Proceeding Paper

Trainable Noise Model as an Explainable Artificial Intelligence Evaluation Method: Application on Sobol for Remote Sensing Image Segmentation [†]

Hossein Shreim [1,2], Abdul Karim Gizzini [3] and Ali J. Ghandour [2,*]

[1]   Scientific Research Center in Engineering (CRSI) , Faculty of Engineering, Lebanese University,
     Hadath P.O. Box 6573, Lebanon; hossein.shreim@hotmail.com
[2]   National Center for Remote Sensing (CNRS), Beirut P.O. Box 11-8281, Lebanon
[3]   Centre for Digital Systems, IMT Nord Europe, Institut Mines-Télécom, University of Lille, 59000 Lille, France;
     abdul-karim.gizzini@imt-nord-europe.fr
[*]   Correspondence: aghandour@cnrs.edu.lb
[†]   Presented at the 5th International Electronic Conference on Remote Sensing, 7–21 November 2023;
     Available online: https://ecrs2023.sciforum.net/.

Abstract: eXplainable Artificial Intelligence (XAI) has emerged as an essential requirement when dealing with mission-critical applications, ensuring transparency and interpretability of the employed black box AI models. The significance of XAI spans various domains, from healthcare to finance, where understanding the decision-making process of deep learning algorithms is essential. Most AI-based computer vision models are often black boxes; hence, providing the explainability of deep neural networks in image processing is crucial for their wide adoption and deployment in medical image analysis, autonomous driving, and remote sensing applications. Existing XAI methods aim to provide insights about the methodology used by the black-box model in making decisions by highlighting the most relevant regions within the input image that contribute to the model's prediction. Recently, several XAI methods for image classification tasks have been introduced. In contrast, image segmentation has received comparatively less attention in the context of explainability, although it is a fundamental task in computer vision applications, especially in remote sensing. Only some research proposes gradient-based XAI algorithms for image segmentation. This paper adapts the recent gradient-free Sobol XAI method for semantic segmentation. To measure the performance of the Sobol method for segmentation, we propose a quantitative XAI evaluation method based on a learnable noise model. The main objective of this model is to induce noise on the explanation maps, where a higher induced noise signifies low accuracy and vice versa. A benchmark analysis is conducted to evaluate and compare the performances of three XAI methods, Seg-Grad-CAM, Seg-Grad-CAM++ and Seg-Sobol, using the proposed noise-based evaluation technique. This constitutes the first attempt to run and evaluate XAI methods using high-resolution satellite images. Our code is publicly available at GitHub.

Keywords: explainable artificial intelligence (XAI); remote sensing; XAI evaluation; semantic segmentation

1. Introduction

Deep neural networks have achieved remarkable success in various computer vision tasks such as classification, detection, and semantic segmentation. However, they lack interpretability because of their black-box-based processing. Consequently, explainable artificial intelligence (XAI) is crucial for understanding and interpreting the decisions made by any deep learning black box model. Numerous XAI methods have been proposed [1–3] to provide valuable insights into the inner workings of the model and help build trust and confidence in its decision-making process. Generally speaking, XAI methods for image processing tasks provide explanations as saliency maps that highlight the most influential

regions of the input that contribute significantly to the model's prediction. The most recent XAI methods are dedicated to classification tasks, where XAI for segmentation is still largely unexplored. There are two main categories of XAI methods [4]: (*i*)) perturbation-based, where the concept is to perturb input features and record the effect of these changes on model performance without diving into the internal architecture of the considered model, and (*ii*) gradient-based methods where the gradients of the output are calculated with respect to the extracted features or the input via backpropagation and used to estimate attribution scores. We note that internal access to the model architecture is essential in these methods.

Motivated by the fact that evaluating the performance and reliability of XAI methods is crucial to determine their efficiency and reliability for real-world applications, in this work, we propose a quantitative XAI evaluation approach that facilitates a deeper understanding of the performance of any XAI method. The proposed XAI evaluation approach is based on the methodology of the U-Noise model [3] that was initially used as an XAI method. The original U-Noise aims to interpret a pre-trained segmentation model by employing an external model that is responsible for adding noise to the input image without harming the accuracy of the pre-trained model. By doing this, the U-Noise model defines the most important pixels contributing towards the target class segmentation as those assigned low noise weights.

