
DANCE: Enhancing saliency maps using decoys

Yang Young Lu^{*1} Wenbo Guo^{*2} Xinyu Xing² William Stafford Noble³

Abstract

Saliency methods can make deep neural network predictions more interpretable by identifying a set of critical features in an input sample, such as pixels that contribute most strongly to a prediction made by an image classifier. Unfortunately, recent evidence suggests that many saliency methods poorly perform, especially in situations where gradients are saturated, inputs contain adversarial perturbations, or predictions rely upon inter-feature dependence. To address these issues, we propose a framework, DANCE, which improves the robustness of saliency methods by following a two-step procedure. First, we introduce a perturbation mechanism that subtly varies the input sample without changing its intermediate representations. Using this approach, we can gather a corpus of perturbed (“decoy”) data samples while ensuring that the perturbed and original input samples follow similar distributions. Second, we compute saliency maps for the decoy samples and propose a new method to aggregate saliency maps. With this design, we offset influence of gradient saturation. From a theoretical perspective, we show that the aggregated saliency map not only captures inter-feature dependence but, more importantly, is robust against previously described adversarial perturbation methods. Our empirical results suggest that, both qualitatively and quantitatively, DANCE outperforms existing methods in a variety of application domains.¹

^{*}Equal contribution ¹Department of Genome Sciences, University of Washington, Seattle, WA, USA ²College of Information Sciences and Technology, The Pennsylvania State University, State College, PA, USA ³Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, WA, USA. Correspondence to: Xinyu Xing <xxing@ist.psu.edu>, William Stafford Noble <william-noble@uw.edu>.

1. Introduction

Deep neural networks (DNNs) deliver remarkable performance in an increasingly wide range of application domains, but they often do so in an inscrutable fashion, delivering predictions without accompanying explanations. In a practical setting such as automated analysis of pathology images, if a patient sample is classified as malignant, then the physician will want to know which parts of the image contribute to this diagnosis. Thus, in general, a DNN that delivers interpretations alongside its predictions will enhance the credibility and utility of its predictions for end users (Lipton, 2016).

In this paper, we focus on a popular branch of explanation methods, often referred to as *saliency methods*, which aim to find input features (*e.g.*, image pixels or words) that strongly influence the network predictions (Simonyan et al., 2013; Selvaraju et al., 2016; Binder et al., 2016; Shrikumar et al., 2017; Smilkov et al., 2017; Sundararajan et al., 2017; Ancona et al., 2018). Saliency methods typically rely on back-propagation from the network’s output back to its input to assign a saliency score to individual features so that higher scores indicate higher importance to the output prediction. Despite attracting increasing attention, saliency methods suffer from several fundamental limitations:

- **Gradient saturation** (Sundararajan et al., 2017; Shrikumar et al., 2017; Smilkov et al., 2017) may lead to the problem that the gradients of important features have small magnitudes, breaking down the implicit assumption that important features, in general, correspond to large gradients. This issue can be triggered when the DNN outputs are flattened in the vicinity of important features.
- **Importance isolation** (Singla et al., 2019) refers to the problem that gradient-based saliency methods evaluate the feature importance in an isolated fashion, implicitly assuming that the other features are fixed.
- **Perturbation sensitivity** (Ghorbani et al., 2017; Kindermans et al., 2017; Levine et al., 2019) refers to the observation that even imperceptible, random perturbations or a simple shift transformation of the input data may lead to a large change in the resulting saliency scores.

In this paper, we propose a novel saliency method, Decoy-

enhANCED saliency (DANCE), to tackle these limitations. At a high level, DANCE generates the saliency score of an input by aggregating the saliency scores of multiple perturbed copies of this input. Specifically, given an input sample of interest, DANCE first generates a population of perturbed samples, referred to as *decoys*, that perfectly mimic the neural network’s intermediate representation of the original input. These decoys are used to model the variation of an input sample originating from either sensor noise or adversarial attacks. The decoy construction procedure draws inspiration from the *knockoffs*, proposed recently by Barber & Candès (2015) in the setting of error-controlled feature selection, where the core idea is to generate knock-off features that perfectly mimic the empirical dependence structure among the original features.

In brief, the current paper makes three primary contributions. First, we propose a framework to perturb input samples to produce corresponding decoys that preserve the input distribution, in the sense that the intermediate representations of the original input data and the decoys are indistinguishable. We formulate decoy generation as an optimization problem, applicable to diverse deep neural network architectures. Second, we develop a decoy-enhanced saliency score by aggregating the saliency maps of generated decoys. By design, this score naturally offsets the impact of gradient saturation. From a theoretical perspective, we show how the proposed score can simultaneously reflect the joint effects of other dependent features and achieve robustness to adversarial perturbations. Third, we demonstrate empirically that DANCE outperforms existing saliency methods, both qualitatively and quantitatively, on three real-world applications. We also quantify DANCE’s advantage over existing saliency methods in terms of robustness against various adversarial attacks.

2. Related work

A variety of saliency methods have been proposed in the literature. Some, such as edge detectors and Guided Back-propagation (Springenberg et al., 2014) are independent of the predictive model (Nie et al., 2018; Adebayo et al., 2018).² Others are designed only for specific architectures (*i.e.*, Grad-CAM (Selvaraju et al., 2016) for CNNs, DeConvNet for CNNs with ReLU activations (Zeiler & Fergus, 2014)). In this paper, instead of exhaustively evaluating all saliency methods, we apply our method to three saliency methods that do depend on the predictor (*i.e.*, passing the sanity checks in Adebayo et al. (2018) and Sixt et al. (2020)) and are applicable to diverse DNN architectures.

- The **vanilla gradient** method (Simonyan et al., 2013)

²Sixt et al. (2020) shows that LRP (Binder et al., 2016) is independent of the parameters of certain layers.

calculates the gradient of the class score with respect to the input \mathbf{x} , defined as

$$E_{grad}(\mathbf{x}; F^c) = \nabla_{\mathbf{x}} F^c(\mathbf{x})$$

- **SmoothGrad** (Smilkov et al., 2017) seeks to reduce noise in the saliency map by averaging over explanations of the noisy copies of an input, defined as

$$E_{sg}(\mathbf{x}; F^c) = \frac{1}{N} \sum_{i=1}^N E_{grad}(\mathbf{x} + g_i; F^c)$$

where $g_i \sim N(0, \sigma^2)$ indicates the noise vectors.

