iGOS++: Integrated Gradient Optimized Saliency by Bilateral Perturbations

Saeed Khorram*

Collaborative Robotics and Intelligent
Systems (CoRIS) Institute, Oregon
State University
khorrams@oregonstate.edu

Tyler Lawson*

Collaborative Robotics and Intelligent
Systems (CoRIS) Institute, Oregon
State University
lawsont@oregonstate.edu

Li Fuxin

Collaborative Robotics and Intelligent
Systems (CoRIS) Institute, Oregon
State University
lif@oregonstate.edu

ABSTRACT

The black-box nature of the deep networks makes the explanation for "why" they make certain predictions extremely challenging. Saliency maps are one of the most widely-used local explanation tools to alleviate this problem. One of the primary approaches for generating saliency maps is by optimizing for a mask over the input dimensions so that the output of the network for a given class is influenced the most. However, prior work only studies such influence by removing evidence from the input. In this paper, we present iGOS++, a framework to generate saliency maps for black-box networks by considering both removal and preservation of evidence. Additionally, we introduce the bilateral total variation term to the optimization that improves the continuity of the saliency map especially under high resolution and with thin object parts. We validate the capabilities of iGOS++ by extensive experiments and comparison against state-of-the-art saliency map methods. Our results show significant improvement in locating salient regions that are directly interpretable by humans. Besides, we showcased the capabilities of our method, iGOS++, in a real-world application of AI on medical data: the task of classifying COVID-19 cases from x-ray images. To our surprise, we discovered that sometimes the classifier is overfitted to the text characters printed on the x-ray images when performing classification rather than focusing on the evidence in the lungs. Fixing this overfitting issue by data cleansing significantly improved the precision and recall of the classifier.

CCS CONCEPTS

• **Computing methodologies** → *Machine learning*; • **Human-centered computing** → *Visualization*; • **Networks** → *Network algorithms*.

KEYWORDS

Explainable AI, Saliency Methods, Integrated-Gradient, COVID-19, Chest X-ray, Medical Imaging

ACM Reference Format:

Saeed Khorram, Tyler Lawson[1], and Li Fuxin. 2021. iGOS++: Integrated Gradient Optimized Saliency by Bilateral Perturbations. In *ACM Conference*

*Equal contribution.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM CHIL '21, April 8–10, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8359-2/21/04.

<https://doi.org/10.1145/3450439.3451865>

on Health, Inference, and Learning (ACM CHIL '21), April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3450439.3451865>

1 INTRODUCTION

As deep networks achieve excellent performance in many tasks, more and more people want to open these black boxes to understand how they make their decisions under the hood. Especially, explaining deep classifiers can help to potentially "debug" them to understand how they make mistakes, and fix those mistakes e.g. by additional data preprocessing. This is increasingly important as deep learning is starting to be used in critical decision-making scenarios such as autonomous driving and medical diagnosis.

Saliency map or heatmap visualization is a fundamental tool for explaining convolutional networks (CNNs). It is mostly used to directly explain classification decisions, but approaches that explain intermediate network nodes or forming new concepts would depend on them as well. In earlier days, heatmap visualizations were mostly based on computing gradient variants of the network with respect to the input [15, 17, 24]. However, due to the highly nonlinear nature of the CNNs, those one-step gradients only account for infinitesimal changes in the function values and do not necessarily correlate to the features that CNNs actually use for decision making [1]. This has caused some of the earlier works to be disillusioned on heatmap research. The most successful of the gradient-based approaches, Grad-CAM [15], partially avoids this issue by not backpropagating into the convolutional layers, hence correlating reasonably well with CNN classification. However, its heatmaps are quite low-resolution since it only works at the final layers where the images are already low-resolution.

Another group of approaches optimizes for a small mask so that CNNs can no longer classify masked images [7, 12]. These approaches provably correlate with the CNN classification, since it can usually be shown that the CNN would no longer predict the original category once the image is masked. However, as we experiment with those approaches, we find that sometimes it is much simpler to "break" the important features CNNs are using to classify, without necessarily capturing all the important features. An intuitive example of this is an object with its important parts being long and thin, as shown in Fig. 1. The mask only needs to cover a small number of points to "break" the legs of the *tiger beetle*, making it disconnected and making the CNN no longer capable of classification. However, those areas, when revealed to the CNN, do not necessarily contain enough information (e.g. the complete legs) for CNNs to come to a correct classification of the image. As a result, those methods do not perform well especially at higher

Figure 1: Comparison of heatmaps generated from [12] (I-GOS) and our proposed approach (iGOS++) for a *tiger beetle* image. I-GOS focuses only on *breaking* existing evidence (e.g. legs of the beetle), hence generated a mask that is highly scattered under high resolution (2nd image from left), and prediction confidence with the masked area (middle image) on the tiger beetle category is very low. However, with the same amount of pixels (6%) from the proposed iGOS++ heatmap, the network has 99.2% confidence (5th image from left)

resolutions, as their masks become disconnected and “adversarial”, focusing only on reducing the CNN prediction rather than locating all the informative areas. Usually, the capability of CNN classifying the image to its original category based on the masked parts drops noticeably when moving from 14×14 heatmap resolution to e.g. 224×224 [12].