In this context, our proposed evaluation methodology is to feed the XAI saliency map multiplied by the input image to the U-Noise model. Therefore, the U-Noise model serves as a tool for assessing and quantifying the fidelity of XAI methods by adding noise to the important highlighted pixels. Inspired by the recent work proposed in [5], where the gradient-weighted class activation mapping (Grad-CAM) XAI method has been adapted from the classification task to the segmentation task, in this work, we adapted the recently proposed perturbation-based Sobol method [2] to segmentation. Rather than calculating the Sobol indices for a single classification output, as performed in the original work [2], we calculated the Seg-Sobol indices with respect to multiple values of the segmentation output mask considering a specific target class.

To demonstrate the effectiveness of our proposed evaluation technique, we performed experiments on two datasets: Cityscapes dataset [6], which contains a diverse set of semantic urban scene labels, and WHU dataset, which contains satellite images focusing on roof buildings segmentation [7]. Our experimental results demonstrate the ability of the proposed evaluation technique to compare the fidelity of different XAI methods, enabling a more comprehensive and objective assessment of any XAI method. Our code is publicly available at this link Repo (accessed on 5 November 2023).To sum up, the contributions of this paper are threefold:

- We propose a quantitative XAI evaluation approach using a learnable noise model. Our evaluation methodology is based on feeding the saliency map combined with the input image to the noise model. Then, on the basis of the generated noise mask, statistical metrics are computed to quantitatively evaluate the performance of any XAI method.
- We adapt the recently proposed perturbation-based Sobol XAI method from classification to semantic segmentation.
- We benchmark the performance of the adapted Sobol with the gradient-based XAI methods Seg-Grad-CAM and Seg-Grad-CAM++ using the WHU dataset for building footprint segmentation.

2. Proposed Trainable Noise Model XAI Evaluation

2.1. Methodology

The saliency map of the XAI method assumes that the highlighted pixels contribute more to the model decision. To validate whether the highlighted pixels are really relevant to the model decision, XAI evaluation is a must. In this context, our proposed XAI evaluation approach is based on combining the saliency map generated by a specific XAI method with

the original image and then feeding the resultant mask, denoted as the explanation map, to a trained U-Noise model. The U-Noise model is responsible for adding noise to the explanation map. A better XAI method would receive less added noise, as it retains the correct important pixels that contribute to the model decision. Figure 1 illustrates the block diagram of the proposed U-Noise XAI evaluation approach.
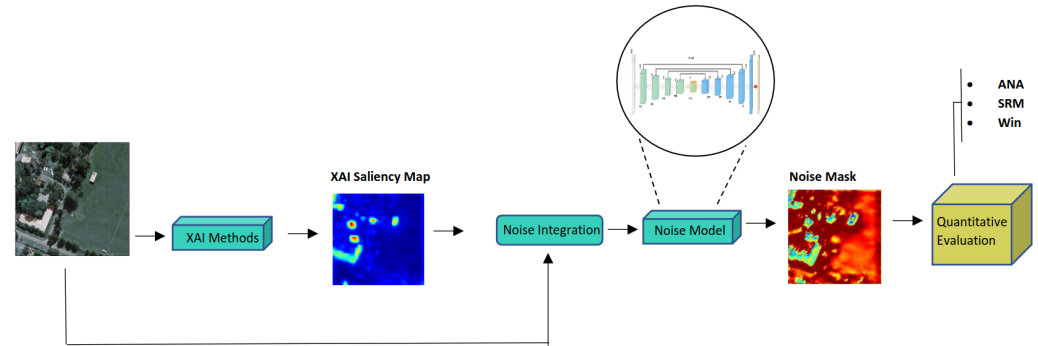
Figure 1. Proposed quantitative evaluation of XAI methods using U-Noise model.

