
Improving Attribution Methods by Learning Submodular Functions

Piyushi Manupriya

Tarun Ram Menta

J. Saketha Nath

Vineeth N Balasubramanian

Indian Institute of Technology, Hyderabad

Abstract

This work explores the novel idea of learning a submodular scoring function to improve the specificity/selectivity of existing feature attribution methods. Submodular scores are natural for attribution as they are known to accurately model the principle of diminishing returns. A new formulation for learning a deep submodular set function that is consistent with the real-valued attribution maps obtained by existing attribution methods is proposed. The final attribution value of a feature is then defined as the marginal gain in the induced submodular score of the feature in the context of other highly attributed features, thus decreasing the attribution of redundant yet discriminatory features. Experiments on multiple datasets illustrate that the proposed attribution method achieves higher specificity along with good discriminative power. The implementation of our method is publicly available at <https://github.com/Piyushi-0/SEA-NN>.

1 INTRODUCTION

Deep neural networks (DNNs) have shown state-of-the-art performance in diverse application domains, including complex tasks such as image recognition, video synthesis, speech-to-text conversion and autonomous navigation, to name a few. While advancements in deep learning have led to increasing accuracy scores, this has often been at the expense of the interpretability of the neural network’s decisions. Our work focuses on attribution algorithms for interpreting neural networks. These algorithms output an attribution score heatmap that represents the feature-wise contribution of each input feature towards the

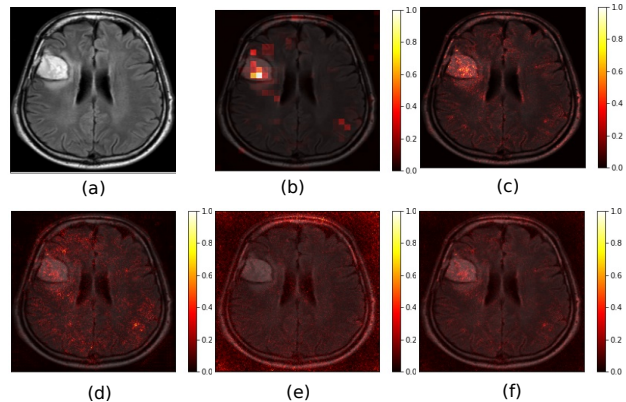


Figure 1: Example image and attribution maps overlaid on the image from the Brain Tumor Dataset: (a) Image with brain tumor; (b) SEA-NN (proposed) map; (c) Integrated Gradient map; (d) Vanilla Gradient map; (e) Smooth Integrated Gradient map; (f) Agg-Mean map of (c)-(e) maps. More such examples are provided in our results section.

prediction, a topic that has attracted significant interest in recent years [Zeiler and Fergus, 2014, Springenberg et al., 2015, Selvaraju et al., 2017, Chattopadhyay et al., 2018, Shrikumar et al., 2017, Lundberg and Lee, 2017, Sundararajan et al., 2017, Chattopadhyay et al., 2019].

As shown in Fig 1, attribution algorithms provide a normalized score for each input feature (pixels in case of image-based problems, where such methods are extensively used), which denotes its contribution to the prediction on the given input. Our approach in this work is based on the observation that attribution heatmaps obtained by existing methods are not sharply focused around the discriminatory pixels and are often spread out across the image (see Fig 1). In other words, though existing algorithms achieve good sensitivity, they do not seem to achieve high specificity/selectivity. In applications like brain tumor detection where the region of interest usually lies in a very small area, specificity can be especially critical. The key idea in our work is to employ a monotone submodular set function to score – and subsequently en-

semble – known heatmaps, such that the scores for the most specific of the given maps saturate (i.e. adding any more pixels to the map will not provide benefit).

