

Reproducibility Study of 'SLICE: Stabilized LIME for Consistent Explanations for Image Classification'

Anonymous authors

Paper under double-blind review

Abstract

This paper presents a reproducibility study of SLICE: Stabilized LIME for Consistent Explanations for Image Classification by Bora et al. (2024). SLICE enhances LIME by incorporating Sign Entropy-based Feature Elimination (SEFE) to remove unstable superpixels and an adaptive perturbation strategy using Gaussian blur to improve consistency in feature importance rankings. The original work claims that SLICE significantly improves explanation stability and fidelity. Our study systematically verifies these claims through extensive experimentation using the Oxford-IIIT Pets, PASCAL VOC, and MS COCO datasets. Our results confirm that SLICE achieves higher consistency than LIME, supporting its ability to reduce instability. However, our fidelity analysis challenges the claim of superior performance, as LIME often achieves higher Ground Truth Overlap (GTO) scores, indicating stronger alignment with object segmentations. To further investigate fidelity, we introduce an alternative AOPC evaluation to ensure a fair comparison across methods. Additionally, we propose GRID-LIME, a structured grid-based alternative to LIME, which improves stability while maintaining computational efficiency. Our findings highlight trade-offs in post-hoc explainability methods and emphasize the need for fairer fidelity evaluations. Our implementation is publicly available at our [GitHub repository](#).

1 Introduction

Explainability in AI has become increasingly important as deep learning models achieve unprecedented success Grobrügge et al. (2024). However, interpreting these models remains a challenging problem, particularly when their decisions are not inherently transparent.

To address this challenge, post-hoc explainability methods aim to explain predictions of trained AI models without requiring access to their internal structures. Local Interpretable Model-agnostic Explanations (LIME) (Ribeiro et al., 2016) is one such post-hoc method. LIME approximates a black-box model’s decision function locally by generating perturbed samples around a given test point and fitting a linear regression model to these perturbations. While LIME generates locally faithful explanations, it suffers from inconsistencies across multiple runs (Gosiewska & Biecek, 2019; Lee & Lee, 2023; Li et al., 2023; Zafar & Khan, 2019; Zhang et al., 2019; Zhao et al., 2021; Zhou et al., 2021).

Stabilized LIME for Consistent Explanations for Image Classification (SLICE) (Bora et al., 2024) enhances LIME by incorporating adaptive Gaussian blurring and sign entropy-based feature elimination to improve explanation consistency. In this reproducibility study, we aim to validate the key claims of SLICE by replicating its primary experiments and results. The implementation is available at ¹.

Beyond verifying the original claims, we propose GRID-LIME, an alternative segmentation approach that enhances LIME’s stability while maintaining efficiency. Unlike LIME, which relies on color-based superpixel segmentations, GRID-LIME employs a structured grid-based segmentation technique incorporating model predictions. Additionally, we introduce the Ground Truth Overlap (GTO) metric, which provides a direct comparison between explainability maps and ground-truth segmentations, addressing the limitations of prior

¹https://anonymous.4open.science/r/slice_reproducibility-8DC6/

evaluation methods. Our findings suggest that while SLICE improves explanation consistency, its claimed fidelity improvements remain debatable.

In this work, we first provide a detailed outline of the specific claims made by Bora et al. (2024) by reproducing the experiments presented in their original paper. We then introduce GRID-LIME, a variation of LIME that leverages a structured grid-based segmentation approach and relies on the black-box model’s predictions to improve stability in explanations without undermining its explainability. Additionally, we conduct a more detailed analysis of SLICE to investigate the relationship between its consistency-improving techniques and the quality of generated explanations. Following this, we extend the evaluation of SLICE by applying it to an additional dataset, providing further insights into its generalizability and performance across diverse settings.

2 Scope of reproducibility

To address the stability issues in LIME-based explanations, Bora et al. (2024) proposed SLICE, a method that enhances LIME by incorporating a novel feature elimination strategy and an adaptive perturbation technique. Specifically, SLICE utilizes Sign Entropy-based Feature Elimination (SEFE) to remove superpixels with high sign entropy, ensuring greater consistency in explanations. Additionally, SLICE introduces an adaptive Gaussian Blur mechanism that dynamically selects the optimal perturbation level, reducing variance in feature importance rankings and improving the overall reliability of explanations.

The focus of our work is on reproducing the following claims made by Bora et al. (2024):

- **Claim 1:** LIME suffers from inconsistencies in its explanations, primarily due to variance in sign and importance ranking of superpixels. These inconsistencies stem from the perturbation method used by LIME.
- **Claim 2:** The Sign Entropy-based Feature Elimination (SEFE) method improves stability in explanations by identifying and eliminating spurious superpixels.
- **Claim 3:** The introduction of Gaussian Blur with an adaptively selected hyperparameter sigma enhances the perturbation method, leading to greater consistency in explanations.
- **Claim 4:** Explanations generated by SLICE demonstrate significantly better fidelity compared to those produced by LIME.

3 Post-Hoc Explainability Methods

Post-hoc explainability methods aim to interpret complex AI models without requiring access to their internal structures. Local Interpretable Model-agnostic Explanations (LIME) (Ribeiro et al., 2016) is a widely used method that approximates a model’s decision function locally by generating perturbed samples around a given instance and fitting a linear surrogate model. However, LIME suffers from inconsistencies across different runs, leading to instability in feature importance rankings and sign variations.

Several extensions have been proposed to improve the stability of LIME. **DLIME** (Rashid et al., 2024) employs hierarchical clustering to select perturbed samples from the cluster nearest to the instance of interest, ensuring greater locality adherence. **S-LIME** (Shi et al., 2020) leverages a hypothesis testing framework based on the Central Limit Theorem to determine the number of samples required for stable explanations. **ALIME** (Knab et al., 2025) introduces an autoencoder-based weighting function to improve coefficient stability by refining sample selection. **BayLIME** (Dehghani et al., 2024) formulates explanations as a Bayesian-weighted sum of prior knowledge and new observations, reducing variance in feature importance.

More recent methods focus on enhancing LIME’s applicability to image-based models. **Beyond Pixels** (Knab et al., 2025) integrates hierarchical feature representations with segmentation foundation models to improve visual explanations. **SMILE** (Dehghani et al., 2024) employs statistical distances based on empirical cumulative distribution functions to enhance interpretability in instruction-based image editing tasks. These advancements contribute to improving LIME’s robustness across diverse applications.

As our study primarily investigates post-hoc explanations for image classification, we focus on evaluating **SLICE** (Bora et al., 2024), which introduces an adaptive Gaussian perturbation method and a feature elimination strategy to improve stability. SLICE to standard LIME, we provide a comprehensive assessment of its effectiveness in producing consistent and reliable explanations.

4 Methodology

The objective of this reproducibility study is to verify the stability and consistency of SLICE which uses Adaptive Gaussian Blur and Sign Entropy based Feature Elimination (SEFE). For re-implementing the code in PyTorch, we referred to the SLICE code that was made publicly available by the authors at ².