- The **integrated gradient** method (Sundararajan et al., 2017) starts from a baseline input \mathbf{x}^0 and sums over the gradient with respect to scaled versions of the input ranging from the baseline to the observed input, defined as

$$E_{ig}(\mathbf{x}; F^c) = (\mathbf{x} - \mathbf{x}^0) \times \int_0^1 \nabla_{\mathbf{x}} F^c(\mathbf{x}^0 + \alpha(\mathbf{x} - \mathbf{x}^0)) d\alpha$$

Note that *input \odot gradient* and DeepLIFT (Shrikumar et al., 2017) are strongly related to the integrated gradient method, as shown by Ancona et al. (2018).

We do not empirically compare to several other categories of methods. *Counterfactual-based methods* work under the same setup as saliency methods, providing explanations for the predictions of a pre-trained DNN model (Sturmfels et al., 2020). These methods identify the important subregions within an input image by perturbing the subregions (by adding noise, rescaling (Sundararajan et al., 2017), blurring (Fong & Vedaldi, 2017), or inpainting (Chang et al., 2019)) and measuring the resulting changes in the predictions (Ribeiro et al., 2016; Lundberg & Lee, 2017; Chen et al., 2018; Fong & Vedaldi, 2017; Dabkowski & Gal, 2017; Chang et al., 2019; Yousefzadeh & O’Leary, 2019; Goyal et al., 2019). Although these methods do identify meaningful subregions in practice, they exhibit several limitations. First, counterfactual-based methods implicitly assume that regions containing the object most contribute to the prediction (Fan et al., 2017). However, Moosavi-Dezfooli et al. (2017) showed that counterfactual-based methods are also vulnerable to adversarial attacks, which force these methods to output unrelated background rather than the meaningful objects as important subregions. Second, the counterfactual images may be potentially far away from the training distribution, causing ill-defined classifier behavior (Burns et al., 2019; Hendrycks & Dietterich, 2019).

In addition to these limitations, counterfactual-based methods and our decoy-based method are fundamentally different in three ways. First, the former seeks the minimum

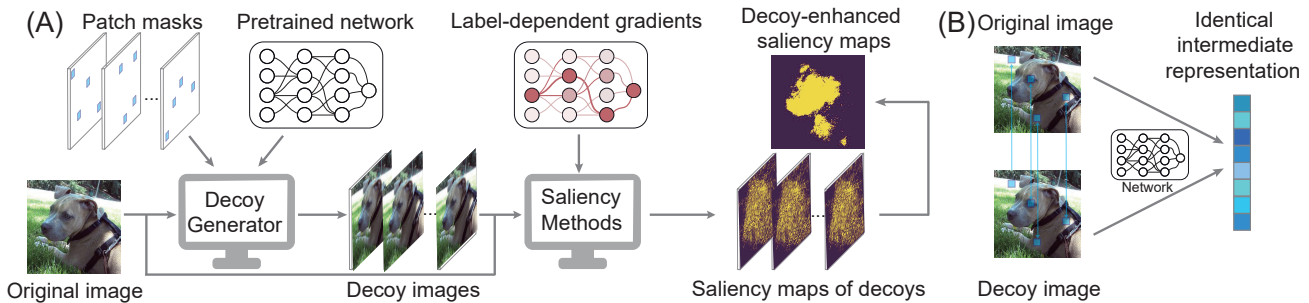


Figure 1. Overview of DANCE. (A) The DANCE workflow. (B) The swapping operation between original and decoy images.

set of features to exclude in order to minimize the prediction score or to include in order to maximize the prediction score (Fong & Vedaldi, 2017), whereas our approach aims to characterize the influence of each feature on the prediction score. Second, counterfactual-based methods explicitly consider the decision boundary by comparing each image to the closest image on the other side of the boundary. In contrast, the proposed method only considers the decision boundary implicitly by calculating the gradient’s variants. Third, unlike counterfactual images, which could potentially be out-of-distribution, decoys are plausibly constructed in the sense that their intermediate representations are indistinguishable from the original input data by design. Because of these limitations and differences, we do not compare our method with counterfactual-based methods.

In addition to saliency methods and counterfactual-based methods, several other types of interpretation methods have been proposed that either aim for a different goal or have a different setup. For example, recent research (e.g., Ribeiro et al. (2016); Lundberg & Lee (2017); Chen et al. (2018; 2019b)) designed techniques to explain a black-box model, where the model’s internal weights are inaccessible. Koh & Liang (2017) and some follow-up work (Yeh et al., 2018; Koh et al., 2019) tried to find the training points that are most influential for a given test sample. Some other efforts have been made to train a more interpretable DNN classifier (Fan et al., 2017; Zolna et al., 2019; Alvarez-Melis & Jaakkola, 2018; Toneva & Wehbe, 2019), synthesize samples that represent the model predictions (Ghorbani et al., 2019; Chen et al., 2019a), or identifying noise-tolerant features (Ikeno & Hara, 2018; Schulz et al., 2020). However, due to the task and setup differences, we do not consider these methods in this paper.

3. Methods

3.1. Problem setup

Consider a multi-label classification task in which a pre-trained neural network model implements a function F :

$\mathbb{R}^d \mapsto \mathbb{R}^C$ that maps from the given input $\mathbf{x} \in \mathbb{R}^d$ to C predicted classes. The score for each class $c \in \{1, \dots, C\}$ is $F^c(\mathbf{x})$, and the predicted class is the one with maximum score, i.e., $\arg \max_{c \in \{1, \dots, C\}} F^c(\mathbf{x})$. A saliency method aims to assign to each feature a saliency score, encoded in a saliency map $E(\mathbf{x}; F^c) : \mathbb{R}^d \mapsto \mathbb{R}^d$, in which the features with higher scores represent higher “importance” relative to the final prediction.

Given a pre-trained neural network model F with L layers, an input \mathbf{x} , and a saliency method E such that $E(\mathbf{x}; F)$ is a saliency map of the same dimensions as \mathbf{x} , DANCE operates in two steps: generating decoys and aggregating the saliency maps of the decoys (Figure 1A).