However, such a submodular function is neither known nor designed in existing literature. We hence propose a novel algorithm that *learns* one based on the heatmaps with existing attribution algorithms. Learning such a submodular function across multiple attribution algorithms provides us a mechanism to combine these methods, as well as provides attribution maps that are consistent across methods. The works closest to ours include [Rieger and Hansen, 2019], which performs pixel-wise averaging of attribution maps, and [Elenberg et al., 2017], which uses a notion of weak submodularity to select top- k important pixels in a streaming setting. We compare against [Rieger and Hansen, 2019] in this work, and show that our ensembled attribution map guided by a learned submodular scoring function outperforms simple linear averaging. [Elenberg et al., 2017] however focuses on interpretability in a streaming setting, which is different from our objective. Besides, they focus on selection of a discrete subset of features, while we focus on scoring the heatmaps. Also, as noted by them, the function that they use is neither expected to be monotone nor submodular.

Based on our new submodular scoring function, we present our model-agnostic attribution algorithm **Submodular Ensembled Attribution for Neural Networks (SEA-NN)**. This algorithm scores features according to their marginal score gains in presence of other critical features, rather than attributing according to the raw scores. For a given input-output pair on a given model, our proposed methodology assigns attribution scores based on the submodular scoring function learned using heatmaps of multiple algorithms. This encourages specificity as the redundant, yet discriminatory, features get lower attribution scores. One could also view our methodology as combining multiple attribution algorithms for a given DNN model which is especially useful as the goodness of attribution scores in existing methods can vary based on one’s perspective. Each such attribution method has different biases in defining what may or may not be important in an input. Moreover, [Rieger and Hansen, 2019] showed that aggregating attribution maps also results in a more robust attribution map, as adversarial attacks do not transfer well across different attribution methods. Additionally, if each of our input attribution method satisfies the axioms of attribution defined in [Sundararajan et al., 2017], our method either inherits them or can be easily modified to satisfy most of them. Importantly, we note that the computational

overhead of learning the submodular score and the ensembled attribution is not prohibitive and is limited to a small number of evaluation and backpropagation calls, which can be efficiently implemented. More discussion of our method’s computational effort is discussed in Section 3, with quantitative analysis in the supplementary materials.

The key contributions of our work can be summarized as follows:

- To achieve the aforementioned objectives of ensembling known attribution heatmaps, we propose a new formulation for learning a submodular set function from real-valued heatmaps. Typical supervision studied in literature for learning submodular functions is of the form of a collection of “important” subsets rather than real-valued importance maps. While proposed in the context of ensembling attribution maps, the proposed formulation may be of independent interest in other applications of submodular learning.
- We propose a new Submodular Ensembled Attribution for Neural Networks (SEA-NN) algorithm, which improves over existing attribution methods in several ways – importantly, their specificity. Furthermore, this method provides a means to consider the marginal gain of a pixel’s attribution in the presence of another pixel (or set of pixels), an aspect that is generally ignored in existing attribution methods.
- We conducted a comprehensive suite of experiments on standard benchmark image datasets such as FMNIST, CUB, Tiny Imagenet as well as the Brain Tumor Detection dataset where specificity is highly desirable. Our experiments and ablation studies validate the usefulness of the proposed method from a qualitative as well as quantitative perspective.

2 RELATED WORK

The last few years have seen a significant number of efforts in enhancing the explainability of machine learning methods – in particular, the predictions of DNN models. These efforts include local and global methods, model-agnostic and model-specific methods, causal and non-causal methods, and so on [Molnar, 2019, Lecue et al., 2021, Tjoa and Guan, 2020]. In this work, our efforts are based on gradient-based attribution methods, which continue to be among the most popular family of methods, especially on image-based problems. A brief overview of such methods is presented below.

Vanilla Gradient (VG) [Simonyan et al., 2014] and Deconvnet [Zeiler and Fergus, 2014] were among

the earliest methods proposed for attribution in images. Many methods developed upon the notion of gradient-based attribution including Guided Backprop [Springenberg et al., 2015], Grad-CAM [Selvaraju et al., 2017] and Grad-CAM++ [Chattopadhyay et al., 2018], which introduced newer ways of computing or using the gradient to get class-discriminative attribution maps. Integrated Gradient (IG) [Sundararajan et al., 2017] was a seminal effort that improved upon gradient-based methods by varying the input x on a linear path from baseline \bar{x} to x and averaging gradients obtained at all these inputs. In order to produce less noisy attribution maps, Smooth Integrated Grad (SG) [Smilkov et al., 2017] averaged attribution maps of IG corresponding to inputs perturbed with Gaussian noise. Several other attribution methods were introduced. Layer-wise Relevant Propagation (LRP) [Bach et al., 2015] and DeepLIFT [Shrikumar et al., 2017] proposed attributions with discrete gradients using a modified form of backpropagation. There has also been some interesting study on relationships between these methods. [Ancona et al., 2018] showed that LRP- ϵ rule is related to *input \circ gradient* (element-wise product of input and gradient) and IG is also strongly related to methods like *input \circ gradient* and DeepLIFT.