It should be noted that the author’s code is not referenced directly in the original paper. However, since the link to the author’s code was provided to us by the authors, we decided to use this implementation as our reference. Furthermore, we added more interpretable intermediate images and a new interpretability metric to enhance the usability of SLICE. The code for this reproducibility study is publicly available on Github ³

4.1 Model descriptions

4.1.1 LIME

Local Interpretable Model-agnostic Explanations (LIME) is a widely used post-hoc interpretability method that approximates the decision boundary of black-box models using a local surrogate model. LIME operates through three key steps: perturbation, weighting, and surrogate model training. Given a black-box model f and an input instance x , LIME generates perturbed images z by masking superpixels. A weighting function π_x assigns importance to these perturbed samples based on their similarity to the original input. Subsequently, a linear surrogate model g is trained on these perturbed samples to approximate f locally. The objective function for LIME is given by:

$$\mathcal{L}(f, g, \pi_x) = \sum_{z, z' \in \mathcal{Z}} \pi_x(z) (f(z) - g(z'))^2 \quad (1)$$

While LIME effectively provides locally faithful explanations, it suffers from instability in feature importance rankings and sign fluctuations across multiple runs, limiting its reliability.

4.1.2 SLICE

SLICE (Stabilized LIME for Consistent Explanations) is designed to mitigate the instability of LIME by introducing two key components: Adaptive Gaussian Blur and Sign Entropy-based Feature Elimination (SEFE). An overview of SLICE is illustrated in Figure 1.

The first component, Adaptive Gaussian Blur, applies a range of perturbation strengths (σ values) to superpixels using Gaussian blurring. Instead of replacing superpixels with a fixed value (as in LIME), SLICE adaptively selects an optimal σ value based on entropy, ensuring perturbations are closer to the original data distribution.

The second component, SEFE, eliminates superpixels that exhibit high sign entropy, meaning their attribution flips between positive and negative across multiple iterations. This ensures greater stability in feature importance rankings and improves the consistency of explanations.

The SLICE algorithm iteratively generates N perturbed samples and trains M Ridge Regression models. SEFE identifies and removes unstable superpixels, while the final ranking of superpixels is based on the learned coefficients from the regression models. Training continues until a predefined number of iterations or a stability tolerance threshold is reached.

²<https://github.com/rebathip/SLICE-Stabilized-LIME-for-Consistent-Explanations-for-Image-Classification/tree/main>

³https://anonymous.4open.science/r/slice_reproducibility-8DC6

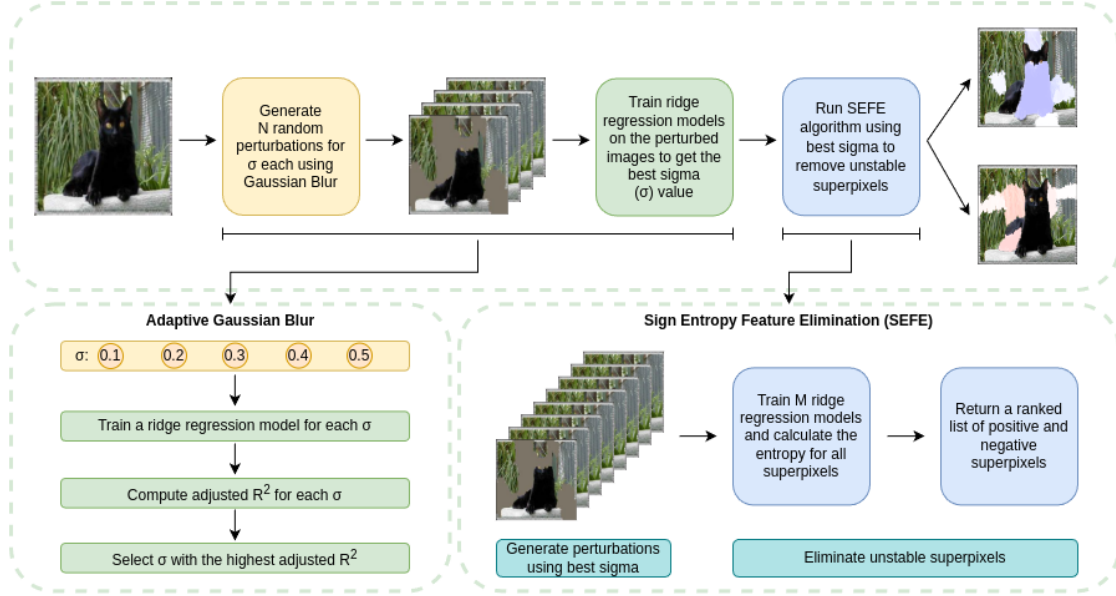


Figure 1: Overview of the SLICE framework. SLICE improves LIME’s stability by incorporating Adaptive Gaussian Blur and Sign Entropy-based Feature Elimination (SEFE).

4.1.3 GRID-LIME

LIME and SLICE both rely on the `quickshift` algorithm (Vedaldi & Soatto, 2008) to generate image superpixels. However, since `quickshift` generates segments based on color-space, we propose GRID-LIME, a variation of LIME that replaces `quickshift` with a structured grid-based segmentation approach that relies on model predictions to create segments of a particular grid-size.

As shown in Figure 2, GRID-LIME divides an input image into uniform grid cells of varying sizes, treating each grid cell as a superpixel. For each grid size, perturbed images are generated by setting all pixels within each grid cell to zero. The black-box model’s predictions on these perturbed images are analyzed, and the optimal grid size is selected based on the coefficient of variation (Zwillinger & Kokoska, 1999) computed from the predicted probabilities. This ensures a balance between segmentation granularity and stability.

Since LIME relies on color-space segmentation (`quickshift`), its explanations suffer from randomness. GRID-LIME eliminates this variability by enforcing consistent superpixel regions that utilize the black-box model’s predictions to determine grid-based superpixels. However, a limitation of GRID-LIME is that grid-based segmentation may not always align with object boundaries, potentially reducing interpretability in datasets with highly irregular object shapes.

Once the optimal grid size is determined, GRID-LIME applies the LIME framework using these fixed grid-based superpixels. By incorporating the black-box model’s predictions, GRID-LIME improves explanation consistency while maintaining interpretability.

4.2 Datasets

To evaluate the reproducibility of SLICE and compare its performance against LIME, we conduct experiments using three widely recognized image classification datasets: **Oxford-IIIT Pet Dataset**, **PASCAL Visual Object Classes (VOC)**, and **MS COCO**. These datasets provide a diverse range of images, enabling a comprehensive assessment of explanation consistency and fidelity.