3.2. Decoy definition

Say that $F_\ell : \mathbb{R}^d \mapsto \mathbb{R}^{d_\ell}$ is the function instantiated by the given network, which maps from an input $\mathbf{x} \in \mathbb{R}^d$ to its intermediate representation $F_\ell(\mathbf{x}) \in \mathbb{R}^{d_\ell}$ at layer $\ell \in \{1, 2, \dots, L\}$. A vector $\tilde{\mathbf{x}} \in \mathbb{R}^d$ is said to be a decoy of $\mathbf{x} \in \mathbb{R}^d$ at a specified layer ℓ if the following swappable condition is satisfied:

$$F_\ell(\mathbf{x}) = F_\ell(\mathbf{x}_{\text{swap}(\tilde{\mathbf{x}}, \mathcal{K})}), \quad (1)$$

for swappable features $\mathcal{K} \subset \{1, \dots, d\}$.

Here, the $\text{swap}(\tilde{\mathbf{x}}, \mathcal{K})$ operation swaps features between \mathbf{x} and $\tilde{\mathbf{x}}$ based on the elements in \mathcal{K} . In this work, \mathcal{K} represents a small meaningful feature set, which represents a small region/segment in an image or a group of words (embeddings) in a sentence. Take an image recognition task for example. Assume $\mathcal{K} = \{10\}$ and $\tilde{\mathbf{x}}$ is a zero matrix, then $\mathbf{x}_{\text{swap}(\tilde{\mathbf{x}}, \mathcal{K})}$ indicates a new image that is identical to \mathbf{x} except that the tenth pixel is set to zero. An illustrative explanation of a swap operator is shown in Figure 1(B).

Using the swappable condition, we aim to ensure that the original image \mathbf{x} and its decoy $\tilde{\mathbf{x}}$ are indistinguishable in terms of the intermediate representation at layer ℓ . Note in particular that the construction of decoys relies solely on the first ℓ layers of the neural network F_1, F_2, \dots, F_ℓ and is independent of the succeeding layers $F_{\ell+1}, \dots, F_L$. As

such, $\tilde{\mathbf{x}}$ is conditionally independent of the classification task $F(\mathbf{x})$ given the input \mathbf{x} ; *i.e.*, $\tilde{\mathbf{x}} \perp\!\!\!\perp F(\mathbf{x})|\mathbf{x}$.

3.3. Decoy generation

To identify decoys satisfying the swappable condition, we solve the following optimization problem:

$$\begin{aligned} & \text{maximize}_{\tilde{\mathbf{x}} \in [\mathbf{x}_{\min}, \mathbf{x}_{\max}]^d} \|((\tilde{\mathbf{x}} - \mathbf{x}) \cdot s)^+\|_1, \\ & \text{s.t.} \quad \begin{cases} \|F_\ell(\tilde{\mathbf{x}}) - F_\ell(\mathbf{x})\|_\infty \leq \epsilon, \\ (\tilde{\mathbf{x}} - \mathbf{x}) \circ (1 - \mathcal{M}) = 0 \end{cases} \end{aligned} \quad (2)$$

Here, $(\cdot)^+ = \max(\cdot, 0)$, and the operators $\|\cdot\|_1$ and $\|\cdot\|_\infty$ correspond to the L_1 and L_∞ norms, respectively. $\mathcal{M} \in \{0, 1\}^d$ is a specified binary mask, where $\mathcal{M}_i = 0$ indicates that the i th features of \mathbf{x} and $\tilde{\mathbf{x}}$ are kept the same (realized by the constraint $(\tilde{\mathbf{x}} - \mathbf{x}) \circ (1 - \mathcal{M}) = 0$). In other words, we take $\tilde{\mathbf{x}}$ and \mathbf{x} to be indistinguishable except for the swappable features indicated by the mask (*i.e.*, $\mathbf{x}_{\text{swap}(\tilde{\mathbf{x}}, (1-\mathcal{M}))} = \tilde{\mathbf{x}}$). The value of each feature in the decoy $\tilde{\mathbf{x}}$ is restricted to lie in a legitimate value range *i.e.*, $[\mathbf{x}_{\min}, \mathbf{x}_{\max}]$ (*e.g.*, the pixel values should lie in $[0, 255]$). We further impose the constraint $\|F_\ell(\tilde{\mathbf{x}}) - F_\ell(\mathbf{x})\|_\infty \leq \epsilon$, which ensures that the generated decoy satisfies the swappable condition described in Equation (1).

As illustrated in Figure 1, a population of n patch masks are constructed subject to the principle that each swappable patch is covered at least once. Because each swappable patch is small (*e.g.*, a small region/segment in an image), assigning each patch mask to a single patch would be computationally expensive. Accordingly, we aggregate multiple patches into a combined patch mask for computational efficiency (see Supplementary Section S1 for details). Empirical results suggest that DANCE is robust to the number of patches that are aggregated into each mask (Figure 2C).

As is shown later in Section 3.4, DANCE aims to capture the range of the saliency maps among all decoys. To achieve this, we first need to estimate the range of values among the decoys by estimating the range of perturbation values that can be added to the input without violating the swappable condition. In other words, we maximize the deviation between $\tilde{\mathbf{x}}$ and \mathbf{x} from both the positive and negative directions, *i.e.*, $s = +1$ and $s = -1$. As shown in Equation (2), for each specified mask \mathcal{M} , we compute two decoys—one for the positive deviation (*i.e.*, $s = +1$) and the other for the negative one (*i.e.*, $s = -1$). More details about how to optimize Equation 2 can be found in Supplementary Section S1.