However, evidently, each of the aforementioned methods focuses on its own perspective to gradient-based attribution, with its own biases. Very little work has been carried out in ensembling such attribution methods. [Rieger and Hansen, 2019] proposed a simple pixel-wise averaging of attribution maps; [Elenberg et al., 2017] used a notion of weak submodularity to select top- k important pixels in a streaming setting; [Bhatt et al., 2020] aggregated attribution maps from the viewpoint of Shapley values, and focused on instance-based aggregation from a data point’s neighbors. The objectives of each of these efforts were different from ours – [Rieger and Hansen, 2019] is the closest to ours, which we compare with in our experiments.

In this work, we learn a new submodular scoring function using heatmaps of given known gradient-based attribution methods. The use of submodularity – and thus marginal gains – in our attribution algorithm provides context-dependent attribution scores (i.e. score of a pixel attribution given other pixels) which differentiates it from most existing methods that compute attribution scores independently for each feature. State-of-the-art DNNs are designed to leverage feature inter-dependencies which are not considered in most popular gradient-based methods. [Tsang et al., 2020] recently developed a framework to show how interactions between features attribute to the predictions

of DNNs, supporting the need for such a context-dependent attribution score. In terms of enhancing specificity of attribution maps, [Wagner et al., 2019] proposed a method for fine-grained visual explanations using an adversarial defense technique. Our objectives are once again different, and our formulation of learning the submodular scoring function can be of independent interest; besides, such a method can also provide an input heatmap for our proposed SEA-NN method.

3 PROPOSED METHODOLOGY

Background. Our work builds upon properties of submodular functions. Submodular functions are a special kind of discrete functions, characterized by the diminishing returns property, and appear naturally in many discrete maximization problems like clustering, sensor placement and document summarization [Krause and Golovin, 2014]. For a set function $f : 2^V \rightarrow \mathbb{R}$ defined on a ground set V , the marginal gain on adding an element e to the set A can be defined as $f(e|A) = f(A \cup e) - f(A)$. f is said to be *submodular* if for any $e \notin B$, for all $A \subseteq V$ and $B \subseteq V$ such that $A \subseteq B$, $f(e|A) \geq f(e|B)$ i.e. the smaller set has a larger gain upon addition of a new element. On the other hand, if both sets have equal marginal gain i.e. $f(e|A) = f(e|B)$, then f is said to be *modular*. In most applications where f acts as a valuation function, f is desired to be non-negative i.e. $f(A) \geq 0$ for all $A \subseteq V$. Additionally, if $f(A) \leq f(B)$ for all $A \subseteq B$, then f is said to be *monotonic*.

Given a neural network with n input features, we assume that we have access to m attribution algorithms that are faithful to it. We denote the heatmap (with normalized values in $[0, 1]$) obtained by the i^{th} algorithm on a given input/image as $\mathcal{H}_i \in [0, 1]^n$. We assume that the algorithms are reasonable baselines (of which we have quite a few methods today such as Integrated Gradients or SmoothGrad); hence, in general, it can be assumed that the features with high attribution values are indeed critical for discriminatory power of the network on a given input image.