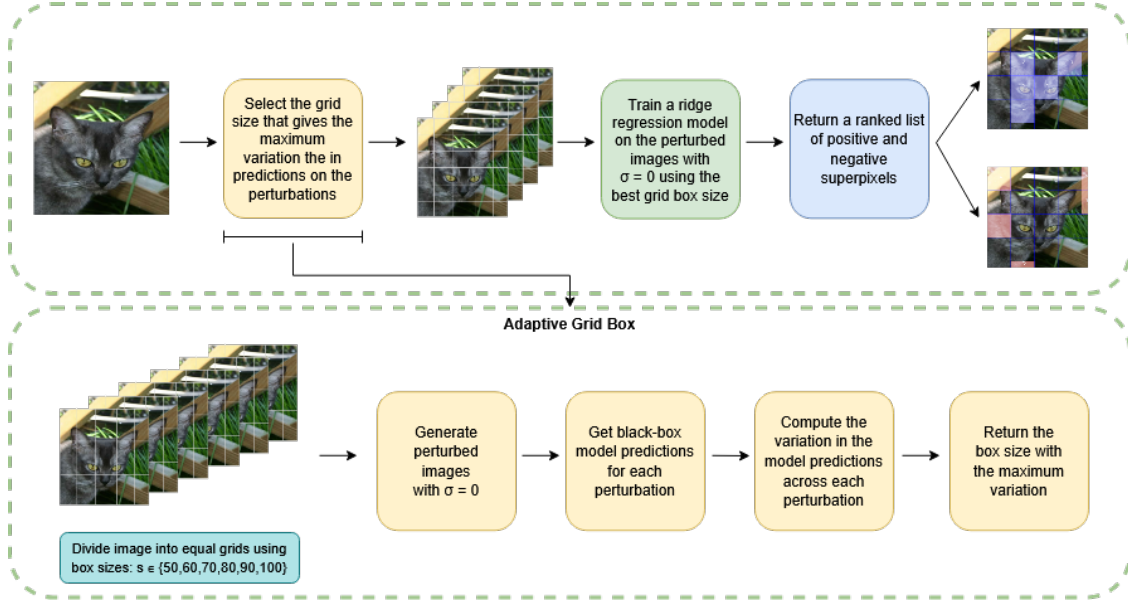


Figure 2: Illustration of GRID-LIME. GRID-LIME replaces color-space-based segmentation with a grid-based approach to ensure stable explanations.

4.2.1 Oxford-IIIT Pet Dataset

The Oxford-IIIT Pet dataset (Parkhi et al., 2012) consists of 37 categories of pet images, covering various breeds of cats and dogs. Each category contains approximately 200 images, leading to a total of 7,349 images. The dataset includes pixel-wise trimap segmentation annotations, enabling precise localization of object regions within the images. This dataset is particularly useful for evaluating explanation methods due to its clear object boundaries and distinct features. In our study, we randomly sample 50 images from this dataset and apply LIME, SLICE, and GRID-LIME to analyze the consistency and fidelity of their explanations.

4.2.2 PASCAL VOC 2009

The PASCAL Visual Object Classes (VOC) 2009 dataset (Everingham et al., 2010) is a benchmark dataset for object recognition tasks. It contains a total of 7054 images in the training and validation set. The dataset includes 20 object categories, covering a broad range of natural and man-made objects, making it suitable for testing explainability methods across different object types. Each image is annotated with bounding boxes and pixel-wise segmentation masks. Given its complexity and variation in object scales, we evaluate whether SLICE and GRID-LIME can improve the stability of LIME explanations. For our experiments, we randomly select 50 images from this dataset for comparative analysis.

4.2.3 MS COCO

The Microsoft Common Objects in Context (MS COCO) dataset (Lin et al., 2015) is one of the largest and most diverse datasets for image recognition and segmentation tasks. It contains over 330,000 images, annotated with 80 object categories, instance-level segmentation masks, and dense captions. This dataset presents a greater challenge due to the presence of multiple objects per image, occlusions, and background clutter, making it an ideal testbed for evaluating robustness in explanation methods. We apply LIME, SLICE, and GRID-LIME to 50 randomly selected images from the validation set to assess their performance in complex visual environments.

4.3 Hyperparameters

To ensure a fair and reproducible comparison between SLICE, LIME, and GRID-LIME, we carefully tune hyperparameters based on prior research and controlled experiments. Each explainer algorithm (SLICE, LIME, or GRID-LIME) is executed for 10 iterations per image to generate the results presented in Section 5.

For SLICE, we train 1000 Ridge regression models to obtain model coefficients for all superpixels. Each Ridge model is trained on 500 randomly generated perturbations of the original test image. The Adaptive Gaussian Blur mechanism in SLICE requires selecting the optimal σ value. We conduct a small-scale hyperparameter search, testing values in the range $[0.1, 0.5]$ with a step size of 0.1, and select the σ value yielding the highest adjusted R^2 .

In the Sign Entropy-based Feature Elimination (SEFE) module, the tolerance limit is set to 3, and the maximum number of iterations is capped at 10 to balance computational efficiency and convergence.

GRID-LIME, our proposed variant, does not require a segmentation algorithm such as quickshift. Instead, it utilizes structured grid-based segmentation, where we empirically determine the optimal grid size by evaluating perturbation consistency across different configurations.

For all experiments, we randomly sample 50 images per dataset and apply the same hyperparameters across SLICE, LIME, and GRID-LIME to maintain consistency. The details of hyperparameter values used in our experiments are provided in Appendix B.

4.4 Experimental setup and code

In this study, we conduct experiments to evaluate the performance of SLICE, LIME, and GRID-LIME in generating interpretable explanations. Our methodology and hyperparameter values follow the original paper, ensuring reproducibility. Additionally, we introduce an auxiliary experiment to assess the alignment of saliency maps with ground-truth segmentation maps. For MS COCO and PASCAL VOC datasets, we sample those images that have only one segmentation object map present for evaluation using Ground Truth Overlap (GTO).

We use ResNet50 (He et al., 2016) and InceptionV3 (Szegedy et al., 2016) models, both pretrained on ImageNet. For superpixel segmentation, we apply `quickshift` (Vedaldi & Soatto, 2008) from the `skimage` library. SLICE iteratively perturbs images by blurring selected superpixels, fits local ridge regression models, and eliminates inconsistent superpixels until convergence. The Gaussian blur parameter σ is optimized via a small hyperparameter search.

To evaluate model explanations, we use multiple metrics assessing consistency and fidelity. Consistency is measured through Average Sign Flip Entropy (ASFE), which quantifies sign stability across multiple runs, and Average Rank Similarity (ARS), which evaluates the consistency of ranked importance scores. We combine these into the Combined Consistency Metric (CCM), a metric introduced by the original paper, providing a unified measure of explanation stability.

Fidelity is assessed using the Area Over Perturbation Curve (AOPC), which tracks classifier confidence changes as important superpixels are perturbed, and the Area Under Curve (AUC), which captures the cumulative impact of these perturbations. Higher AOPC and higher AUC values for insertion, and higher AOPC and lower AUC values for deletion of superpixels indicate that the explainer correctly identifies the most influential features.

As an additional metric, we introduce Ground Truth Overlap (GTO) to evaluate the alignment between model-generated saliency maps and ground-truth segmentation maps. Unlike previous metrics relying on classifier confidence, GTO directly compares the predicted importance regions with known object boundaries. Figure 3 illustrates the computation of the GTO metric and provides an example of its calculation.

$$GTO = \frac{\text{Area of Ground Truth} \cap \text{Area of top-k positive superpixels}}{\text{Area of top-k positive superpixels}} \quad (2)$$

Lastly, our code is available on GitHub⁴. A detailed analysis of the computational requirements and environmental impact of our experiments, including runtime and estimated carbon emissions, is provided in Appendix C.




Figure 3: Ground Truth Overlap (GTO) metric. The intersection between the top-k positive superpixels and the segmentation map is computed and normalized by the area of top-k positive superpixels. In this example, $GTO = 0.58$.