In this paper, we propose novel improvements to address this gap. The first novelty is to optimize for an additional insertion loss which aims for the CNN to correctly classify the object even if only a small informative part of the image is revealed. Instead of just adding a simple loss term, we found that separating the deletion mask and insertion mask during the optimization process helps the visualization performance. Our second novelty is to propose a new bilateral total variation term that makes the mask smooth on image areas with similar color, alleviating heatmap diffusion at higher resolutions. A combination of these techniques enables us to obtain heatmap visualizations with significantly better performance, especially when the heatmaps are generated at a high resolution – which are more prone to adversarial solutions if the masks are derived solely on deletion loss. In the tiger-beetle example (Fig.1), it can be seen that iGOS++ is capable of capturing entire legs of the beetle hence capable of generating a high-confidence prediction even with only 6% of the pixels in the original image. We evaluate iGOS++ through extensive experiments and compare our results against the state-of-the-art baselines on non-medical datasets – ImageNet and FashionMNIST as standard benchmarks for evaluation of saliency maps.

We utilized iGOS++ in a real-world problem of detecting COVID-19 patients based on chest x-ray imaging. Interestingly, we found that in some cases, the classifier is overfitted to characters printed on the x-ray images – which clearly should not be related to the underlying indicators of the disease. This illustrates one of the major problems of the current deep network models: their *lack of interpretability*; These black box classifiers can overfit to class priors not expected by humans, which could lead to poor generalization and arbitrary decisions, particularly in high-stake tasks such as in medical diagnosis. Once we pre-processed all images to remove written characters, meaningful performance improvements on the COVID detection task were observed. This shows the utility of the heatmap visualization algorithms in realistic tasks to open up the

black box and reveal some of the biases the classifiers may have learned.

2 RELATED WORK

We only review related work in heatmap visualizations rather than the broader problem of explaining deep networks. Most visualization approaches can be categorized into gradient-based approaches and perturbation-based approaches.

Gradient-based approaches for generating saliency maps commonly use different backpropagation heuristics to derive the sensitivity of the output score with respect to the input. Deconv and Saliency Maps [17, 23] attach special deconvolutional layers to the convolutional layers. Guided-backprop [19] works in a similar, yet different, method to [23], masking values based on negative gradients. [16] multiplies the gradient with the image RGB values. [20] proposes integrated gradients, which compute multiple gradients along a straight line in the image space and average them. [2] computes the relevance across different layers and assigns an importance to each pixel that is used to create a saliency map. [24] uses a winner take all probabilistic model to send top-down signals through the network and generate probabilities based on the weights that are used to create the saliencies. Grad-CAM [15] is the most popular visualization method in this category, it generalizes the existing class activation method [26] to work on any CNN-based neural network and maps class scores back to the previous convolution layer.

However, gradients reflect *one-step* infinitesimal changes in the input, which do not necessarily correspond to a direction in which the output score from a deep network would drop significantly – particularly for deep networks that are highly non-linear functions. In addition, some of these saliency maps were shown to be completely or somewhat independent of the category, only showing strong image edges [1].

Perturbation-based approaches work by modifying the input in some way, e.g. masking, and testing how the output of the network changes. [25] is a method that iteratively removes sections of the image using the gradient until the image contains only the information needed for classification of the target class. In [4], a new network is trained to generate saliency maps. RISE [11] and LIME [13] are similar perturbation based methods that treat the

model as a black-box, and thus do not use gradients at all. They both involve randomly perturbing the image. RISE weighs all of the random masks by the model's confidence and combines them, taking into account each pixels distribution in the random masks. In LIME, the random masks are used to fit a linear model to the black-box model in the local space around the image. The distance from the original image is used in the loss function and the final weights of the linear model are used to generate an explanation for the image. LIME has a reliance on super-pixels to avoid adversarial masks. Some methods utilize optimization with multiple iterations to generate a heatmap visualization [7, 12]. Here the main challenge is the highly non-convex nature of the optimization problem. [7] optimizes a mask to reduce the prediction confidence of a target class. Following [7], [6] uses a fixed-area binary mask that maximally effects the output. This is advantageous as it mitigates the balancing issue in the original mask optimization. [12] is the most related to our work. It combines the algorithms in [7] and [20], by utilizing integrated gradient to optimize the mask. This is shown to significantly improve the performance of the optimization.

3 BILATERALLY-OPTIMIZED PERTURBATIONS

3.1 Background

3.1.1 Heatmap Visualizations by Optimization. We consider the well-known image classification task, where a black-box network f predicts a score $f_c(I_0)$ on class c for input image I_0 . Let $\odot$ denote the Hadamard product.