In order to achieve a comprehensive evaluation analysis of XAI, the explanation maps are generated according to the following methodology.

Given an original image I and its corresponding saliency map L_c generated by an XAI method where c denotes the target class, the explanation map I' can be manipulated as follows:

1. **Multiplication**: The original input image is directly multiplied by the saliency map, highlighting regions of the image assumed important by the XAI method, as shown in Equation (1):

$$I'_{\mathrm{mul}} = I \times L_c. \tag{1}$$

2. **Addition**: By adding the saliency map to the original image, we augment the image with importance scores, potentially highlighting regions of interest, as shown in Equation (2):

$$I'_{\mathrm{add}} = I + L_c. \tag{2}$$

3. **Normal sampling with Multiplication**: Similar to the "Normal Sampling with Addition" method, but with multiplication instead of addition. This method emphasizes or de-emphasizes regions based on the importance scores and the sampled noise, as shown in Equation (3):

$$I'_{\mathrm{nsm}} = I \times N(L_c). \tag{3}$$

4. **Normal sampling with Addition**: To introduce variability in the pixels of the explanation map, L_c is sampled from a normal distribution. The resulting sampled values are then added to the original image, as shown in Equation (4):

$$I'_{\mathrm{nsa}} = I + N(L_c). \tag{4}$$

Figure 2 illustrates the proposed explanation map generation methods. We can clearly notice the impact of each method on generating the explanation map. The use of normal sampling with multiplication (Equation (3)) is expected to not provide a reasonable evaluation, as the U-Noise model was not trained on images with such a distribution. For the scope of this work, we will mainly rely on the multiplication method with no sampling introduced in Equation (1).
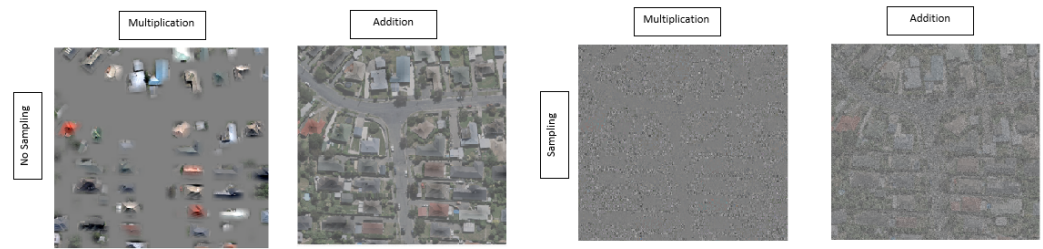
Figure 2. Different integration techniques.

2.2. Metrics

In this work, we propose the following two metrics in order to quantitatively report the results of the U-Noise model:

1. **Average Noise Added (ANA)**: This metric computes the mean value of the output of the U-noise model denoted by $O \in \mathbb{R}^{u \times v}$. A higher ANA indicates that the XAI method introduces more noise to the input image, which means the lower this metric is, the better.

$$ANA = \frac{1}{N} \sum_{(u,v)} O_{i,j}, \quad N = uv. \tag{5}$$

2. **Second raw moment (SRM)**: This metric represents the variance of the noise distribution. A higher SRM suggests that the noise introduced by the trained noise model is spread further away from zero, which also means that the lower this metric is, the better.

$$SRM = \frac{1}{N} \sum_{(u,v)} (N_{i,j})^2 \tag{6}$$

3. Results

This section presents a quantitative evaluation of the U-Noise-based XAI evaluation method using the Cityscapes and WHU datasets.