5 Results

We have performed extensive experimentation to test the claims made by Bora et al. (2024). Our evaluation focuses on the interpretability, consistency, and fidelity improvements claimed for SLICE over LIME. We also introduce GRID-LIME as an alternative method and assess its effectiveness.

In this section, we provide a detailed analysis of the consistency and fidelity of explanations produced by these methods. We examine the impact of Sign Entropy-based Feature Elimination (SEFE) and Adaptive Gaussian Blur, and conduct an ablation study to determine their individual contributions. Additionally, we extend our analysis beyond the original paper by introducing the Ground Truth Overlap (GTO) metric to compare explanation maps with segmentation labels.

Overall, our results partially support the claims of Bora et al. (2024). While we confirm that SLICE produces significantly more consistent explanations than LIME, we find that its fidelity improvements are less conclusive. Additionally, we identify potential biases in the evaluation methodology, particularly in the choice of perturbation strategies. Our findings suggest that while SLICE improves consistency, further work is required to ensure fair fidelity comparisons.

Furthermore, we provide additional qualitative examples in Appendix D, illustrating the differences in explanation maps generated by LIME, SLICE, and GRID-LIME. These visualizations help to better understand the distinct attribution patterns produced by each method. In Appendix A, we also present an alternative AOPC evaluation where SLICE is tested under the same perturbation conditions as LIME and GRID-LIME, leading to more comparable fidelity results.

In the following subsections, we analyze each claim in detail and present supporting evidence.

5.1 Claim 1: LIME suffers from inconsistencies due to sign flip and rank variance

To assess explanation stability, we compare the Combined Consistency Metric (CCM) distributions across methods. Figure 4 presents the Kernel Density Estimation (KDE) plots of CCM scores for SLICE, LIME, and GRID-LIME across different datasets and models.

The results indicate that SLICE produces significantly more stable explanations than LIME, as evidenced by a strong peak near 1.0, meaning that sign flips and ranking inconsistencies are minimized. In contrast, LIME’s CCM values are more widely spread, confirming instability. GRID-LIME exhibits a CCM distribution between that of LIME and SLICE, indicating improved stability while maintaining flexibility.

⁴https://anonymous.4open.science/r/slice_reproducibility-8DC6

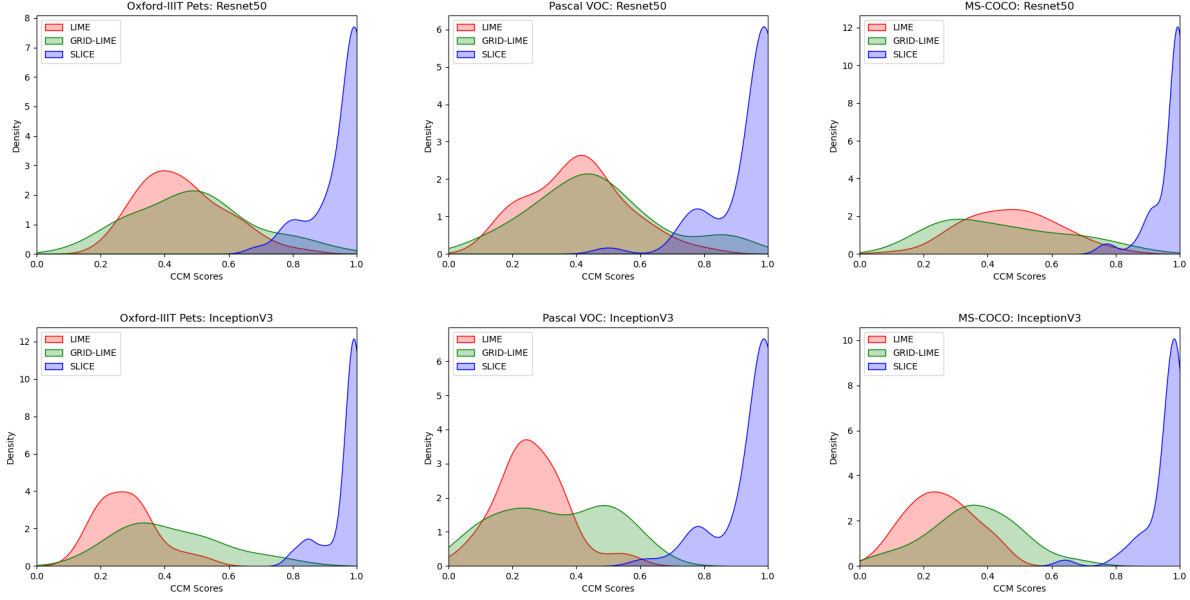


Figure 4: Kernel Density Estimation (KDE) plots of Combined Consistency Metric (CCM) scores for SLICE, LIME, and GRID-LIME. A higher peak in CCM indicates greater explanation consistency, while a wider spread reflects higher variability.

5.2 Claim 2: SEFE improves stability in explanations

The impact of Sign Entropy-based Feature Elimination (SEFE) is evaluated using the Combined Consistency Metric (CCM), which measures explanation stability. Figure 7 consists of six plots, each containing three KDE curves representing different methods. Across these plots, SLICE consistently achieves the highest CCM, with values closest to 1, confirming that it provides the most stable explanations. SLICE_BLUR follows with a CCM around 0.9, indicating strong stability. In contrast, SLICE_SEFE alone exhibits much lower stability, with a CCM ranging between 0.2 and 0.3. While SEFE reduces sign flips and ranking inconsistencies, its direct impact on stability is limited. The improvement in CCM is primarily observed when SEFE is combined with Adaptive Gaussian Blur, rather than when used in isolation. These results suggest that SEFE alone is insufficient for improving stability in explanations and is most effective when integrated with Adaptive Gaussian Blur.

5.3 Claim 3: Adaptive Gaussian Blur enhances explanation consistency

To evaluate the effect of Adaptive Gaussian Blur, we compare the adjusted R^2 values of Ridge regression models trained with different σ values. Our results show that dynamically selecting σ based on entropy increases R^2 , suggesting that perturbations remain realistic while improving consistency. However, the improvement is not always significant across all cases. This can also be observed in Figure 7, where AdaBlur achieves relatively high CCM values, but the difference compared to SLICE without Adaptive Gaussian Blur is not always significant. These findings indicate that while Adaptive Gaussian Blur contributes to stability, its impact depends on the specific setting and dataset.

5.4 Claim 4: SLICE provides better fidelity than LIME

To assess fidelity, we use the Area Over Perturbation Curve (AOPC) and Area Under Curve (AUC) metrics. Figure 5 presents the ECDF plots of AOPC scores.

In the insertion task (first row of Figure 5), SLICE consistently achieves higher AOPC scores than LIME and GRID-LIME, demonstrating better preservation of important features. However, in the deletion task (second row), SLICE performs worse than both methods, suggesting its deletion strategy is less effective.

Figure 6 presents ECDF plots of AUC scores, where SLICE outperforms the other methods in both insertion (first row) and deletion (second row). A higher AUC value is preferable for insertion, while a lower AUC value is better for deletion, confirming SLICE’s strong performance in both tasks.