The idea of optimization-based heatmap visualization is to locate the regions in the input image I_0 that are most important for the network f in outputting $f_c(I_0)$. These local perturbations can be formulated by the inner-product of a real-valued *mask* M to the image $f(I_0 \odot M)$ [7]. Prior work optimizes for the *deletion task*, namely masking the image so that it has *e.g.* low predictive confidence for class c . Afterward, one can visualize the mask to find the salient regions that *caused* the output confidence to decrease. Mathematically:

$$\begin{aligned} \min_M F_c(I_0, M) &= f_c(\Phi(I_0, \tilde{I}_0, M)) + g(M), \\ \text{s.t. } g(M) &= \lambda_1 \|1 - M\|_1 + \lambda_2 \text{TV}(M), \\ \Phi(I_0, \tilde{I}_0, M) &= I_0 \odot M + \tilde{I}_0 \odot (1 - M), \quad 0 \leq M \leq 1, \end{aligned} \quad (1)$$

where $\tilde{I}_0$ is a *baseline* image with near zero evidence about the target class c , $f_c(\tilde{I}_0) \approx \min_I f_c(I)$. It is often chosen to be a constant, white noise, or a blurred version of the image I_0 [7]. The masking operator $\Phi(I_0, \tilde{I}_0, M)$ uses a weighted version between I_0 and $\tilde{I}_0$ to block the influence of certain pixels, decided by the mask M values. The aim of optimization problem (1) is to locate a small and smooth mask M which identifies the regions that are most informative to the black-box f by maximally reducing the output score (predictive confidence on class c) $f_c(\Phi(I_0, \tilde{I}_0, M)) \ll f_c(I_0)$. The regularization term $g(M)$ encourages the mask to be small, by penalizing the magnitude of M with coefficient λ_1 , as well as to be smooth, by penalizing the total-variation (TV) in M [7] with coefficient λ_2 .

3.1.2 Integrated Gradient. Eqn. (1) is a complicated non-convex optimization problem. [7] optimizes the mask by gradient descent. However, this is slow and can take hundreds of iterations to converge. In addition, gradient descent can converge to a local optimum and is not able to jump out of it. [12] has alleviated this issue by using *Integrated-gradient* (IG) [20] rather than conventional gradient as the descent direction for solving the mask optimization. The IG of $f_c(M)$ with respect to M can be formulated as follows,

$$\nabla_{I_0}^{IG} f_c(M) = \frac{1}{S} \sum_{s=1}^S \frac{\partial f_c(\Phi(I_0, \tilde{I}_0, \frac{s}{S}M))}{\partial M}, \quad (2)$$

where it accumulates the gradients along the straight-line path from the perturbed image $\Phi(I_0, \tilde{I}_0, M)$ towards the baseline $\tilde{I}_0$, which approximately solves the global optimum to the unconstrained problem eq. (1). Equivalently, IG can be thought of as performing gradient descent to simultaneously optimize the performance of multiple masks:

$$\min_M \frac{1}{S} \sum_{s=1}^S f_c(\Phi(I_0, \tilde{I}_0, \frac{s}{S}M)) + g(M) \quad (3)$$

which makes it a proper optimization algorithm. Practically, [12] has shown that it improves the optimization performance of eq. (1) as well.

3.2 Bilateral Minimal Evidence

As shown in Fig. 1 and as discussed in the introduction: the capability of destroying a CNN feature does not by itself fully explain the features a CNN used. In this section we propose I-GOS++, which improves upon [12] by also optimizing for the *insertion task*. This would make the CNN predict the original class given information from only a *small* and *smooth* area. We believe that the deletion task and the *insertion task* are complementary and need to be considered simultaneously, since both of them contain important information that introduces a novel "look" into the network behavior by perturbing the input. Besides, considering one task alone is prone to reach *adversarial* solutions [21], particularly when removing evidence. However, it would be unlikely to find adversarial solutions that can satisfy both aforementioned criteria.

A naïve approach to implement this would be to add an insertion loss $-f_c(\Phi(I_0, \tilde{I}_0, 1 - M))$ directly to eq. (1). However, empirically we found that optimizing separate masks performed better than the direct approach. Our method aims to optimize separate deletion and insertion masks over the input with the constraint that the product of the two also satisfies both the deletion and insertion criteria for a target class c , i.e. deletion of the evidence from the input I_0 drastically reduces the output score while retaining the same evidence preserves the initial output score $f_c(I_0)$. Formally, we solve the optimization problem with 3 masks:

$$\begin{aligned} \min_{M=(M_x, M_y)} F_c(I_0, M) &= f_c(\Phi(I_0, \tilde{I}_0, M_x)) \\ &\quad - f_c(\Phi(I_0, \tilde{I}_0, 1 - M_y)) + f_c(\Phi(I_0, \tilde{I}_0, M_{xy})) \\ &\quad - f_c(\Phi(I_0, \tilde{I}_0, 1 - M_{xy})) + g(M_{xy}) \\ \text{subject to } g(M_{xy}) &= \lambda_1 \|1 - M_{xy}\|_1 + \lambda_2 \text{BTV}(M_{xy}); \\ M_{xy} &= M_x \odot M_y; \quad 0 \leq M_x, M_y \leq 1 \end{aligned} \quad (4)$$

where M_x is the deletion mask, M_y is the insertion mask, and their dot product, M_{xy} , is taken as the final solution to the above optimization problem. The BTV term stands for *Bilateral Total Variation* which is explained later in section 3.2.1.

In the above formulation eq. (4), the resolution of the masks is flexible. If the mask has lower resolutions than the original input I_0 , it is first *up-sampled* to the input size using bi-linear interpolation, and then is applied over the image. The choice of the mask resolution depends on the application and the amount of detail desired in the output mask. Commonly, lower resolution masks e.g. 14×14 tend to generate more coarse and smooth saliency maps, while higher resolution masks, e.g. 224×224 , generate more detailed and scattered ones. Lower resolution masks also have the advantage of being more robust against adversarial solutions [12, 21]. In addition, one can add regularization terms on the individual masks M_x and M_y , but we have found regularization only on the final mask M_{xy} to be sufficient.

