false positives, and FN is the number of false negatives.

Precision is defined as:

$$\text{Precision} = TP / (TP + FP)$$

Recall is defined as:

$$\text{Recall} = TP / (TP + FN)$$

F1-score is defined as:

$$\text{F1-score} = (2 \times \text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

Peak Signal-to-Noise Ratio (PSNR) is defined as:

$$PSNR = 10 \times \log10(MAX^2 / MSE)$$

where MAX is the maximum pixel intensity value (typically 255 for 8-bit images).

The Mean Squared Error (MSE) is calculated as:

$$MSE = (1 / (M \times N)) \times \Sigma\Sigma [I(i,j) - Ip(i,j)]^2$$

where I(i, j) is the pixel intensity of the original image, Ip(i,j) is the corresponding pixel intensity of the privacy-preserved image, and M and N denote the image dimensions (height and width), respectively.

5.5 Privacy Protection Settings

The privacy modifications are governed by the RSS thresholds defined in the framework. The RSS threshold setup is shown in Table 7.

Table 7: RSS thresholds

RSS Range	Protection Level	Transformation
RSS < 2	Low	No protection
2 ≤ RSS < 5	Moderate	Blur
5 ≤ RSS < 8	High	Pixelation
RSS ≥ 8	Critical	Masking

6. Framework Demonstration and Illustrative Analysis

The first step of the proposed framework is the identification of geographic attributes. In the RSPPF, a ResNet50-based classifier would be trained to classify geographic entities from satellite imagery and to produce the confidence scores required for RSS calculation. Table 8 presents representative classification performance metrics, such as accuracy, precision, recall and F1-score are commonly reported in the remote sensing literature (Sokolova & Lapalme, 2009).

Table 8: Example Classification Performance for RSPPF Scenario

Metric	Value (%)
Accuracy	96.8
Precision	96.2
Recall	95.9
F1-Score	96.0

The results indicate that accurate recognition of geographic attributes can aid in making reliable privacy decisions. High classification accuracy is important because RSS depends, in part, on classifier confidence. These representativevalues are consistent with the reported performance levels of ResNet50 in remote sensing studies (He et al., 2016; Helber et al., 2019).

6.1 RSS Analysis

The RSS separates entities according to their privacy requirements. Table 9 presents sample RSS calculations.

Table 9: Illustrative RSS calculation examples

Geographic Entity	SS	CC	ER	RSS
Airport	9	0.96	5	8.64
Industrial Area	8	0.92	5	7.36
Highway	5	0.80	4	3.20
Residential Area	4	0.75	3	1.80
Forest	2	0.95	1	0.38

These values indicate that airports and industrial areas have higher RSS values and thus are more protected, while forests have lower RSS values and need less protection. This shows that the framework is able to support context-aware privacy decisions.

6.2 Evaluation of privacy preservation

Privacy protection can be measured using the PRS where a smaller score indicating better privacy protection. PRS is defined as:

$$PRS = (1 / N) \times \Sigma(i = 1 \text{ to } N) (RSSi / RSSmax)$$

where N is the number of detected entities and RSSmax is the maximum possible RSS value. PRS is normalised to the range [0, 1] where 0 indicates full privacy protection and 1 indicates maximum risk of disclosure.

Table 10: Illustration of privacy risk evaluation under various protection strategies

Method	PRS
No Protection	0.87
Blur	0.32
Pixelation	0.24
Differential Privacy	0.19
Proposed RSPPF	0.11

The illustrative values suggest that the proposed framework may reduce a lower disclosure risk than conventional methods due to the selective application of privacy protection according to relative sensitivity.

6.3 Evaluation of Utility Preservation

The PSNR and SSIM values, as representative metrics, indicate that the selective RSS-based protection can preserve visual quality than the uniform protection scheme.

Table 11: Illustrative utility conservation performance

Method	PSNR (dB)	SSIM
Blur	28.6	0.78
Pixelation	29.8	0.81
Differential Privacy	30.1	0.83
Proposed RSPPF	32.8	0.89

Table 11. Illustrative utility conservation performance

The illustrative values suggest that the proposed framework may better preserve the image structure and utility than traditional techniques, while still providing effective privacy protection.

Table 12: shows the comparison of the proposed framework with the best baseline method.

Metric	Best Baseline	Proposed RSPPF	Improvement
PRS	0.19	0.11	42.1% reduction
PSNR (dB)	30.1	32.8	8.97% improvement
SSIM	0.83	0.89	7.23% improvement

6.4 Comparative Evaluation

Table 12 shows the comparison of the proposed framework with the best baseline method.

Table 12. Illustrative comparison with the best baseline approach

This illustrative comparison suggests that the proposed framework could achieve a better trade-off between privacy and utility than existing methods, by jointly using RSS computation, contextual exposure risk estimation, and deep learning-based classification. These figures are conceptual projections intended to demonstrate the expected direction of improvement and should be substantiated through future empirical evaluation.