One potential reason for SLICE’s mixed fidelity performance is the difference in perturbation strategies. SLICE applies Gaussian blur-based perturbations, whereas LIME and GRID-LIME use randomized occlusions. This methodological difference affects how feature importance is evaluated, potentially biasing AOPC and AUC scores. To address this, we introduce an alternative AOPC evaluation (Appendix A), where SLICE is tested under the same perturbation method as LIME and GRID-LIME. This ensures a fair comparison of fidelity by eliminating potential biases introduced by differing perturbation techniques.

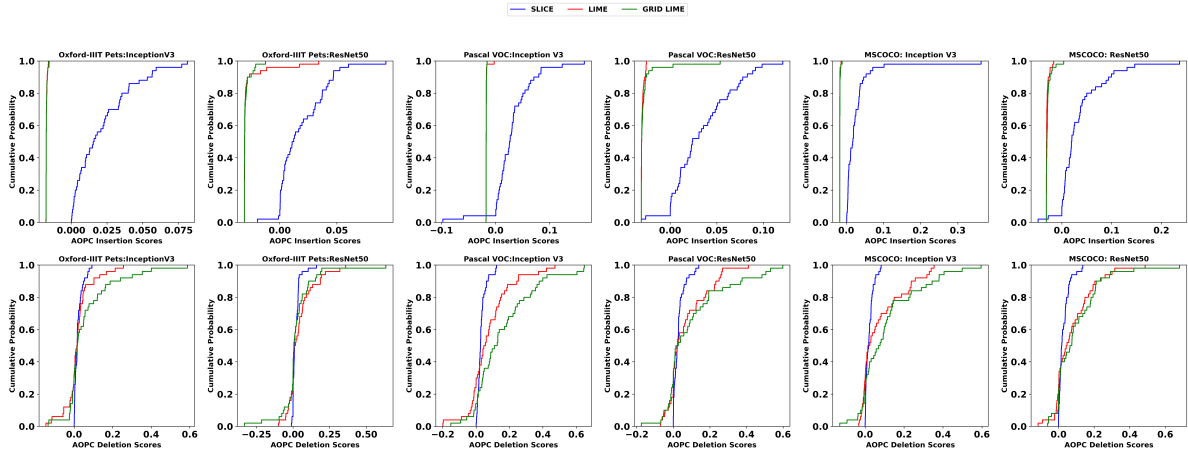


Figure 5: Empirical Cumulative Distribution Function (ECDF) plot of AOPC scores for different methods. The first row corresponds to the insertion task and the second row corresponds to the deletion task. A higher AOPC score indicates better fidelity in capturing important features.

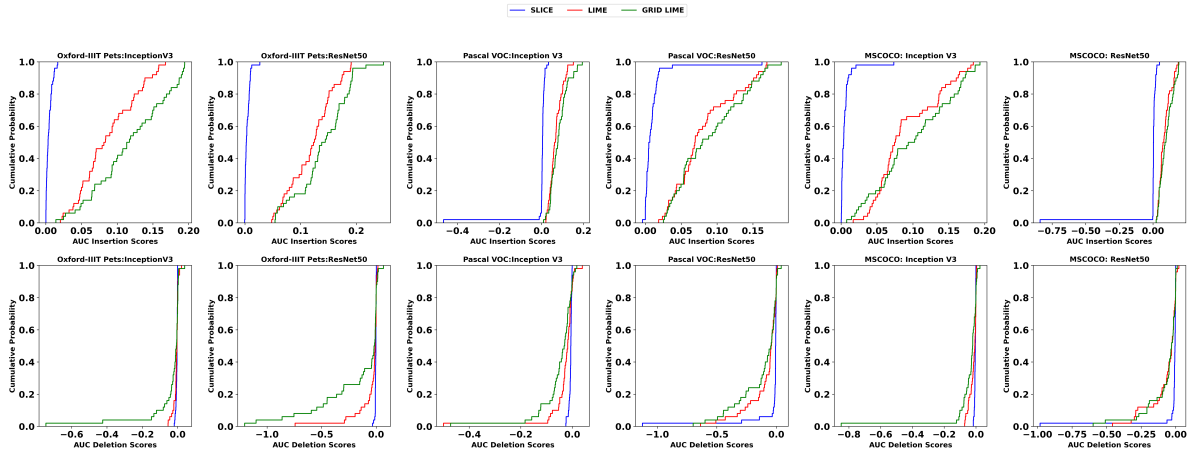


Figure 6: Empirical Cumulative Distribution Function (ECDF) plot of AUC scores for insertion (row 1) and deletion (row 2) tasks. For insertion, a higher AUC value is preferable, indicating that adding back important features improves model confidence. For deletion, a lower AUC value is better, as it suggests that removing key features leads to a significant drop in confidence.

Table 1: Wilcoxon rank test results for comparison of LIME(L), SLICE(S), and GRID-LIME(G) on AOPC. AOPC(x,y) tests whether the median difference AOPCscore(x) - AOPCscore(y) is zero.

Test	D:M	W	p-value	M_Δ	N.C
AOPC(S,L)	O:I	1275.0	8.88e-16	0.032863	0
AOPC(S,L)	O:R	1267.0	2.22e-14	0.041655	2
AOPC(S,L)	P:I	1209.0	1.74e-10	0.042798	2
AOPC(S,G)	O:I	1275.0	8.88e-16	0.032545	0
AOPC(S,G)	O:R	1275.0	8.88e-16	0.04088	0

Table 2: Wilcoxon rank test results for comparison of LIME(L), SLICE(S), and GRID-LIME(G) on AUC. AUC(x,y) tests whether the median difference in AUC scores is zero.

Test	D:M	W	p-value	M_Δ	N.C
AUC(S,L)	O:I	0.0	1.00e+00	-0.075461	50
AUC(S,L)	O:R	0.0	1.00e+00	-0.119034	50
AUC(S,L)	P:I	0.0	1.00e+00	-0.058183	50
AUC(S,G)	O:I	0.0	1.00e+00	-0.110187	50
AUC(S,G)	O:R	0.0	1.00e+00	-0.134202	50

To confirm significance, we conducted a Wilcoxon rank test on AOPC (Table 1) and AUC (Table 2), where D:M denotes Dataset:Model. O represents Oxford-IIT Pets, P refers to Pascal VOC, I indicates InceptionV3, and R denotes ResNet50. While SLICE significantly outperforms LIME in AOPC, the AUC test does not show consistent differences across all comparisons.

5.5 Ablation Study

We conducted an ablation study to replicate the findings of (Bora et al., 2024), evaluating the individual contributions of SEFE and Adaptive Gaussian Blur (AdaBlur) in SLICE. Since SLICE consists of these two key components, we performed the ablation study by isolating each module, as shown in Figure 7.

Our results align with the original ablation study in (Bora et al., 2024). SLICE achieves the highest consistency when both components are included, confirming their complementary roles. Among the individual modules, AdaBlur contributes more to consistency than SEFE. The absence of AdaBlur leads to perturbed images that deviate further from the original, reducing consistency. These findings highlight the importance of using both components together for optimal stability.