Note that the integrated gradient for the insertion mask $-f_c(\Phi(I_0, \tilde{I}_0, 1 - M))$ is slightly different from Eq. (2) as it calculates the negative IG along the straight-line path from the image perturbed using the inverse mask $1 - M$, i.e., $\Phi(I_0, \tilde{I}_0, 1 - M)$, toward the original image I_0 ,

$$\nabla_{I_0}^{IG} h_c(M) = -\nabla_{I_0}^{IG} f_c(1 - M) = -\frac{1}{S} \sum_{s=1}^S \frac{\partial f_c\left(\Phi(\tilde{I}_0, I_0, \frac{s}{S}M)\right)}{\partial M}. \quad (5)$$

We use IG to substitute the conventional gradient for the partial objective $f_c(\cdot)$ in eq. (4) and the simple gradient for the convex regularization terms $g(\cdot)$. The following is the total-gradient (TG) for each iteration of the mask update:

$$TG(M) = \nabla_{I_0}^{IG} f_c(M_x) + \nabla_{I_0}^{IG} h_c(M_y) + \nabla_{I_0}^{IG} f_c(M_{xy}) + \nabla_{I_0}^{IG} h_c(M_{xy}) + \nabla g(M_{xy}). \quad (6)$$

TG contains separate integrated gradients w.r.t the deletion and insertion masks (M_x and M_y). $\nabla_{I_0}^{IG}$ is indicative of the direction to the unconstrained problem eq. (4) while $\nabla g(M_{xy})$ regularizes the gradients toward a local and smooth mask and discards unrelated information. Moreover, to make the masks to generalize better and be less dependent over the individual dimensions, we add noise to the perturbed image [7, 12] during each step of the IG calculation in $\nabla_{I_0}^{IG} f_c(M_x) + \nabla_{I_0}^{IG} h_c(M_y)$.

3.2.1 Bilateral Total-Variance. To further alleviate the problem of scattered heatmaps shown in Fig. 1, we introduce a new variation of the TV loss [7], called *Bilateral Total-Variance (BTV)*,

$$BTV = \sum_{u \in \Lambda} e^{-\nabla I(u)^2 / \sigma^2} \|\nabla M(u)\|_{\beta}^{\beta} \quad (7)$$

where $M(u)$ and $I(u)$ are the mask and the input image value at pixel (u) , and β and σ are hyperparameters. This enforces the mask to not only be smooth in its own space but also to consider the pixel value differences in the image space. This is intuitive since BTV would discourage mask value changes when the input image pixels

have similar color. In other words, BTV penalizes the variation in the mask when it is over a single part of an object. This helps particularly in high-resolution mask optimizations and prevents having scattered and adversarial masks.

3.2.2 Backtracking Line Search for Computing the Step Size. Similar to [12], we use backtracking line search at each iteration of mask update. Appropriate step size plays a significant role in avoiding local optimum and accelerates convergence. To that end, we revised the Armijo-Goldstein condition as follows,

$$\sum_{s=1}^S F_c\left(\frac{s}{S}(M^k - \alpha^k \cdot TG(M^k))\right) - \sum_{s=1}^S F_c\left(\frac{s}{S}M^k\right) \leq -\alpha^k \cdot \beta \cdot TG(M^k)^T TG(M^k), \quad (8)$$

where α^k is the step size at time step k and $\beta \in (0, 1)$ is a control parameter. This attempts to determine the maximum movement in the search direction that corresponds to an adequate decrease in the objective function $F_c(M^k)$. Note, this is slightly different from the revised condition in [12] in that the objective function decrease is calculated over the IG intervals rather than at the mask M^k . This is due to the fact that IG actually solves an optimization problem similar to eq. (3).

4 EXPERIMENTS

In this section we will validate the algorithm, first through an experiment in the natural image domain in order to compare with other baselines. Then, we will show the application of the algorithm to the analysis of COVID-19 X-ray images.

4.1 Metrics and Evaluation Setup

We first evaluate the algorithm in the natural image domain to validate it against other baselines. Although visual assessment of the saliency maps might seem straight forward, quantitative comparisons still pose a challenge. For example, one of the widely-used metrics is the *pointing game* [24] which assesses how accurate the saliency maps can locate the objects using the bounding-boxes annotated by *humans*. However, this is not necessarily correlated with the underlying decision-making process of the deep networks, and localization of the objects is only intelligible for humans. Also, pointing game requires correct localization even for *misclassified* examples, and its results correlate poorly with other metrics (e.g. Excitation-BP, which has excellent pointing game scores, suffers heavily from the fallacies pointed out in [1]). This casts doubt on the validity of this metric. Due to these flaws, we do not utilize this metric in our evaluations.