3.4. Decoy-enhanced saliency scores

We denote the generated decoys as $\{\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2, \dots, \tilde{\mathbf{x}}^{2n}\}$, *i.e.*, n decoys in the positive direction and n in the negative direction. For these decoys, we then apply a

given saliency method E to yield the corresponding decoy saliency maps $\{E(\tilde{\mathbf{x}}^1; F), E(\tilde{\mathbf{x}}^2; F), \dots, E(\tilde{\mathbf{x}}^{2n}; F)\}$. With these decoy saliency maps in hand, for each feature \mathbf{x}_i in \mathbf{x} , we characterize its saliency score variation by using a population of saliency scores $\tilde{E}_i = \{E(\tilde{\mathbf{x}}^1; F^c)_i, E(\tilde{\mathbf{x}}^2; F^c)_i, \dots, E(\tilde{\mathbf{x}}^{2n}; F^c)_i\}$. Here we define the decoy-enhanced saliency score Z_i for each feature \mathbf{x}_i as

$$Z_i = \max(\tilde{E}_i) - \min(\tilde{E}_i). \quad (3)$$

Here, Z_i is determined by the empirical range of the decoy saliency scores. Ideally, important features will have large values and unimportant ones will have small values.

3.5. Theoretical insights

In this section, we analyze the saliency score method in a theoretical fashion. For expedience of exposition, we carry out the theoretical analysis using the vanilla gradient as the base saliency method. In particular, we take a convolutional neural network with the ReLU activation function as an example to discuss why the proposed interpretation method can account for inter-feature dependence while also improving explanatory robustness. It should be noted that, while we conduct our theoretical analysis in the setting of convolutional neural networks (CNNs) with a specific activation function, the conclusions drawn from the theoretical analysis can be extended to other feed-forward neural architectures and other activation functions (See Supplementary Section S4 for more details).

Consider a CNN with L hidden blocks, with each layer ℓ containing a convolutional layer with a filter of size $\sqrt{s_\ell} \times \sqrt{s_\ell}$ and a max pooling layer with pooling size $\sqrt{s_\ell} \times \sqrt{s_\ell}$. (We set the pooling size the same as the kernel size in each block for simplicity.) The input to this CNN is $\mathbf{x} \in \mathbb{R}^d$, unrolled from a $\sqrt{d} \times \sqrt{d}$ matrix. Similarly, we also unroll each convolutional filter into $\mathbf{g}_\ell \in \mathbb{R}^{s_\ell}$, where \mathbf{g}_ℓ is indexed as $(\mathbf{g}_\ell)_j$ for $j \in \mathcal{J}_\ell$. Here, \mathcal{J}_ℓ corresponds to the index shift in matrix form from the top-left to bottom-right element. For example, a 3×3 convolutional filter (*i.e.*, $s_\ell = 9$) is indexed by $\mathcal{J}_\ell = \{-\sqrt{d}-1, -\sqrt{d}, -\sqrt{d}+1, -1, 0, 1, \sqrt{d}-1, \sqrt{d}, \sqrt{d}+1\}$. The output of the network is the probability vector $\mathbf{p} \in \mathbb{R}^C$ generated by the softmax function, where C is the total number of classes. Such a network can be represented as

$$\begin{aligned} \mathbf{m}_\ell &= \text{pool}(\text{relu}(\mathbf{g}_\ell * \mathbf{m}_{\ell-1})) \text{ for } \ell = 1, 2, 3, \dots, L, \\ \mathbf{o} &= \mathbf{W}_{L+1}^T \mathbf{m}_L + \mathbf{b}_{L+1}, \\ \mathbf{p} &= \text{softmax}(\mathbf{o}), \end{aligned}$$

where $\text{relu}(\cdot)$ and $\text{pool}(\cdot)$ indicate the ReLU and pooling operators, $\mathbf{m}_\ell \in \mathbb{R}^{d_\ell}$ is the output of the block ℓ ($\mathbf{m}_0 = \mathbf{x}$), and $(\mathbf{g}_\ell * \mathbf{m}_{\ell-1}) \in \mathbb{R}^{d_{\ell-1}}$ represents a convolutional operation on that block. We assume for simplicity that the convolution retains the input shape.

Consider an input \mathbf{x} and its decoy $\tilde{\mathbf{x}}$, generated by swapping features in \mathcal{K} . For each feature $i \in \mathcal{K}$, we have the following theorem for the decoy-enhanced saliency score Z_i :

Theorem 1. In the aforementioned setting, Z_i is bounded by

$$\left| Z_i - \frac{1}{2} \left| \sum_{k \in \mathcal{K}} (\tilde{\mathbf{x}}_k^+ - \tilde{\mathbf{x}}_k^-) (\mathbf{H}_{\mathbf{x}})_{k,i} \right| \right| \leq C_1. \quad (4)$$

Here, $C_1 > 0$ is a bounded constant and $\mathbf{H}_{\mathbf{x}}$ is the Hessian of $F^c(\mathbf{x})$ on \mathbf{x} where $(\mathbf{H}_{\mathbf{x}})_{i,k} = \frac{\partial^2 F^c}{\partial \mathbf{x}_i \partial \mathbf{x}_k}$. $\tilde{\mathbf{x}}^+$ and $\tilde{\mathbf{x}}^-$ refer to the decoy that maximizes and minimizes $E(\tilde{\mathbf{x}}; F^c)$, respectively. Theorem 1 implies that the proposed saliency score is determined by the second-order Hessian $((\mathbf{H}_{\mathbf{x}})_{i,k})$ in the same swappable feature set. The score explicitly models the feature dependencies in the swappable feature set via this second-order Hessian, potentially capturing meaningful patterns such as edges, texture, etc.

In addition to enabling representation of inter-feature dependence, Theorem 1 sheds light on the robustness of the proposed saliency score against adversarial attack. To illustrate the robustness improvement of our method, we introduce the following proposition.