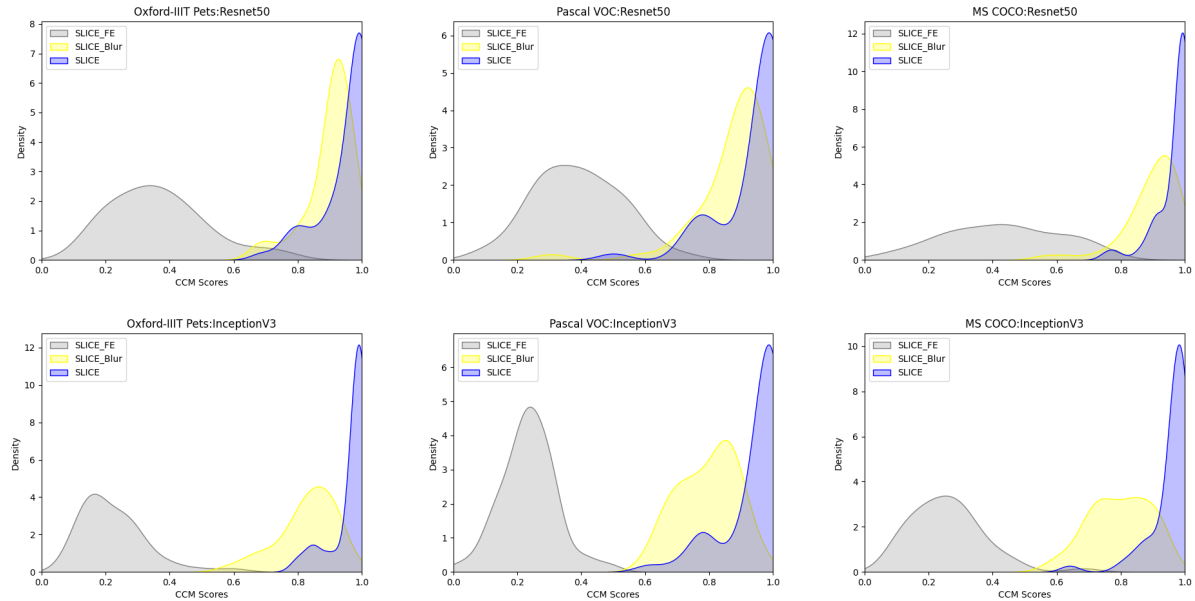


Figure 7: Distribution of CCM scores for ablation experiment. SLICE_FE shows the results for the SEFE while SLICE_Blur shows the results for the Adaptive Gaussian Blur component. SLICE shows the results for both components included. (higher is better).

5.6 Results Beyond the Original Paper

Beyond reproducing the original results, we introduce the Ground Truth Overlap (GTO) metric to evaluate how well explanations align with segmentation maps. Figure 8 presents the results.

LIME consistently achieves the highest GTO scores, indicating better alignment with ground truth segmentations. SLICE generally performs worse across most cases, while GRID-LIME performs better in about half the cases but remains further from LIME. This trend is expected since GTO is based on the Intersection over Union (IoU) of segmentation masks, and grid-based methods naturally achieve higher overlap due to their grid boxes.

Additionally, we introduce an alternative variation of the AOPC ECDF plot in Appendix A, where the insertion logic is modified. Unlike the standard evaluation, this version uses the perturbation method of LIME and GRID-LIME for SLICE instead of applying Gaussian blur. This adjustment ensures that SLICE is evaluated under the same perturbation conditions as the other methods, removing potential biases introduced by differing perturbation strategies. As a result, this version provides a more direct and fair comparison of fidelity. Under this setting, SLICE performs similarly to LIME and GRID-LIME for insertion but slightly worse for deletion.

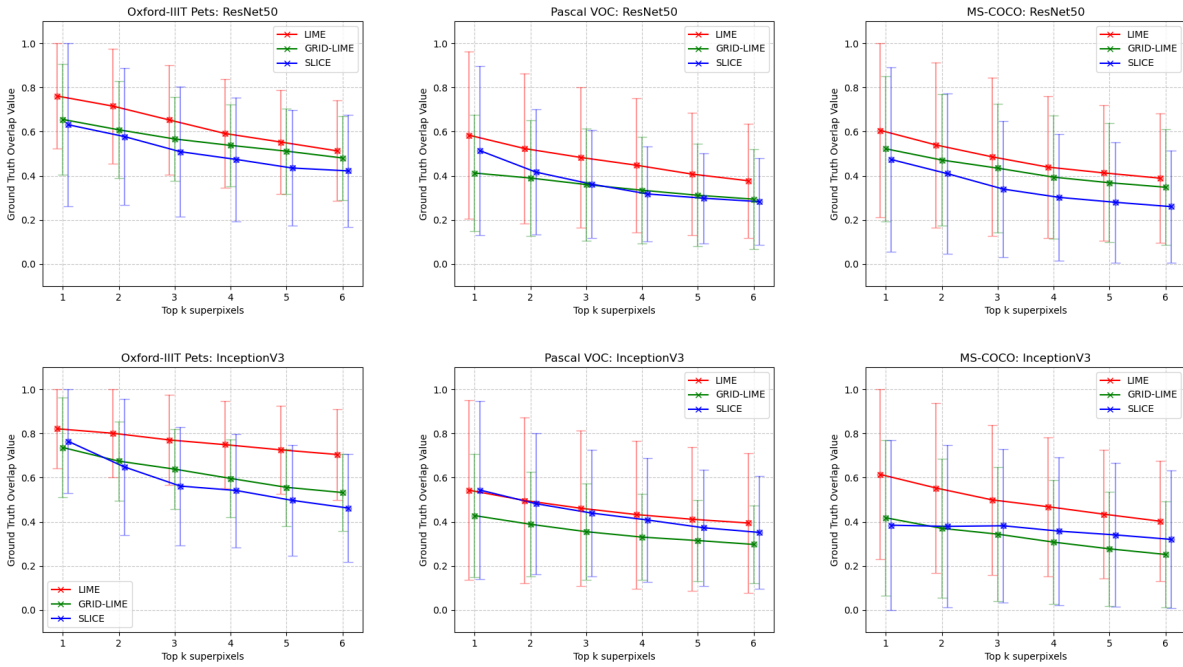


Figure 8: Ground Truth Overlap (GTO) plots, where the x-axis represents the number of top-ranked superpixels, and the y-axis indicates the GTO score (higher is better). LIME achieves the highest alignment with segmentation maps, GRID-LIME performs moderately well, and SLICE lags behind in most cases.

6 Discussion

Our results partially support the claims made by Bora et al. (2024), confirming that SLICE improves explanation consistency but challenging the assertion that it provides superior fidelity. Additionally, we identified methodological limitations in the original study, particularly in how perturbation strategies affect fidelity evaluations.

Across the Oxford-IIIT Pets, PASCAL VOC, and MS COCO datasets, SLICE consistently produced more stable explanations than LIME, as reflected in its higher Combined Consistency Metric (CCM) scores. However, our fidelity analysis using AOPC and AUC showed that SLICE did not always outperform LIME.

Our newly introduced Ground Truth Overlap (GTO) metric further suggested that LIME often produces more segmentation-aligned attributions, despite its instability. This discrepancy likely arises from SLICE’s use of Gaussian blur-based perturbations, which affect how feature importance is preserved compared to LIME’s direct occlusions. To address this, we introduced an alternative AOPC evaluation (Appendix A) where SLICE was tested using the same perturbation strategy as LIME and GRID-LIME, ensuring a fairer fidelity comparison. Under these conditions, SLICE performed similarly to the other methods for insertion but slightly worse for deletion, suggesting that its fidelity advantage in the original setup may be influenced by differences in perturbation techniques rather than intrinsic method effectiveness.