We opted to follow the causal metrics introduced in [11], which is the evolved version of the "deletion game" introduced in [7]. Although not perfect, we find this to be a better evaluation metric, as it performs *interventions* on the input image, a necessary approach to understand causality. These metrics are: the *deletion metric* that evaluates how sharply the confidence of the network drops as regions are removed from the input. Starting from the original image, relevant pixels as indicated by the heatmap are gradually deleted from the image and are replaced by pixels from the highly-blurred baseline. This goes on until all pixels from the original image are

Figure 2: Heatmap visualization from different methods. It can be seen that iGOS++ has better insertion and deletion curves than the baselines. Note the difference in detail between the 28×28 and 224×224 masks (Best viewed in Color)

removed and network has near-zero confidence. The deletion score is the area under the curve (AUC) of the classification confidences. The lower the deletion score is the better. The *Insertion metric*, which is complementary to the deletion one, shows how quickly the original confidence of the network can be reached if relevant evidence are presented. Starting from a baseline image with near-zero confidence, relevant pixels from the image, based on the ranking provided by the heatmap, are gradually inserted into the baseline image. This goes on until all pixels in the baseline are replaced by the original image. The insertion score is also the AUC of the classification scores, with a higher insertion score indicating better performance. While adversarial examples can break the deletion metric with ease and achieve near perfect score, they usually do poorly on the insertion score. Yet only relying on the insertion score will sometimes invite irrelevant background regions to be ranked higher than relevant areas. That is why one needs to jointly look at both scores. See supplementary materials for the set of parameters used in our experiments.

4.2 Results and Analysis

Insertion and Deletion Scores. Table 1 and Table 2 show our results on the Insertion and Deletion metrics, averaged over 5,000 random ImageNet images, for ResNet50 and VGG19 architectures respectively. We experimented with masks of the size 224×224 , 28×28 , 7×7 (only on Resnet50), and 14×14 (only on VGG19) in our experiments. The choice of the masks is to provide fair comparison with other baselines, e.g GradCam [15] that for ResNet50 is best on its “layer4” (7×7 resolution). In addition, Integrated Gradients visualization [20] has heatmaps directly on the input (224×224). Also note that binary masks for RISE [11] are generated at 7×7 and then up-sampled to 224×224 .

As shown in Tables 1 and 2, iGOS++ has superior performance at all resolutions compared to the other approaches. Particularly, for the insertion score, our approach is showing 10%–25% improvement over prior work. This is mainly due to the novel incorporation of the insertion objective in our eq. (4). Note that, for all the methods, the insertion score tends to be lower on higher resolution heatmaps, since it is easier for CNNs to extract features if a significant contiguous part is inserted, rather than isolated pixels in high-resolution masks. On the other hand, on higher resolutions, it is easier to break an image feature by destroying a small number of pixels, hence the deletion metric is usually better. We note that the drop of our insertion performance from 7×7 (or 14×14) to 224×224 is significantly smaller than that of I-GOS, hence it is safer to choose a higher-resolution with iGOS++. Note, we did not notice high variance in our reported results. For instance, on the reported numbers in Table 1 for iGOS++ (224×224 resolution), we obtained the standard deviations of 0.032 ± 0.00023 and 0.722 ± 0.00097 for the insertion and deletion scores, respectively. The low variance shows validity of the experiment setup.

Furthermore, Fig.2 shows visualizations from multiple saliency map methods along with their corresponding deletion and insertion curves and scores. Saliency map for I-GOS and iGOS++ are generated at high (224×224) and medium (28×28) resolutions. It can be seen that iGOS++ performs well even if the object has long and thin parts, whereas I-GOS generates much more scattered heatmaps at high resolutions. GradCAM works well on insertion but included too many irrelevant regions due to its low resolution, a similar issue shared with RISE. For additional visual comparison of our method against various baselines please refer to the supplementary materials.

Figure 3: Sanity check using the Fashion MNIST dataset. A simple CNN model was trained first using ground truth labels and then with random labels. The masks on the second row were generated with the ground truth trained model, while the third row are masks generated for the random label trained model

Sanity Check. Following the work of [1], we perform sanity checks on our method to justify its validity, as only relying on visual assessment can be fallacious. First is the *model randomization* test which checks whether the generated heatmaps are indeed independent of the model parameters or not. Since by randomizing the model (even over one layer), the output score of the network will go near-zero, our method simply returns the initialization $M \approx 1$

Table 1: Quantitative comparison in terms of deletion (lower is better) and insertion (higher is better) metrics using ResNet50. The top row shows the different resolutions

ResNet50	224×224		28×28		7×7	
	Deletion	Insertion	Deletion	Insertion	Deletion	Insertion
GradCam [15]	--	--	--	--	0.1675	0.6521
Integrated Gradients [20]	0.0907	0.2921	--	--	--	--
RISE [11]	0.1196	0.5637	--	--	--	--
Mask [7]	0.0468	0.4962	0.1151	0.5559	0.2259	0.6003
IGOS [12]	0.0420	0.5846	0.1059	0.5986	0.1607	0.6632
iGOS++ (ours)	0.0328	0.7261	0.0929	0.7284	0.1810	0.7332

Table 2: Quantitative comparison in terms of deletion/insertion metrics using VGG19. The top row shows the different resolutions.