Proposition 1. Given an input \mathbf{x} and the corresponding adversarial sample $\hat{\mathbf{x}}$, if both $|\mathbf{x}_i - \tilde{\mathbf{x}}_i| \leq C_2 \delta_i$ and $|\hat{\mathbf{x}}_i - \tilde{\mathbf{x}}_i| \leq C_2 \delta_i$ can be obtain where $C_2 > 0$ is a bounded constant and $\delta_i = |E(\hat{\mathbf{x}}, F)_i - E(\mathbf{x}, F)_i|$, then the following relation can be guaranteed.

$$|(Z_{\hat{\mathbf{x}}})_i - (Z_{\mathbf{x}})_i| \leq |E(\hat{\mathbf{x}}, F)_i - E(\mathbf{x}, F)_i|. \quad (5)$$

Given an adversarial sample $\hat{\mathbf{x}}$ (*i.e.*, the perturbed \mathbf{x}), we say a saliency method is not robust against $\hat{\mathbf{x}}$ if the deviation of the corresponding explanation $\delta_i = |E(\hat{\mathbf{x}}, F)_i - E(\mathbf{x}, F)_i|$ (for all $i \in \{1, 2, \dots, d\}$) is large. According to the proposition above, we can easily discover that the deviation of our decoy-enhanced saliency score is always no larger than that of other saliency methods when a certain condition is satisfied. This indicates that, when the condition holds, our saliency method can guarantee a stronger resistance to the adversarial perturbation. To ensure the conditions $|\mathbf{x}_i - \tilde{\mathbf{x}}_i| \leq C_2 \delta_i$ and $|\hat{\mathbf{x}}_i - \tilde{\mathbf{x}}_i| \leq C_2 \delta_i$, we can further introduce the corresponding condition as a constraint to Equation (2). In the following section, without further clarification, the saliency scores used in our evaluation are all derived with this constraint imposed. The proof and in-depth analysis of Theorem 1 and Proposition 1 can be found in the Supplementary Section S2 and S3.

4. Experiments

To evaluate the effectiveness of DANCE, we perform extensive experiments on deep learning models that target three

tasks: image classification, sentiment analysis, and network intrusion detection. Our results suggest that that DANCE, in conjunction with state-of-the-art saliency methods, makes already good saliency maps even more intuitively coherent. DANCE also quantitatively achieves better alignment to truly important features and demonstrates stronger robustness to adversarial manipulation. The description of the datasets and experimental setup can be found in the Supplementary Section S5.

4.1. Baseline methods

We applied DANCE in conjunction with three state-of-the-art saliency methods: vanilla gradient, integrated gradient, and SmoothGrad. (See Supplementary Section S9 for results from more saliency methods such as ExpGrad (Sturmfels et al., 2020), VarGrad (Hooker et al., 2019), and Grad-CAM (Selvaraju et al., 2016).) As claimed in Section 2, a prerequisite of saliency methods is the dependency on the predictor. To confirm that the conjunction with DANCE does not violate this prerequisite, we carried out a sanity check on the ImageNet dataset. The results (Supplementary Section S6) show that our method does indeed depend on the predictor.

One significant challenge when comparing different saliency methods is that each method produces a raw saliency map with its own distribution. Therefore, to facilitate a fair comparison among all methods, we used a consistent post-processing scheme to normalize all methods. Specifically, we selected the top- K normalization, *i.e.*, constructing a binary saliency map by retaining only the top- K features ranked by each method. We then set the saliency value of the selected features equal to 1 and the remaining features equal to 0. Here we chose K as the top 20% of all features, and we show that our results are robust to variation in the choice of K in Supplementary Section S11. It is worth mentioning that in this paper we do not consider another common normalization scheme, 0-1 normalization (*i.e.*, linearly rescaling the saliency values to the range $[0, 1]$), because 0-1 normalization leads to a biased estimation of the evaluation metric (See Section 4.2).

4.2. Evaluation metric

Intuitively, we prefer a saliency method that highlights features that align closely with the predictions (*e.g.*, highlights the object of interest in an image or the words indicating the sentiment of the sentence). To measure how well a saliency map achieves qualitative coherence, we use the fidelity metric (Dabkowski & Gal, 2017), defined as

$$SF(E(\cdot; F^c), \mathbf{x}) = -\log F^c(E(\mathbf{x}; F^c) \circ \mathbf{x}) \quad (6)$$

where c indicates the predicted class of input \mathbf{x} , and $E(\mathbf{x}; F^c)$ is the top- K -retained binary saliency map described above. $E(\mathbf{x}; F^c) \circ \mathbf{x}$ performs entry-wise multi-

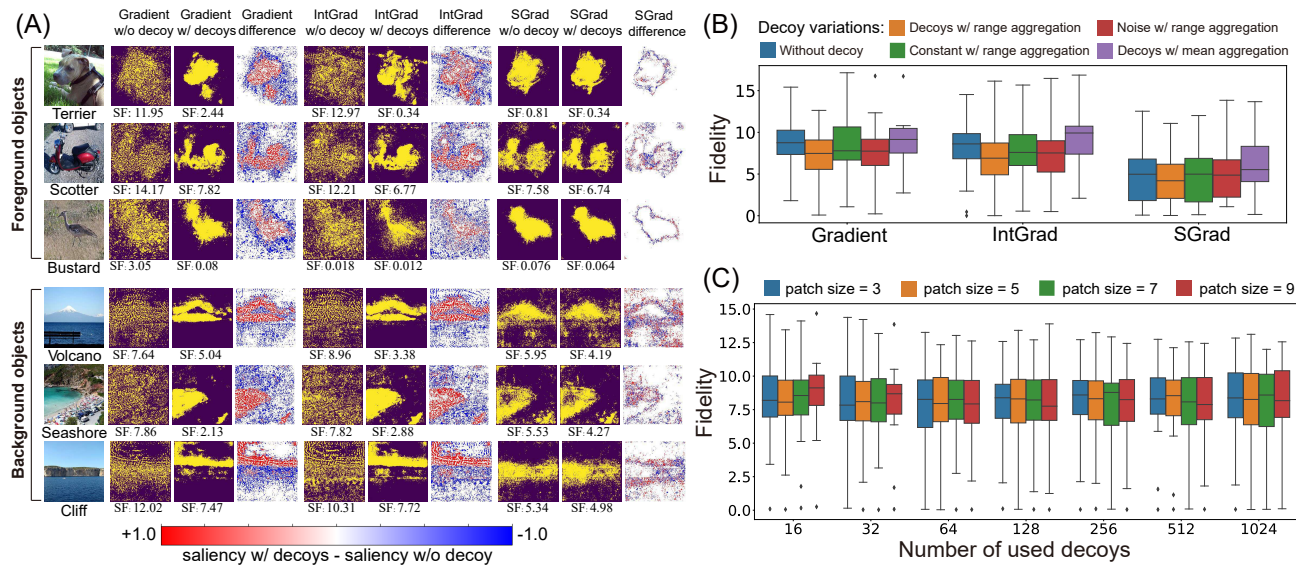


Figure 2. **Performance evaluation on ImageNet.** (A) Visualization of saliency maps on foreground and background objects. (B) Fidelity comparison of original saliency method (i.e., "Without decoys"), our method (i.e., "Decoys w/ range aggregation"), and its alternatives: replacing the decoy generation (Equation (2)) with constant perturbation (i.e., "Constant w/ range aggregation") or noise perturbation (i.e., "Noise w/ range aggregation"); replacing the decoy aggregation (Equation (3)) with mean aggregation (i.e., "Decoys w/ mean aggregation"). See Supplementary Section S14 for more statistics about the performance differences between our method and the baselines. (C) Performance with regard to variant patch size and different numbers of decoys.