Another important limitation of SLICE is its computational cost. Our runtime analysis revealed that SLICE is significantly slower than LIME, requiring up to 13 times more processing time per experiment. While SLICE improves explanation stability, its increased computational complexity and energy consumption make it less practical for large-scale applications. As an alternative, we introduced GRID-LIME, which replaces LIME’s segmentation-based superpixels with a structured grid. Our results showed that GRID-LIME improves explanation stability while maintaining a computational cost closer to LIME than SLICE. This suggests that part of LIME’s inconsistency stems from its segmentation process rather than its perturbation mechanism. However, GRID-LIME’s rigid grid-based approach does not always align well with object boundaries, potentially affecting interpretability in certain cases.

A key limitation of our study is that our implementation of SLICE may differ slightly from the authors’ original setup. While we referenced their official SLICE code, certain hyperparameter choices and evaluation metric implementations (such as ASFE, ARS, and CCM) had to be re-implemented based on the paper’s descriptions. Despite these challenges, our extensive experimentation and sensitivity analysis suggest that our findings remain robust.

Overall, our study highlights trade-offs in post-hoc explainability methods. While SLICE improves consistency, its computational cost and fidelity claims require further scrutiny. Future research should focus on refining perturbation techniques, optimizing computational efficiency, and developing alternative fidelity metrics that do not rely solely on perturbation-based evaluations.

6.1 What was easy

The authors provided a GitHub repository containing an implementation of SLICE, which facilitated the reproduction of initial results. We were able to validate the consistency of sign and ranking of important superpixels for two images using their provided pickle files. Communication with the authors was helpful, as they provided missing metric code for validation upon request.

6.2 What was difficult

Several challenges emerged during reproduction. The codebase was incomplete, lacking implementations for key metrics such as ASFE, ARS, and CCM. Additionally, AOPC scores were not initially included. While the authors later provided an implementation for AOPC, the other metrics had to be inferred based on the paper’s descriptions.

Another major limitation was computational efficiency. Running the original SLICE implementation took approximately 15 minutes per image, making large-scale experiments infeasible. We improved runtime by rewriting the implementation in PyTorch and incorporating multi-threading, which enabled us to conduct more comprehensive evaluations.

6.3 Communication with original authors

After reviewing the GitHub repository, we identified missing components and contacted the authors for clarification. They responded positively and provided missing code for AOPC calculations. However, ASFE, ARS, and CCM metric implementations were never made available, requiring us to re-implement them independently. Despite this, the provided assistance facilitated certain aspects of reproduction.

References

- Revoti Prasad Bora, Philipp Terhörst, Raymond Veldhuis, Raghavendra Ramachandra, and Kiran Raja. Slice: Stabilized lime for consistent explanations for image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10988–10996, 2024.
- Zeinab Dehghani, Koorosh Aslansefat, Adil Khan, Adín Ramírez Rivera, Franky George, and Muhammad Khalid. Mapping the mind of an instruction-based image editing using smile, 2024. URL <https://arxiv.org/abs/2412.16277>.
- Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010.
- Alicja Gosiewska and Przemyslaw Biecek. Do not trust additive explanations, 2019. URL <https://arxiv.org/abs/1903.11420>.
- Arne Grobrügge, Nidhi Mishra, Johannes Jakubik, and Gerhard Satzger. Explainability in ai-based applications – a framework for comparing different techniques. In *2024 26th International Conference on Business Informatics (CBI)*, pp. 70–79. IEEE, September 2024. doi: 10.1109/cbi62504.2024.00018. URL <http://dx.doi.org/10.1109/CBI62504.2024.00018>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Patrick Knab, Sascha Marton, and Christian Bartelt. Beyond pixels: Enhancing lime with hierarchical features and segmentation foundation models, 2025. URL <https://arxiv.org/abs/2403.07733>.
- Gichan Lee and Scott Uk-Jin Lee. Towards reliable software analytics: Systematic integration of explanations from different model-agnostic techniques. *IEEE Software*, 2023. URL <https://doi.org/10.1109/MS.2023.3241458>.
- Xuhong Li, Haoyi Xiong, Xingjian Li, Xiao Zhang, Ji Liu, Haiyan Jiang, Zeyu Chen, and Dejing Dou. G-lime: Statistical learning for local interpretations of deep neural networks using global priors. *Artificial Intelligence*, 314:103823, 2023. URL <https://doi.org/10.1016/j.artint.2023.103823>.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context, 2015. URL <https://arxiv.org/abs/1405.0312>.
- Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 3498–3505. IEEE, 2012.
- Muhammad Rashid, Elvio G. Amparore, Enrico Ferrari, and Damiano Verda. Using stratified sampling to improve lime image explanations, 2024. URL <https://arxiv.org/abs/2403.17742>.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144, 2016.
- Sheng Shi, Yangzhou Du, and Wei Fan. An extension of lime with improvement of interpretability and fidelity, 2020. URL <https://arxiv.org/abs/2004.12277>.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Andrea Vedaldi and Stefano Soatto. Quick shift and kernel methods for mode seeking. In *Computer Vision—ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part IV 10*, pp. 705–718. Springer, 2008.

- Muhammad Rehman Zafar and Naimul Mefraz Khan. Dlime: A deterministic local interpretable model-agnostic explanations approach for computer-aided diagnosis systems, 2019. URL <https://arxiv.org/abs/1906.10263>.
- Yujia Zhang, Kuangyan Song, Yiming Sun, Sarah Tan, and Madeleine Udell. Why should you trust my explanation? understanding uncertainty in lime explanations, 2019. URL <https://arxiv.org/abs/1904.12991>.
- Xingyu Zhao, Wei Huang, Xiaowei Huang, Valentin Robu, and David Flynn. Baylime: Bayesian local interpretable model-agnostic explanations. In *Uncertainty in artificial intelligence*, pp. 887–896. PMLR, 2021.
- Zhengze Zhou, Giles Hooker, and Fei Wang. S-lime: Stabilized-lime for model explanation. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 2429–2438, 2021.
- Daniel Zwillinger and Stephen Kokoska. *CRC standard probability and statistics tables and formulae*. Crc Press, 1999.

A Alternative AOPC Scores

A modified version of the ECDF plot of AOPC scores with $\sigma = 0$ is presented in Figure 9. In this version, the insertion logic is adjusted such that SLICE uses the perturbation method of LIME and GRID-LIME rather than blurring images with Gaussian blur. This ensures a fairer evaluation by placing SLICE on the same evaluation framework as the other methods.