3.1. Cityscapes

The utility model used was trained to segment the *Road* class of the Cityscapes dataset. It is worth mentioning that to efficiently evaluate the benchmarked XAI methods, a thresholding operation should be applied to the generated noise mask. This is due to the presence of gray regions within the explanation map, as illustrated in Figure 3.
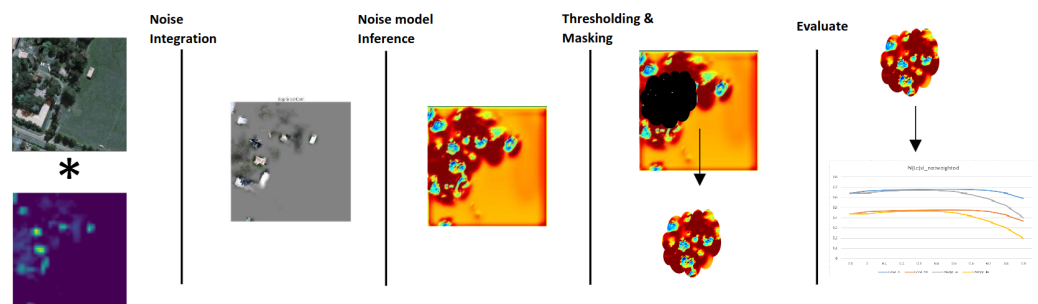

Figure 3. Thresholding operation as an additional step to overcome gray areas effect: We first integrate the saliency map of the XAI method with the original image. Then, we run inference through the noise model and apply thresholding before we calculate the evaluation metrics.

Figure 4 shows the saliency maps of Seg-Grad-CAM [5] and Seg-Grad-CAM++ [8], multiplied by the original image. Figure 5 shows the average and second raw moment of

the added noise mask for the two compared XAI methods, where the x-axis corresponds to the masking threshold and the y-axis represents the metrics ANA and SRM, introduced in Equations (5) and (6). Starting with a threshold of -0.1, which dictates that no thresholding was performed, the evaluation metrics were calculated on the entire noise mask. Seg-Grad-CAM++ shows lower ANA and SRM than Seg-Grad-CAM, indicating that Seg-Grad-CAM++ provides a better explanation of the utility model, which is consistent with the literature.

(**a**)   (**b**)

Figure 4. (**a**) Saliency Maps for Seg-Grad-CAM and (**b**) Saliency Maps for Seg-Grad-CAM++, using Equation (1) (multipliclation with no sampling integration technique) over a sample image from the Cityscapes dataset.
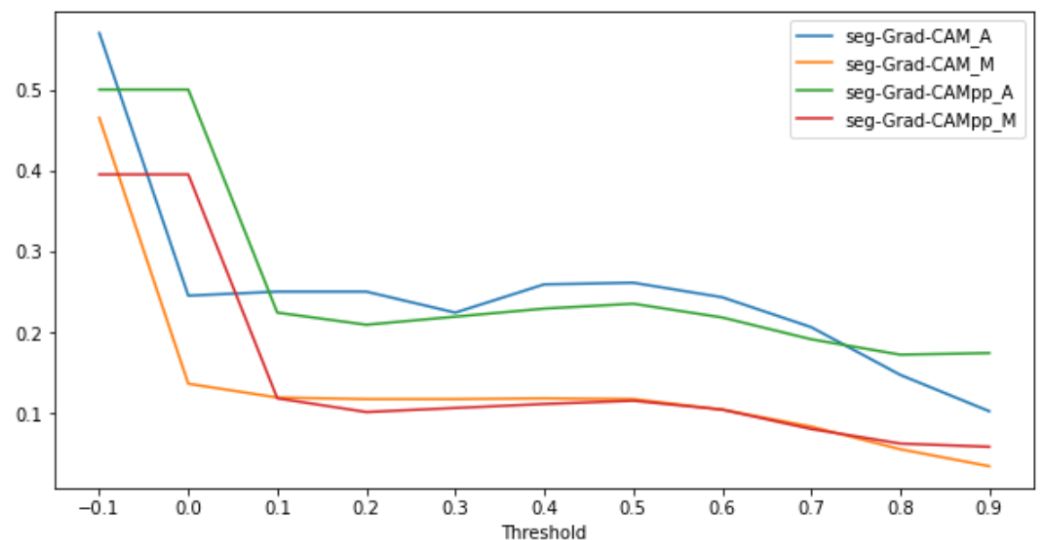
Figure 5. Results for the two benchmarked XAI methods over different threshold values: Seg-Grad-CAM_A and Seg-Grad-CAM++_A are the average noise added on Seg-Grad-CAM and Seg-Grad-CAM++, respectively. Seg-Grad-CAM_M and Seg-Grad-CAM++_M are the second raw moment for noise added on Seg-Grad-CAM and Seg-Grad-CAM++, respectively.