plication between $E(\mathbf{x}; F^c)$ and \mathbf{x} , encoding the overlap between the object of interest and the concentration of the saliency map. The rationale behind this metric is that, by viewing the saliency score of a feature as its contribution to the predicted class, a good saliency method will highlight more important features and thus give rise to higher predicted class scores and lower metric values.

Note that, to guarantee a fair comparison among different saliency methods, it is important to retain the same number of important features for evaluation. Without such a scheme, pathologic cases such as $E(\mathbf{x}; F^c) = \mathbf{1}$ (i.e., all saliency values equal to 1) would lead to highest fidelity score unexpectedly, which may be particularly problematic for alternative scheme such as 0-1 normalization.

4.3. Performance in various applications

4.3.1. PERFORMANCE ON THE IMAGENET DATASET

To evaluate the effectiveness of DANCE, we first applied DANCE to randomly sampled images from the ImageNet dataset (Russakovsky et al., 2015), with a pretrained VGG16 model (Simonyan & Zisserman, 2014) (See Supplementary S7 for the applicability to diverse CNN architectures such as AlexNet (Krizhevsky et al., 2012) and ResNet (He et al., 2016)). The 3×3 image patches are treated as swappable

features in generating decoys. A side-by-side comparison (Figure 2(A)) suggests that decoys consistently help to reduce noise and produce more visually coherent saliency maps. For example, the original integrated gradient method highlights the region of the dog's head in a scattered format, which is also revealed by the difference plot. In contrast, the decoy-enhanced integrated gradient method not only highlights the missing body but also identifies the dog's head with more details such as ears, cheek, and nose (See Supplementary Section S13 for more visualization examples). The visual coherence is also quantitatively supported by the saliency fidelity score.

To further evaluate the necessity of the two steps in our method (i.e., decoy generation and aggregation), we carried out a control experiment by replacing each step with alternatives. Specifically, as alternatives to the decoy generation, we used an image in which all pixel values are either replaced with a single mean pixel value or contaminated with Gaussian white noise. For the decoy aggregation, we calculated the mean saliency score as the alternative. As shown in Figure 2(B), our method, which incorporate both steps, yields the best performance. This validates the effectiveness of our two-step approach.

Thirdly, to evaluate the computational efficiency of DANCE, we carried out a fidelity comparison with respect to the num-

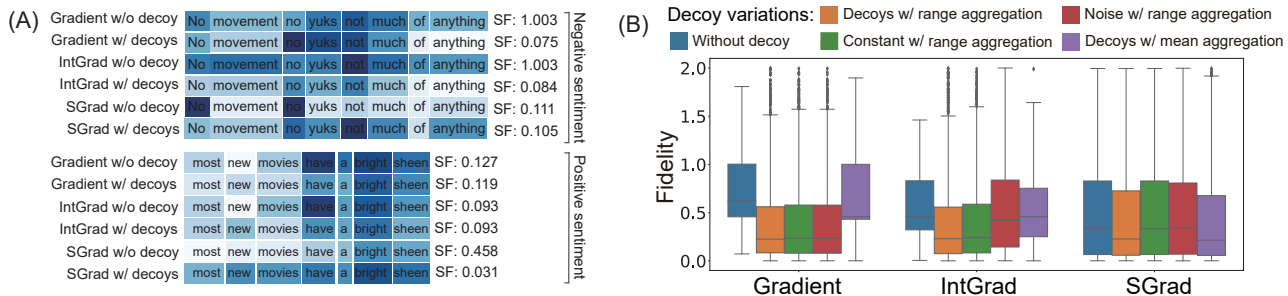


Figure 3. Results obtained from the SST dataset. (A) Visualization of saliency maps in each word, where the normalized saliency values are shown for better distinction. (B) Fidelity comparison of the original saliency method, our method, and its alternatives. Here, the alternative methods represent the practice of replacing the decoy generation (Equation (2)) with constant perturbation or noise perturbation as well as the practice of replacing the decoy aggregation (Equation (3)) with mean aggregation. See Supplementary Section S14 for more statistics about the performance differences between our method and the baselines.

ber of decoys to optimize. As discussed in Section 3.3, multiple swappable patches are aggregated into one combined patch mask for computational efficiency. Consequently, the mask multiplicity (*i.e.*, the number of swappable patches per mask) is inversely proportional to the number of decoys to optimize. Figure 2(C) shows that our method achieves stable fidelity scores across a wide range of decoy numbers. Furthermore, as shown in Supplementary Section S10, the computational cost to optimize a single decoy in DANCE is negligible compared to even the fastest vanilla gradient-based saliency method. This analysis result confirms that our system could give reasonably good saliency maps without introducing too much computational cost.

Finally, In Supplementary Section S11, we run a sensitivity test on other hyper-parameters (*i.e.*, the swappable feature size P , the targeted network layer ℓ , and the initial Lagrange multiplier λ). The results show that our method is insensitive to substantial variation of these hyperparameters. This is an important property because users do not need to extensively tune the hyper-parameters when using our method.