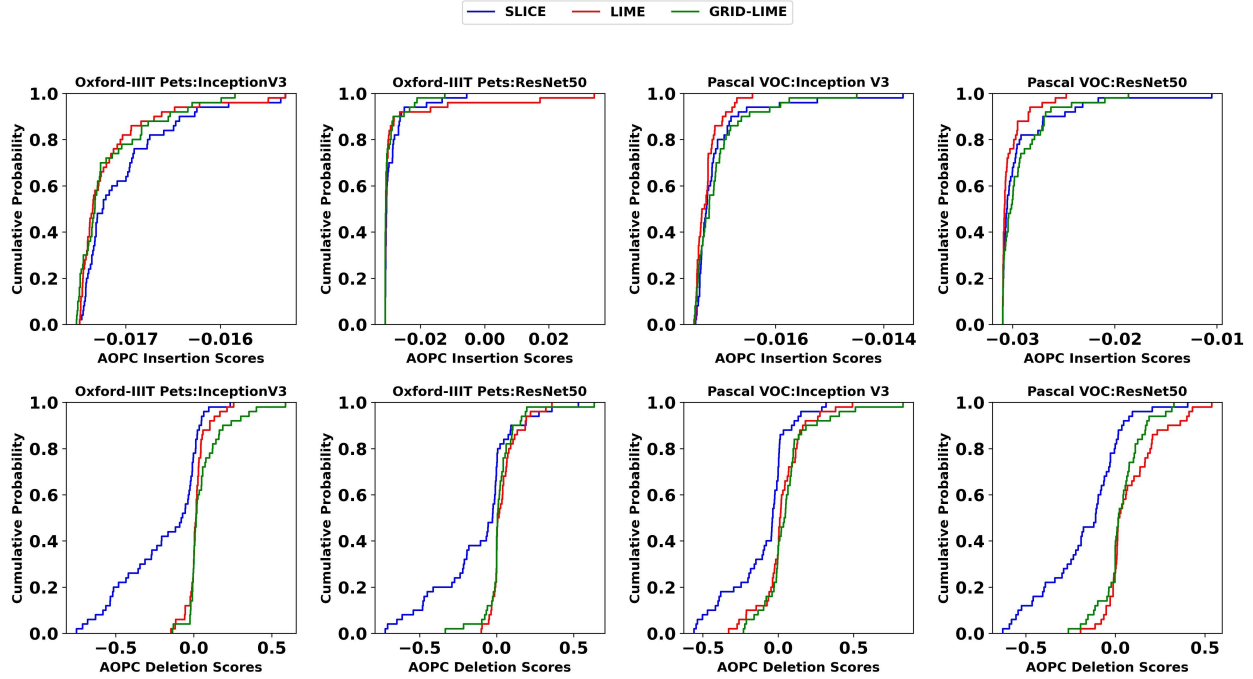


Figure 9: Alternative ECDF plot of AOPC scores with $\sigma = 0$, using LIME and GRID-LIME’s perturbation method for SLICE instead of Gaussian blur (higher is better).

B Hyperparameters

Table 3: Hyperparameter values used in our experiments for SLICE, LIME, and GRID-LIME. The values were determined through empirical validation and small-scale hyperparameter searches to ensure fair and reproducible comparisons. For SLICE, we optimized the Adaptive Gaussian Blur σ using an iterative search strategy, while GRID-LIME’s grid box size is determined dynamically based on perturbation consistency.

Hyperparameter	Value
Adaptive Gaussian Blur σ search range	{0.1, 0.2, 0.3, 0.4, 0.5}
Number of Ridge Regression Models (M)	1000
Number of perturbations per image (N)	500
SEFE Tolerance	3
Number of iterations per image	10
Number of images sampled per dataset	50
Quickshift kernel size	5
Quickshift max distance	200
Quickshift ratio	0.2

C Computational requirements and environmental impact

The table below presents an analysis of computational runtime and associated carbon emissions for different model-explainer-dataset combinations. The explainers considered include LIME, SLICE, and GRID-LIME, applied to ResNet50 and InceptionV3 models trained on three datasets: Oxford-IIIT Pets, Pascal VOC, and MS COCO. The objective is to assess the environmental impact of these machine learning experiments in terms of both execution time and carbon emissions.

The values in the table were estimated using the `ML CO2 Impact calculator`⁵, a tool that provides carbon emission estimates based on hardware type, runtime, and cloud provider. One NVIDIA A100 GPU was used for all computations. The estimated carbon footprint for each experiment is measured in kilograms of CO₂ equivalent.

The `ML CO2 Impact calculator` determines emissions based on multiple factors, including hardware efficiency, runtime, and the energy grid’s carbon intensity. The results indicate clear differences between the explainability methods. LIME and GRID-LIME have consistently lower runtime and emissions compared to SLICE. InceptionV3 requires more computational resources than ResNet50, leading to increased carbon emissions. Dataset size also influences results, with larger datasets requiring more processing time and energy.

Table 4: Computational runtime and carbon emissions per experiment.

Explainer	Model	Dataset	Time (min)	Carbon emissions (kg CO ₂ e)
LIME	ResNet50	Oxford-IIIT Pets	34	0.05
LIME	ResNet50	Pascal VOC	35	0.05
LIME	ResNet50	MSCOCO	37	0.06
LIME	InceptionV3	Oxford-IIIT Pets	51	0.11
LIME	InceptionV3	Pascal VOC	52	0.11
LIME	InceptionV3	MS COCO	52	0.11
SLICE	ResNet50	Oxford-IIIT Pets	332	0.65
SLICE	ResNet50	Pascal VOC	328	0.65
SLICE	ResNet50	MS COCO	334	0.65
SLICE	InceptionV3	Oxford-IIIT Pets	576	1.08
SLICE	InceptionV3	Pascal VOC	581	1.08
SLICE	InceptionV3	MS COCO	593	1.08
GRID-LIME	ResNet50	Oxford-IIIT Pets	35	0.05
GRID-LIME	ResNet50	Pascal VOC	39	0.06
GRID-LIME	ResNet50	MS COCO	41	0.07
GRID-LIME	InceptionV3	Oxford-IIIT Pets	52	0.11
GRID-LIME	InceptionV3	Pascal VOC	53	0.11
GRID-LIME	InceptionV3	MS COCO	53	0.11

⁵<https://mlco2.github.io/impact/>

D Qualitative Examples

To provide additional insight into the explanations generated by each method, we present qualitative examples of LIME, SLICE, and GRID-LIME. These visualizations highlight the ranked positive and negative superpixels, showcasing how each explainer method attributes importance to different image regions.

D.1 LIME Explanations




Figure 10: Qualitative examples of LIME explanations, displaying ranked positive (highlighted in warm colors) and negative (highlighted in cool colors) superpixels. LIME assigns importance by sampling perturbed versions of the image and fitting a linear model to approximate the model’s decision boundary.

D.2 SLICE Explanations




Figure 11: Qualitative examples of SLICE explanations, showing ranked positive and negative superpixels. SLICE applies spatially-aware perturbations and adaptive feature elimination, improving local consistency in attributions.

D.3 GRID-LIME Explanations




Figure 12: Qualitative examples of GRID-LIME explanations, where attributions are structured based on a predefined grid. Positive and negative superpixels are ranked similarly to LIME but constrained within fixed spatial regions, leading to different interpretations.

D.4 Consistency of LIME, SLICE, and GRID-LIME




Figure 13: Examples of positive superpixels returned over multiple runs of LIME, SLICE, and GRID-LIME. LIME exhibits the highest variability, while SLICE demonstrates the most consistent attributions.




Figure 14: Examples of negative superpixels returned over multiple runs of LIME, SLICE, and GRID-LIME.