Learning the Scoring Function. In order to ensemble the provided attribution maps, we first define a scoring function, $f : 2^V \rightarrow \mathbb{R}^+$, where V is the set of pixels. This function assigns a non-negative score to a subset of pixels chosen from the input image consisting of n pixels. We denote f ’s real-valued extension by $f_{ext} : [0, 1]^n \mapsto \mathbb{R}^+$ that scores attribution heatmaps. The desirable properties of such a scoring function, given m attribution heatmaps as input, are: (1) The score for pixels belonging to baseline heatmaps must be high; (2) The score for a heatmap with high specificity/selectivity is high, i.e. further increasing attribution values in a selective heatmap does not increase

the score significantly; and (3) A discriminatory feature’s value in presence of other useful features must be lower than that in their absence. Note that properties (2) and (3) are critical, as attributions based on such scoring functions will promote heatmaps with high specificity and low redundancy. Here, we propose to learn an appropriate scoring function for this purpose.

While ensuring property (1) above is straight-forward, properties (2) and (3) suggest that the proposed scoring function should satisfy the diminishing returns property. Accordingly, we model f by monotone submodular functions and f_{ext} by monotone concave functions. This is because submodular functions and their concave extensions are known to accurately model the principle of diminishing returns [Iyer and Bilmes, 2020]. And together with monotonicity, such functions allow us to satisfy Property (2) and facilitate early enough saturation of values. Another pragmatic reason for making this choice is that models inducing monotone submodular functions have been well-studied in recent years [Dolhansky and Bilmes, 2016]. This work [Dolhansky and Bilmes, 2016] proposed a class of functions known as Deep Submodular Functions (DSFs), which can be efficiently learned using appropriate neural network training techniques. We leverage this approach to obtain a submodular score function for our work. (We provide more details on DSFs in the supplementary section.)

Let $f_{\mathbf{w}}$ denote a DSF learned as a neural network (NN), parameterized by network parameters, $\mathbf{w} > 0$. As shown in [Dolhansky and Bilmes, 2016], $f_{\mathbf{w}}$ restricted to binary-valued inputs, $\{0, 1\}^n$, is indeed a monotone submodular function. We also use non-negative, increasing, concave activation functions in our neural network, similar to [Dolhansky and Bilmes, 2016]. More details of the architecture details of this NN are provided in Sec 4 and the supplementary section. We now make a critical observation that when such an $f_{\mathbf{w}}$ is extended to the domain $[0, 1]^n$, it is indeed a concave extension of this submodular function. This is because the activation functions in the network defining the DSF are all non-negative, monotone (increasing), and concave functions. We hence propose to learn the parameters \mathbf{w} by solving the following novel (intermediate) formulation:

$$\min_{\mathbf{w} \geq 0} \frac{\lambda}{2} \|\mathbf{w}\|^2 + \sum_{i=1}^m [f_{\mathbf{w}}(\mathcal{H}^*) - f_{\mathbf{w}}(\mathcal{H}_i)] + (1 - f_{\mathbf{w}}(\mathcal{H}^*))^+ \quad (1)$$

where $(a)^+ \equiv \max(0, a)$, $\lambda > 0$ is a regularization hyperparameter, and \mathcal{H}^* is a heatmap with all attri-

butions as unity. This NN-based DSF is obtained by training using the m attribution maps as input. Since \mathcal{H}^* is the heatmap where any monotone function is maximized, the second term in the objective of Eqn (1) promotes functions/parameters that give high score to pixels in all baseline heatmaps. As a result, the scoring function, $f_{\mathbf{w}}$ will saturate at heatmaps that resemble the most specific heatmaps among the given attribution maps that we ensemble. The last term in the objective of Eqn (1) prevents the trivial solution of $\mathbf{w} = \mathbf{0}$.

Since m , the number of input attribution maps, can be limited in practice, we include further supervision in order to improve the generalization of Eqn (1). To this end, we assume that each input heatmap is binarized using k different hard thresholds (resulting in $k \times m$ inputs to learn the DSF). We denote such binarized heatmaps as $\mathcal{H}_i^j, i \in \{1, \dots, m\}, j \in \{1, \dots, k\}$. Let B_i^j denote the number of features selected in \mathcal{H}_i^j by the binary thresholding. We now propose our final formulation for learning the scoring function as:

$$\min_{\mathbf{w} \geq 0} \frac{\lambda}{2} \|\mathbf{w}\|^2 + \lambda_1 \left(\sum_{i=1}^m \left(f_{\mathbf{w}}(\mathcal{H}^*) - f_{\mathbf{w}}(\mathcal{H}_i) \right) \right) + \lambda_2 \left(\sum_{i=1}^m \sum_{j=1}^k \left(\delta + \max_{|A| \leq B_i^j} f_{\mathbf{w}}(A) - f_{\mathbf{w}}(\mathcal{H}_i^j) \right)^+ \right) \quad (2)$$

where $(a)^+ \equiv \max(0, a)$; $\lambda > 0$, $\lambda_1 > 0$ & $\lambda_2 > 0$ are hyperparameters; and $\delta > 0$ also a hyperparameter controlling the margin. In practice, it is possible to have $(\delta + \max_{|A| \leq B_i^j} f_{\mathbf{w}}(A) - f_{\mathbf{w}}(\mathcal{H}^*))^+ < \delta$ when we are solving the maximization using the approximate greedy algorithm. Hence, we did not need the term to prevent $\mathbf{w} = \mathbf{0}$ as a trivial solution in our implementation. The new loss terms in Eqn (2) promote scores such that the binarized heatmaps are scored high among maps with the same number of selected features. Note that the term $\max_{|A| \leq B_i^j} f_{\mathbf{w}}(A) - f_{\mathbf{w}}(\mathcal{H}_i^j)$ may turn out to be negative whenever the submodular maximization is solved approximately (e.g., when solved via the greedy algorithm). Hence the truncation based on the $(\cdot)^+$ is necessary in the above formulation (Eqn 2).

To the best of our knowledge, formulations like Eqn (2) that jointly learn the submodular function as well as its concave extension are not well-studied in literature. We hence hope that the proposed formulation may be of independent interest in other applications (e.g. weighting samples instead of selecting samples in active learning, imbalanced classification or long-tailed recognition), where early saturating functions need to be induced for similar problems beyond attribution maps.

The proposed formulation (2) can be solved using Mirror Descent or Projected Subgradient Descent. Any subgradient of the objective in Eqn (2) is given by: $\lambda \mathbf{w} + \sum_{i=1}^m [\nabla_{\mathbf{w}} f_w(\mathcal{H}^*) - \nabla_{\mathbf{w}} f_w(\mathcal{H}_i) + \sum_{j=1}^k \rho_i^j (\nabla_{\mathbf{w}} f_w(\bar{A}_i^j) - \nabla_{\mathbf{w}} f_w(\mathcal{H}_i^j))]$, where \bar{A}_i^j is a solution to the inner submodular maximization problem subject to a cardinality constraint, and $\rho_i^j > 0$ iff $\delta + \nabla_{\mathbf{w}} f_w(\bar{A}_i^j) - \nabla_{\mathbf{w}} f_w(\mathcal{H}_i^j) > 0$ and $\rho_i^j = 0$ otherwise. We solve for \bar{A}_i^j using the well-known constant-factor greedy approximation algorithm for maximizing a non-negative monotone submodular function [Nemhauser et al., 1978]. This subgradient can be efficiently computed by backpropagating through the NN modeling the DSF.

Notably, the function learnt, $f_{\mathbf{w}}$, exhibits certain degree of stability to noise. This is because $f_{\mathbf{w}}$, when restricted to the domain $\{0, 1\}^n$, is a non-negative, monotone, submodular function. Such functions are known to be noise-stable [Cheraghchi et al., 2011]. Also, on the compact domain $[0, 1]^n$, f_{ext} is a concave function and hence is Lipschitz continuous (for e.g., see theorem C.4.1 in [Ben-Tal and Nemirovski, 2021]). Hence the proposed scoring function is expected to be fairly robust to noise in baseline attribution maps.

Lemma 1. *Given input x and a classifier neural network function F^c corresponding to class c , $|f_{ext}(\mathcal{H}_i(x, F^c)) - f_{ext}(\mathcal{H}_i(x + \delta, F^c))| \leq L \|\mathcal{H}_i(x, F^c) - \mathcal{H}_i(x + \delta, F^c)\| \forall \mathcal{H}_i, i \in \{1, \dots, m\}$, L being a constant.*

This lemma highlights that the deviation in f_{ext} caused by an additively perturbed input is not arbitrarily high but depends on the deviations that arise from the baseline heatmaps.