VGG19	224×224		28×28		14×14	
	Deletion	Insertion	Deletion	Insertion	Deletion	Insertion
GradCam [15]	--	--	--	--	0.1527	0.5938
Integrated Gradients [20]	0.0663	0.2551	--	--	--	--
RISE [11]	0.1082	0.5139	--	--	--	--
Mask [7]	0.0482	0.4158	0.1056	0.5335	0.1753	0.5647
I-GOS [12]	0.0336	0.5246	0.0899	0.5701	0.1213	0.6387
iGOS++ (ours)	0.0344	0.6537	0.0796	0.7066	0.1123	0.7061

which is different from the final mask we generate. Second, for the *label randomization* test, we trained a simple Convolutional Neural Network (CNN) over the Fashion-MNIST dataset. The CNN was trained with two sets of data. First it was trained using the ground truth labels of the dataset until it reached 100% accuracy. Heatmaps were generated for this model and then compared to maps generated for a model trained with random labels. The idea is if the heatmaps are the same, then the method might be interpreting the image rather than explaining the models decision. Figure 3 shows that iGOS++ generates meaningful heatmaps on real classes that are significantly different from random classes.

Adversarial Examples In addition to the sanity check presented in [1], the visualization from our method on adversarial examples have been analyzed. Figure 4 shows four image samples from ImageNet Validation set where the heatmaps generated on the adversarial images have been visualized along with the original natural images. In this experiment, VGG19 architecture and two different mask resolutions: a *low-resolution* (28×28) and a *high-resolution* (224×224) are used. To generate adversarial examples (target class: *persian cat*), we used the MI-FGSM [5] on the VGG19 model. There are two main takeaways from this analysis: *first*, for adversarial examples, the heatmaps generated by iGOS++ does not provide *meaningful* explanations. *Second*, it can be observed that for adversarial examples, the insertion score is significantly lower and it only reaches the original score after almost all the pixels have been inserted. This shows that the insertion score is a good indication of whether the generated mask is adversarial or not.

Ablation Study. The results from the ablation study are presented in Table 3. The experiments are run at 224×224 and 28×28 resolutions using the ResNet50 model on ImageNet. We observe that optimizing a mask by replacing the deletion loss in [12] with

Table 3: Results from ablation study on ResNet50. The top row shows the different mask resolutions.

Ablation	224×224		28×28	
	Deletion	Insertion	Deletion	Insertion
I-GOS	0.0420	0.5846	0.1059	0.5986
Insertion	0.0760	0.6192	0.1321	0.7231
I-GOS + Insertion (naïve)	0.0322	0.6175	0.2037	0.5103
iGOS++ (no noise)	0.0490	0.5943	0.0904	0.7108
iGOS++ (fix step size)	0.0332	0.5695	0.1052	0.7060
iGOS++ (no BTV)	0.0245	0.6742	0.0813	0.6825
iGOS++	0.0328	0.7261	0.0929	0.7284

an insertion one obtains a good insertion score on both resolutions. However, it significantly hurts the deletion score. On the other hand, naïvely incorporating the insertion loss to eq. (1) by just adding an insertion loss term clearly does not work as well. In fact, it is performing worst than either of the [12] and insertion optimization alone. Further, we find that adding noise makes a clear difference in high resolution (224×224). Moreover, having a fixed step size is worse than using an adaptive step size with line-search. Finally, by removing the bilateral TV term we observed that the deletion score decreases, particularly in high resolution while the insertion score also was impacted negatively. This shows the benefit of the bilateral TV term in avoiding adversarial solutions. In addition, for a comparison between the deletion and insertion masks, please refer to the supplementary materials.

Convergence Behavior. Throughout our experiments, we found the convergence behavior of our method robust against the choice

Figure 4: Comparison of heatmap visualizations from iGOS++ for natural and adversarial examples on two different resolutions: a) 28×28 b) 224×224. The insertion curve only goes up at the end on adversarial examples, indicating that the full image is needed to explain adversarial predictions, while for natural images, the insertion curve goes up quickly. The heatmaps are also different, and for adversarial examples, they do not represent *meaningful* explanations.

of hyperparameters. The choice of hyperparameters were also transferable between all the datasets and models as stated in the *Set of Hyperparameters* section. Table 4, shows the final value of the combined deletion and insertion loss $f_c(\Phi(I_0, \tilde{I}_0, M)) - f_c(\Phi(I_0, \tilde{I}_0, 1-M))$ (eq. (4) in the main paper) after optimizing with iGOS++ and the naïve extension of I-GOS [12] — when the insertion loss is directly added to it. The reported numbers are averaged over 500 images for different choices of hyperparameters λ_1 and λ_2 . Although, the naïve extension is directly optimized over the combined deletion and insertion loss, iGOS++ shows superior performance in all settings. Note the loss value can be negative since the insertion loss is maximized during the optimization. This validates the capability of iGOS++ in achieving a lower objective compared to using the naïve extension of I-GOS for the same objective, showing that it found better optima in our difficult non-convex optimization

problem. Running time comparison of iGOS++ is provided in the supplementary materials.