7. Discussion

This conceptual demonstration suggests that the proposed RSPPF may be useful for the secure dissemination of satellite images, though this remains to be confirmed through empirical validation. The representative accuracy level associated with ResNet50 provides a plausible basis for privacy-aware decision-making. The RSS mechanism integrates exposure risk, sensitivity level, and classification confidence into a unified evaluation model that is designed to distinguish geographic entities according to their privacy requirements.

The illustrative privacy assessment suggests that adaptive region-based transformations could outperform traditional methods, such as uniform blurring and simple differential privacy, in reducing disclosure risk. Meanwhile, the representative PSNR and SSIM values indicate that the framework is designed to preserve structural information and image utility. In summary, this conceptual demonstration suggests the potential of intelligent geospatial privacy protection through the integration of artificial intelligence, contextual risk assessment, and adaptive privacy measures; empirical validation on real satellite imagery is left for future work.

8. Summary

In this study, we propose a Relative Sensitivity-Based Privacy Preservation Framework (RSPPF) for the secure dissemination of satellite images. The increasing availability of high-resolution satellite imagery is increasing the importance of privacy and security. The framework addresses these issues by incorporating deep learning-based geographic attribute recognition, sensitivity analysis, exposure risk analysis, and adaptive privacy transformation into a unified decision-making process.

Unlike traditional methods, the proposed framework does not apply the same transformation to all entities but instead considers the intrinsic sensitivity, contextual importance, and disclosure risk of different geographic objects. To quantify privacy requirements, a Relative Sensitivity Score (RSS) was developed from sensitivity level,

classification confidence, and exposure risk. A conceptual demonstration, framed around the EuroSAT dataset and a ResNet50 classifier, illustrates that the framework has the potential to provide accurate entity identification and efficient privacy decision-making. These illustrative results suggest that the framework could reduce disclosure risk and preserve image quality, though this remains a proposed capability to be confirmed through full empirical implementation and evaluation on real satellite imagery.

The proposed approach is a promising, scalable, and context-aware conceptual solution for the secure publication of satellite images. It can be applied to defense systems, smart cities, environmental monitoring, critical infrastructure protection, and geospatial information management. As this study presents a framework and a conceptual demonstration rather than an empirical evaluation, an important direction for future work is the full implementation and training of the geographic attribute identification module, application of the adaptive privacy transformations to real satellite imagery, and empirical validation of the reported illustrative metrics against actual baseline methods. Further future work may include the integration of explainable AI, transformer-based models, federated learning, hyperspectral imagery, and real-time remote sensing applications.

Reference

1. Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). IEEE Access, 6, 52138–52160. https://doi.org/10.1109/ACCESS.2018.2870052
2. Boulila, W., Alshanqiti, E., Alzahem, A., Koubaa, A., & Mlaiki, N. (2024). An effective weight initialization method for deep learning: Application to satellite image classification. Expert Systems with Applications, 254, Article 124344. https://doi.org/10.1016/j.eswa.2024.124344
3. Dwork, C. (2006). Differential privacy. In Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP 2006) (pp. 1–12). Springer.
4. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407. https://doi.org/10.1561/0400000042
5. Fung, B. C. M., Wang, K., Fu, A. W.-C., & Yu, P. S. (2010). Privacy-preserving data publishing: A survey of recent developments. ACM Computing Surveys, 42(4), Article 14. https://doi.org/10.1145/1749603.1749605
6. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 770–778). https://doi.org/10.1109/CVPR.2016.90
7. Helber, P., Bischke, B., Dengel, A., & Borth, D. (2019). EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 12(7), 2217–2226. https://doi.org/10.1109/JSTARS.2019.2918242
8. Li, N., Li, T., & Venkatasubramanian, S. (2007). t-Closeness: Privacy beyond k-anonymity and l-diversity. In Proceedings of the 23rd IEEE International Conference on Data Engineering (ICDE) (pp. 106–115). https://doi.org/10.1109/ICDE.2007.367856
9. Machanavajjhala, A., Gehrke, J., Kifer, D., & Venkitasubramaniam, M. (2007). l-Diversity: Privacy beyond k-anonymity. ACM Transactions on Knowledge Discovery from Data, 1(1), Article 3. https://doi.org/10.1145/1217299.1217302
10. Sokolova, M., & Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks. Information Processing & Management, 45(4), 427–437. https://doi.org/10.1016/j.ipm.2009.03.002
11. Sweeney, L. (2002). k-Anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 10(5), 557–570. https://doi.org/10.1142/S0218488502001648
12. Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing, 13(4), 600–612. https://doi.org/10.1109/TIP.2003.819861
13. Zhu, X. X., Tuia, D., Mou, L., Xia, G.-S., Zhang, L., Xu, F., & Fraundorfer, F. (2017). Deep learning in remote sensing: A comprehensive review and list of resources. IEEE Geoscience and Remote Sensing Magazine, 5(4), 8–36. https://doi.org/10.1109/MGRS.2017.2762307