3.2. WHU

Using the WHU dataset, we benchmark two recent gradient-based XAI methods, Seg-Grad-CAM and Seg-Grad-CAM++, in addition to our adapted Seg-Sobol method.

The Sobol XAI method [2] was initially developed for classification models, in which the idea is to perturb the image with several noisy masks and calculate the Sobol indices for each input feature with respect to the output of the classification model, taking into account the applied perturbation. The calculated Sobol indices reflect the impact of the applied perturbations on the prediction of the black-box model. For semantic segmentation, the Sobol indices should be calculated with respect to the summation of target-class pixels within the output probability mask. Sobol has the advantage of not needing to have access

to the model's internal architecture. Figure 6 shows the steps taken to adapt the Sobol method to semantic segmentation, which we refer to as Seg-Sobol.

The Seg-Sobol saliency map highlights the building's surroundings with different intensities as important regions in segmenting building pixels. The results in Figure 7 are qualitatively plausible; the highlighted buildings and regions are thought to be important for the segmentation process.

Figure 8 shows the average and the second raw moment of the added noise mask for the three benchmarked XAI methods, where the x-axis corresponds to the masking threshold and the y-axis represents ANA and SRM metrics, introduced in Equations (5) and (6). Seg-Grad-CAM++ shows the lowest noise average, followed by Seg-Sobol and Seg-Grad-CAM. This is also the case for the second raw moment metric. The same results are also observed for the threshold value of zero. For threshold = 0.1, Seg-Grad-CAM receives the lowest noise average and thus outperforms the other two methods. Future work will investigate means to improve the Seg-Sobol explanation outcome for earth observation segmentation use cases.
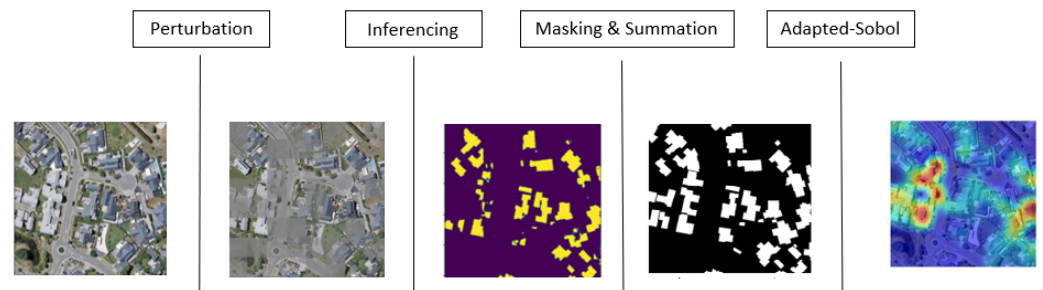
Figure 6. Seg-Sobol: Adaptation of Sobol method from classification to segmentation.