Table 4: Optimization objective comparison with naïve addition of the insertion loss at 28×28 resolution, averaged over 500 images. Our algorithm resulted a lower objective than the naïve version for the same objective, showing that it found better optima in this difficult non-convex optimization problem. Note that $\lambda_2 = 200$ was never used in the actual experiments and only included here for completeness. Also note the loss could be negative because there is a negative sign on the insertion loss in eq. (4) in the main paper

Loss Value	$\lambda_1=1$ $\lambda_2=2$	$\lambda_1=1$ $\lambda_2=20$	$\lambda_1=10$ $\lambda_2=20$	$\lambda_1=10$ $\lambda_2=200$
I-GOS + Ins (naïve)	-0.869	-0.374	0.346	0.693
iGOS++	-0.971	-0.649	0.158	0.654

4.3 COVID-19 Detection from X-ray Imaging

COVID-19 has been devastating to human lives throughout the world in 2020. Currently, diagnostic tools such as RT-PCR has non-negligible false-negative rates hence it is desirable to be able to diagnose COVID-19 cases directly from chest imaging. X-ray imaging is significantly cheaper than CT or other more high-resolution imaging tools, hence there would be significant socio-economical benefits if one can diagnose COVID-19 reliably from X-ray imaging data, especially in an explainable manner. To that end, we used the COVIDx dataset [9] which is one of the largest publicly-available COVID-19 dataset with 13,786 training samples and 1,572 validation samples comprised of X-ray images from *Normal*, *Pneumonia*, and COVID-19 patients. Following the setting from [9], we trained a classifier over these images.

Somewhat surprisingly, when we applied iGOS++ on the trained classifier, we noticed that in occasional cases the classifier seems to have overfitted to singleton characters printed on the x-ray image (Fig. 5), such that even when only the character region is available, there is a non-negligible chance of classifying for COVID-19. Note that the higher resolution explanation from iGOS++ is important in pinpointing the heatmap to the character whereas low-resolution alternatives such as GradCAM were not informative. For further examples and the case when only text region is revealed to the classifier, please refer to the supplementary materials.

Noticing this “bug” of the classifier, we utilized a state-of-the-art character detector, CRAFT [3], and removed the spurious characters from all the x-ray images in both the training and the testing sets by cleaning and in-painting the detected regions. Examples from the original and cleaned dataset, referred to as *COVIDx++*, can be found in the supplementary materials. The results shown in Table 5 show that the recall of COVID-19 detection on the validation dataset improved by 2.5% and the F1 score improved by more than 1%. This small exercise showcases that “bugs” do exist in deep network classifiers as they do not have common sense on what part of the data is definitely noise, and that heatmap visualizations can help humans locate these bugs as a useful debugging tool. We hope

Figure 5: Showcase of the capability of iGOS++ in detecting bugs in a COVID-19 classification pipeline. It can be noted that unlike Grad-Cam which provides coarse (8×8 resolution) and ineffective explanation, iGOS++ can generate a more-detailed explanation (32×32 resolution) to discover the most salient regions. In the right-most image, only by inserting the top 6% of the pixels from iGOS++ heatmap (highlighting the character “R”), the classifier predicts it as COVID-19 (confidence 43 %).

to dig more into this experiment in the future and obtain more meaningful knowledge from this data.

Table 5: Classification performance on the validation set of the COVIDx and COVIDx++ datasets. COVID-NET [9] is used as the classifier.

Dataset	Accuracy	F1-Score	Precision	Recall
COVIDx	95.19	93.81	95.75	91.85
COVIDx++	95.93	95.08	95.70	94.49

Figure 6 showcases further examples of our visualization method giving insights to the underlying decision making process of the classifier and potentially debugging and improving it. It can be seen that the iGOS++ visualizations (second column) are highly weighting in the text region in the images, i.e. the letter “R”. The third column shows images when only a small fraction of the pixels (e.g. 6%), referred to as *Pixel Ratio* (PR), were inserted back into the baseline image (a Gaussian-blurred of the original image). However, these small amount of pixels, which mostly include the text region, are predicted as COVID-19 classes with their corresponding confidence (shown in green color). On the other hand, using the same iGOS++ visualizations, when only a small fraction of the pixels (e.g. PR: 12%) are removed from the original image, the classifier shows different predictions (than COVID-19). The corresponding confidence (shown in red color) is depicted on the top of the perturbed images. It should be noted that this confidence is not high enough so that the images get classify as COVID-19.

Figure 7 shows examples of the case when the sole insertion of the *text region* into a baseline is enough for the classifier to change its prediction, *falsely to COVID-19*. For the purpose of this analysis, we used a few examples of X-ray images from COVID-19 (Fig. 7a) and *Pneumonia* (Fig. 7b) patients, depicted in the left columns. When feeding the highly-blurred version of these images to the classifier, they were all predicted as *Normal* (the second columns). However, when only revealing the text regions back into the baselines (using the bounding boxed from the CRAFT text detector [3]), the classifier falsely predicts all those images as COVID-19, shown by red color (third columns). This suggests that the presence of text regions is influencing the decisions of the classifier — a bug that

Figure 6: Showcase for capability of the iGOS++ in detecting bugs in a COVID-19 classification pipeline. The X-ray images in the first column belong to COVID-19 patients. The second column shows iGOS++ visualizations at 32×32 resolution where our method highlights the text region, the letter “R”. The third column shows images when only a small fraction of the pixels (e.g. 6%), referred to as *Pixel Ratio* (PR), are inserted back into the baseline. Interestingly, the insertion of these small amount of pixels, which mostly include the text region, is enough for the perturbed images to be classified as COVID-19 (The corresponding confidence scores are shown in color green). The fourth column shows the opposite scenario where only a small fraction of the pixels (e.g. PR: 12%) are removed (blurred) from the original image and that is enough for the perturbed images to be *mis-classified*, i.e. not COVID-19. Obviously, the confidence score for these images (shown in color red) are lower than of Normal or Pneumonia classes.

we hypothesized using our visualizing method. Nevertheless, in reality, one assumes these should be regarded as noise in the data and should be discarded by the classifier. This is an example of the fragility of trusting the decision making of *black-box* classifiers, particularly in tasks where human life is at stake such as medical diagnosis.