4.3.2. PERFORMANCE ON THE STANFORD SENTIMENT TREEBANK (SST) DATASET

To further evaluate the effectiveness of DANCE, we applied the method to randomly sampled sentences from the Stanford Sentiment Treebank (SST) (Russakovsky et al., 2015). We trained a two-layer CNN (Kim, 2014) which takes the pretrained word embeddings as input (Pennington et al., 2014) (See Supplementary Section S6 for more details about the experimental setup). As suggested by Guan et al. (2019), the average saliency value of all dimensions of a word embedding is regarded as the word-level saliency value. The embeddings of the words are treated as swappable features when generating decoys. As shown in Figure 3(A), a side-

by-side comparison suggests that our method consistently helps to produce semantically more meaningful saliency maps. For example, in a sentence with negative sentiment, keywords associated with negation, such as “no” and “not,” are more strongly highlighted by decoy-enhanced saliency methods. The semantic coherence is also quantitatively supported by the saliency fidelity (Figure 3(B)). We also tested the alternatives mentioned above: constant (replacing the decoy generation with the mean embedding of the whole dictionary) and noise perturbation with range aggregation, and decoys with mean aggregation. Figure 3(B) shows that our method outperforms these alternatives.

To demonstrate the effectiveness of DANCE on models other than CNNs, we carried out experiments on a multi-layer perceptron trained with a network intrusion dataset. The results (Supplementary Section S8) are consistent with those on CNNs, thereby confirming our method’s applicability to non-CNN architectures.

4.4. Robustness to adversarial attacks

An important design philosophy of DANCE is to model the variation of an input sample originating from either sensor noise or unknown perturbations by using decoys. We therefore hypothesized that DANCE may be particularly robust to adversarial manipulations of images. To test this hypothesis, we evaluated the robustness of our method to adversarial manipulations of images subject to three popular attacks (Ghorbani et al., 2017): (1) the top- k attack, which seeks to decrease the scores of the top k most important features, (2) the target attack, which aims to increase the importance of a pre-specified region in the input image, and (3) the mass-center attack, which aims to spatially change the mass center of the original saliency map. Here, we specify the bottom-right 4×4 region of the original image




Figure 4. **Robustness to adversarial attacks on images.** (A) Visualization of saliency maps under adversarial attacks. (B)–(D) The decoy-enhanced saliency score is compared to the original saliency score under adversarial attacks, evaluated by sensitivity. See Supplementary Section S14 for more statistics about the performance differences between our method and the baselines.

for the target attack and select $k = 5000$ in the top- k attack (See Supplementary Section S6 for detailed setups). We use the sensitivity metric (Alvarez-Melis & Jaakkola, 2018) to quantify the robustness of a saliency method E to attack, defined as:

$$SS(E(\cdot, F^c), \mathbf{x}, \hat{\mathbf{x}}) = \frac{\|E(\mathbf{x}, F^c) - E(\hat{\mathbf{x}}, F^c)\|_2}{\|\mathbf{x} - \hat{\mathbf{x}}\|_2} \quad (7)$$

where $\hat{\mathbf{x}}$ is the perturbed image of \mathbf{x} . A small SS value means that similar inputs do not lead to substantially different saliency maps. As shown in Figure 4(A), a side-by-side comparison suggests that decoys consistently yield low sensitivity scores and help to produce more visually coherent saliency maps, mitigating the impact of various adversarial attacks (See the Supplementary material for more examples). The visual coherence and robustness to adversarial attacks are also quantitatively supported by Figure 4(B)–(D).

5. Discussion and conclusion

In this work, we propose DANCE, a method for computing, from a given saliency method, decoy-enhanced saliency scores that yield more accurate and robust saliency maps. We formulate the decoy generation as an optimization problem, applicable to diverse DNN architectures. We demonstrate the superior performance of our method relative to three standard saliency methods, both qualitatively and quan-

titatively, even in the presence of various adversarial perturbations to the image. From a theoretical perspective, by deriving a closed-form solution, we show that the proposed score can provably compensate for the limitations of existing saliency methods by reflecting the joint effects from other dependent features and maintaining robustness to adversarial perturbations. We also demonstrate the computational efficiency of DANCE, and we show that the cost to optimize a single decoy is small, indicating that our technique can improve upon existing saliency methods without introducing too much computational overhead.

This work points to several promising directions for future research. First, DANCE is designed for non-linear models such as feedforward DNNs which are most in need of interpretation. Future work will explore the extension of our method to other models (*e.g.*, linear model and recurrent neural networks) and to inputs with categorical or discrete features. Second, recent work (Etmann et al., 2019; Chen et al., 2019c; Chalasani et al., 2020) shows that adversarial training can improve a DNN’s interpretability. It is worth exploring whether DANCE could further enhance the quality of saliency maps derived from these adversarially retrained classifiers. Finally, a promising direction could be reframing interpretability as hypothesis testing and using decoys to deliver a set of salient features, subject to false discovery rate control at some pre-specified level (Burns et al., 2019; Lu et al., 2018).

Acknowledgments

We would like to thank the anonymous reviewers and Meta reviewer for their helpful comments. This project was supported in part by NSF grant CNS-1718459, by NSF grant CNS-1954466, and by ONR grant N00014-20-1-2008.