Figure 7. Seg-Sobol results with grid size = 11 using sample from the WHU dataset.
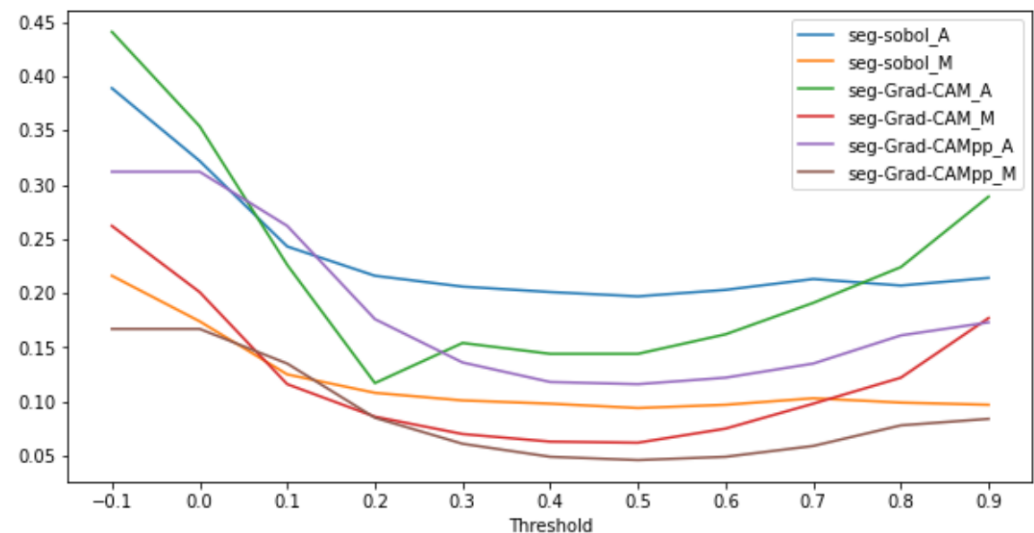
Figure 8. Quantitative metrics results for the benchmarked XAI methods using Equation (1) (multiplcation with no sampling) over different threshold values. Seg-Sobol_A, Seg-Grad-CAM_A and Seg-Grad-CAM++_A are the average noise added on Seg-Sobol, Seg-Grad-CAM, and Seg-Grad-CAM++, respectively. Seg-Sobol_M, Seg-Grad-CAM_M and Seg-Grad-CAM++_M are the second raw moment for noise added on Seg-Sobol, Seg-Grad-CAM, and Seg-Grad-CAM++, respectively.

4. Conclusions

In our research, we successfully adapted the Sobol XAI method to better understand image segmentation tasks. To evaluate its effectiveness, we introduced a unique noise model technique. When we compare Seg-Sobol with other methods such as Seg-Grad-CAM and Seg-Grad-CAM++, it showed promising results. Furthermore, using high-resolution satellite images for our tests was a new and important step. These findings are crucial because they make AI-driven earth observation applications more transparent and easier to understand, paving the way for safer and more reliable real-world applications.

References

1. Jung, H.; Oh, Y. Towards better explanations of class activation mapping. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, BC, Canada, 11–17 October 2021 ; pp. 1336–1344.
2. Fel, T.; Cadène, R.; Chalvidal, M.; Cord, M.; Vigouroux, D.; Serre, T. Look at the variance! efficient black-box explanations with sobol-based sensitivity analysis. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 26005–26014.
3. Koker, T.; Mireshghallah, F.; Titcombe, T.; Kaissis, G. U-noise: Learnable noise masks for interpretable image segmentation. In Proceedings of the 2021 IEEE International Conference on Image Processing (ICIP), Anchorage, Alaska, USA, 19–22 September 2021; IEEE: New York, NY, USA, 2021; pp. 394–398.
4. Nielsen, I.E.; Dera, D.; Rasool, G.; Ramachandran, R.P.; Bouaynaya, N.C. Robust explainability: A tutorial on gradient-based attribution methods for deep neural networks. *IEEE Signal Process. Mag.* **2022**, *39*, 73–84. [CrossRef]

5. Vinogradova, K.; Dibrov, A.; Myers, G. Towards interpretable semantic segmentation via gradient-weighted class activation mapping (student abstract). In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 13943–13944.

6. Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; Schiele, B. The Cityscapes Dataset for Semantic Urban Scene Understanding. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, Nevada, USA, 27–30 June 2016.

7. Nasrallah, H.; Samhat, A.E.; Shi, Y.; Zhu, X.X.; Faour, G.; Ghandour, A.J. Lebanon Solar Rooftop Potential Assessment Using Buildings Segmentation From Aerial Images. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.* **2022**, *15*, 4909–4918. [CrossRef]

8. Chattopadhay, A.; Sarkar, A.; Howlader, P.; Balasubramanian, V.N. Grad-CAM++: Generalized Gradient-Based Visual Explanations for Deep Convolutional Networks. In Proceedings of the 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), Lake Tahoe, NV, USA, 12–15 March 2018; pp. 839–847. [CrossRef]