5 CONCLUSION AND DISCUSSION

We propose a new approach, iGOS++, for creating a heatmap visualization to explain decisions made by the CNNs. iGOS++ differs from

Figure 7: The prediction of the classifier for each image is shown on the top. X-ray images of a) COVID-19 b) Pneumonia patients from the COVIDx dataset [9] are depicted in the left columns. When the highly-blurred version of these images are fed to the classifier, they are all predicted as *Normal* (middle columns). However, when only the text regions were inserted back into the baselines, the classifier falsely predicts all those images as COVID-19, shown by red color (third columns). This suggests that the text regions are directly influencing the decision making of the classifier.

other approaches by optimizing for separate deletion and insertion masks which are tied together to output a single explanation. We show empirically that with this approach, significantly better insertion performance at all resolutions can be achieved. Besides, we introduce a new term for regularization using the bilateral total variance of the image. This is shown to improve the smoothness of the generated heatmaps as well. As a real-life example, we showed that in a task of classifying COVID-19 patients from x-ray images, sometimes the classifier would overfit to the characters printed on the image. Removing this bug improved the classifier performance meaningfully. We hope the high-fidelity heatmaps generated with iGOS++ will be helpful for downstream tasks in explainable deep learning in both the natural image domain and medical imaging domain.

ACKNOWLEDGEMENTS

We thank Dr. Alan Fern for helpful discussions. This work is supported in part by DARPA contract N66001-17-2-4030.

REFERENCES

- [1] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. 2018. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*. 9505–9515.
- [2] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus Robert Müller, and Wojciech Samek. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS One* 10 (10 7 2015). <https://doi.org/10.1371/journal.pone.0130140>
- [3] Youngmin Baek, Bado Lee, Dongyoon Han, Sangdoo Yun, and Hwalsuk Lee. 2019. Character Region Awareness for Text Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 9365–9374.

- [4] Piotr Dabkowski and Yarín Gal. 2017. Real time image saliency for black box classifiers. In *Advances in Neural Information Processing Systems*. 6967–6976.
- [5] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. 2018. Boosting Adversarial Attacks With Momentum. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [6] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. 2019. Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE International Conference on Computer Vision*. 2950–2958.
- [7] R. C. Fong and A. Vedaldi. 2017. Interpretable Explanations of Black Boxes by Meaningful Perturbation. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 3449–3457. <https://doi.org/10.1109/ICCV.2017.371>
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [9] Zhong Qiu Lin Linda Wang and Alexander Wong. 2020. COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest Radiography Images. *arXiv:2003.09871 [cs.CV]*
- [10] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. [n.d.]. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32.
- [11] Vitali Petsiuk, Abir Das, and Kate Saenko. 2018. RISE: Randomized Input Sampling for Explanation of Black-box Models. In *Proceedings of the British Machine Vision Conference (BMVC)*.
- [12] Zhongang Qi, Saeed Khorram, and Li Fuxin. 2020. Visualizing deep networks by optimizing with integrated gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34.
- [13] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. 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*. ACM, 1135–1144.
- [14] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* (2015).
- [15] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. 2017. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 618–626.
- [16] Avanti Shrikumar, Peyton Greenside, Anna Shcherbina, and Anshul Kundaje. 2016. Not Just a Black Box: Learning Important Features Through Propagating Activation Differences. *CoRR abs/1605.01713* (2016). [arXiv:1605.01713](https://arxiv.org/abs/1605.01713)
- [17] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *ICLR Workshop* (2014).
- [18] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. *International Conference on Learning Representations* (2015).
- [19] J.T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. 2015. Striving for Simplicity: The All Convolutional Net. In *ICLR Workshop*. <http://lmb.informatik.uni-freiburg.de/Publications/2015/DB15a>
- [20] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic Attribution for Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 3319–3328.
- [21] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *International Conference on Learning Representations*. <http://arxiv.org/abs/1312.6199>
- [22] Alexandru Telea. 2004. An image inpainting technique based on the fast marching method. *Journal of graphics tools* 9, 1 (2004), 23–34.
- [23] Matthew D. Zeiler and Rob Fergus. 2014. Visualizing and Understanding Convolutional Networks. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 818–833.
- [24] Jianming Zhang, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. 2016. Top-down neural attention by excitation backprop. In *European Conference on Computer Vision*. Springer, 543–559.
- [25] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2014. Object Detectors Emerge in Deep Scene CNNs. *CoRR abs/1412.6856* (2014). [arXiv:1412.6856](https://arxiv.org/abs/1412.6856) <https://arxiv.org/abs/1412.6856>
- [26] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2921–2929.