References

- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., and Kim, B. Sanity checks for saliency maps. In *Proc. of NeurIPS*, 2018.
- Alvarez-Melis, D. and Jaakkola, T. S. Towards robust interpretability with self-explaining neural networks. In *Proc. of NeurIPS*, 2018.
- Ancona, M., Ceolini, E., Öztireli, C., and Gross, M. Towards better understanding of gradient-based attribution methods for deep neural networks. In *Proc. of ICLR*, 2018.
- Barber, R. F. and Candès, E. J. Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 2015.
- Binder, A., Montavon, G., Lapuschkin, S., Müller, K.-R., and Samek, W. Layer-wise relevance propagation for neural networks with local renormalization layers. In *Proc. of ICANN*, 2016.
- Burns, C., Thomason, J., and Tansey, W. Interpreting black box models via hypothesis testing. *arXiv:1904.00045*, 2019.
- Chalasanani, P., Chen, J., Chowdhury, A. R., Wu, X., and Jha, S. Concise explanations of neural networks using adversarial training. In *Proc. of ICML*, 2020.
- Chang, C.-H., Creager, E., Goldenberg, A., and Duvenaud, D. Explaining image classifiers by counterfactual generation. In *Proc. of ICLR*, 2019.
- Chen, C., Li, O., Tao, D., Barnett, A., Rudin, C., and Su, J. K. This looks like that: deep learning for interpretable image recognition. In *Proc. of NeurIPS*, 2019a.
- Chen, J., Song, L., Wainwright, M. J., and Jordan, M. I. Learning to explain: An information-theoretic perspective on model interpretation. In *Proc. of ICML*, 2018.
- Chen, J., Song, L., Wainwright, M. J., and Jordan, M. I. L-shapley and c-shapley: Efficient model interpretation for structured data. In *Proc. of ICLR*, 2019b.
- Chen, J., Wu, X., Rastogi, V., Liang, Y., and Jha, S. Robust attribution regularization. In *Proc. of NeurIPS*, 2019c.
- Dabkowski, P. and Gal, Y. Real time image saliency for black box classifiers. In *Proc. of NeurIPS*, 2017.
- Etmann, C., Lunz, S., Maass, P., and Schönlieb, C.-B. On the connection between adversarial robustness and saliency map interpretability. *arXiv preprint arXiv:1905.04172*, 2019.
- Fan, L., Zhao, S., and Ermon, S. Adversarial localization network. In *Proc. of NeurIPS LLD Workshop*, 2017.
- Fong, R. C. and Vedaldi, A. Interpretable explanations of black boxes by meaningful perturbation. In *Proc. of ICCV*, 2017.
- Ghorbani, A., Abid, A., and Zou, J. Interpretation of neural networks is fragile. *arXiv:1710.10547*, 2017.
- Ghorbani, A., Wexler, J., Zou, J. Y., and Kim, B. Towards automatic concept-based explanations. In *Proc. of NeurIPS*, 2019.
- Goyal, Y., Wu, Z., Ernst, J., Batra, D., Parikh, D., and Lee, S. Counterfactual visual explanations. *Proc. of ICML*, 2019.
- Guan, C., Wang, X., Zhang, Q., Chen, R., He, D., and Xie, X. Towards a deep and unified understanding of deep neural models in nlp. In *Proc. of ICML*, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proc. of CVPR*, 2016.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. In *Proc. of ICLR*, 2019.
- Hooker, S., Erhan, D., Kindermans, P.-J., and Kim, B. A benchmark for interpretability methods in deep neural networks. In *Proc. of NeurIPS*, 2019.
- Ikeno, K. and Hara, S. Maximizing invariant data perturbation with stochastic optimization. *arXiv preprint arXiv:1807.05077*, 2018.
- Kim, Y. Convolutional neural networks for sentence classification. *Proc. of EMNLP*, 2014.
- Kindermans, P.-J., Hooker, S., Adebayo, J., Alber, M., Schütt, K. T., Dähne, S., Erhan, D., and Kim, B. The (Un) reliability of saliency methods. *arXiv:1711.00867*, 2017.
- Koh, P. W. and Liang, P. Understanding black-box predictions via influence functions. *Proc. of ICML*, 2017.
- Koh, P. W. W., Ang, K.-S., Teo, H., and Liang, P. S. On the accuracy of influence functions for measuring group effects. In *Proc. of NeurIPS*, 2019.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Proc. of NeurIPS*, 2012.

- Levine, A., Singla, S., and Feizi, S. Certifiably robust interpretation in deep learning. *arXiv preprint arXiv:1905.12105*, 2019.
- Lipton, Z. C. The mythos of model interpretability. *arXiv:1606.03490*, 2016.
- Lu, Y., Fan, Y., Lv, J., and Noble, W. S. DeepPINK: reproducible feature selection in deep neural networks. In *Proc. of NeurIPS*, 2018.
- Lundberg, S. M. and Lee, S.-I. A unified approach to interpreting model predictions. In *Proc. of NeurIPS*, 2017.
- Moosavi-Dezfooli, S.-M., Fawzi, A., Fawzi, O., and Frossard, P. Universal adversarial perturbations. In *Proc. of CVPR*, 2017.
- Nie, W., Zhang, Y., and Patel, A. A theoretical explanation for perplexing behaviors of backpropagation-based visualizations. In *Proc. of ICML*, 2018.
- Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *Proc. of EMNLP*, 2014.
- Ribeiro, M. T., Singh, S., and Guestrin, C. Why should i trust you?: Explaining the predictions of any classifier. In *Proc. of KDD*, 2016.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 2015.
- Schulz, K., Sixt, L., Tombari, F., and Landgraf, T. Restricting the flow: Information bottlenecks for attribution. 2020.
- Selvaraju, R. R., Das, A., Vedantam, R., Cogswell, M., Parikh, D., and Batra, D. Grad-cam: Visual explanations from deep networks via gradient-based localization. *arXiv:1611.07450*, 2016.
- Shrikumar, A., Greenside, P., and Kundaje, A. Learning important features through propagating activation differences. In *Proc. of ICML*, 2017.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- Simonyan, K., Vedaldi, A., and Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv:1312.6034*, 2013.
- Singla, S., Wallace, E., Feng, S., and Feizi, S. Understanding impacts of high-order loss approximations and features in deep learning interpretation. *arXiv:1902.00407*, 2019.
- Sixt, L., Granz, M., and Landgraf, T. When explanations lie: Why many modified bp attributions fail. In *Proc. of ICML*, 2020.
- Smilkov, D., Thorat, N., Kim, B., Viégas, F., and Wattenberg, M. Smoothgrad: removing noise by adding noise. *arXiv:1706.03825*, 2017.
- Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.
- Sturmfels, P., Lundberg, S., and Lee, S.-I. Visualizing the impact of feature attribution baselines. *Distill*, 2020.
- Sundararajan, M., Taly, A., and Yan, Q. Axiomatic attribution for deep networks. In *Proc. of ICML*, 2017.
- Toneva, M. and Wehbe, L. Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain). In *Proc. of NeurIPS*, 2019.
- Yeh, C.-K., Kim, J., Yen, I. E.-H., and Ravikumar, P. K. Representer point selection for explaining deep neural networks. In *Proc. of NeurIPS*, 2018.
- Yousefzadeh, R. and O’Leary, D. P. Interpreting neural networks using flip points. *arXiv preprint arXiv:1903.08789*, 2019.
- Zeiler, M. D. and Fergus, R. Visualizing and understanding convolutional networks. In *Proc. of ECCV*, 2014.
- Zoła, K., Geras, K. J., and Cho, K. Classifier-agnostic saliency map extraction. In *Proceedings of AAAI*, 2019.