We now discuss our algorithm to compute a new heatmap that ensembles baseline heatmaps using this induced scoring function.

Submodular Ensemble Attribution Algorithm.

We begin by noting that high specificity in attribution heatmaps may not be easy to achieve by arbitrary linear/non-linear combinations of the given heatmaps (due to the interactions between pixels, as stated in Sec 2). Thus we propose to learn $f_{\mathbf{w}}$ that attributes features based on their marginal gain in the context of other features (i.e. $f_{\mathbf{w}}$ scores a subset of pixels rather than each individual pixel, which is a key difference from a linear combination of these heatmaps). The proposed attribution algorithm is detailed in Algorithm (1). Our algorithm is designed with the following key features/advantages: (i) It improves as the given heatmaps that we ensemble improve; (ii) It is expected to achieve high specificity; and (iii) By inheritance, the axiom of ‘‘Implementation Invariance’’ defined in [Sundararajan et al., 2017] is satisfied whenever the given heatmaps satisfy the same axiom.

Algorithm 1 Attribution Algorithm

Input: Trained DSF f , Set of features to attribute V
 Initialize feature subset $A = \{\}$
 Initialize $n = |V|$
 Initialize attribution map $G[i] = 0$ for $i = 0, 1, \dots, n$
 $M = \{\arg \max_{v \in V \setminus A} f(v|A)\}$, Pick $e \in M$
while $|A| < n$ and $f(e|A) > 0$ **do**
 for $p \in M$ **do**
 $G[p] = f(e|A)$
 end for
 $A = A \cup \{e\}$; $V = V \setminus M$
 $M = \{\arg \max_{v \in V} f(v|A)\}$, Pick $e \in M$
end while
return G

Many types of other prior information can also be easily encoded as constraints in the proposed formulation (Eqn 2). In particular, one can insist that the proposed attribution satisfies the axioms Sensitivity(a) and/or Sensitivity(b) [Sundararajan et al., 2017] by simply adding appropriate constraints in Eqn (2). For e.g., if the i^{th} feature needs to have non-zero attribution, one may include the constraint: $f_{\mathbf{w}}(\mathbf{1}) - f_{\mathbf{w}}(\mathbf{i}) \geq \epsilon$, where $\mathbf{1}$ is vector of all ones, and \mathbf{i} is vector of all ones except that the i^{th} entry is zero, and ϵ is a small tolerance. If this constraint is satisfied, the marginal gain of this feature w.r.t. any subset of features and hence the attribution value will be greater than or equal to ϵ , again because of the diminishing returns property. Analogously, if a feature, say v , must have zero attribution, then adding the constraint $f_{\mathbf{w}}(\{v\}) = f_{\mathbf{w}}(\mathbf{0})$ should suffice, where $\mathbf{0}$ is the vector of all zeros. Lastly, the proposed methodology offers some unique features like: (a) one can attribute a set of images (say, all examples of a particular class) by simply pooling all the attributions of all the images and learning the submodular score. This will be interesting because now features important at set (class) level will be revealed; (b) If one needs to find p features that are most influential, then instead of picking the top- p attributed features, we can now solve for $\arg \max_{A \subseteq \{1, \dots, n\}, |A| \leq p} f_{\mathbf{w}}(A)$. This would give the top p discriminatory, yet non-redundant set of features.

Computational Effort. The main computational overhead for the projected subgradient algorithm is the computation of $\bar{A}_i^j \forall i \in \{1, \dots, m\} \forall j \in \{1, \dots, k\}$. Interestingly, the use of the greedy algorithm ensures that computing \bar{A} for larger thresholds gives those for the smaller ones too. Hence at every iteration of the descent, only one call to the greedy algorithm is required with cardinality bound as $B \equiv \max_{i \in \{1, \dots, m\}, j \in \{1, \dots, k\}} B_i^j$, leading to $O(Bn)$ evaluations of DSF output. Typically, $B \ll n$, thus leading to a marginal overhead of $O(n)$ DSF evaluations. In implementation, our module for cardinality-constrained submodular maximization performs $O(B)$

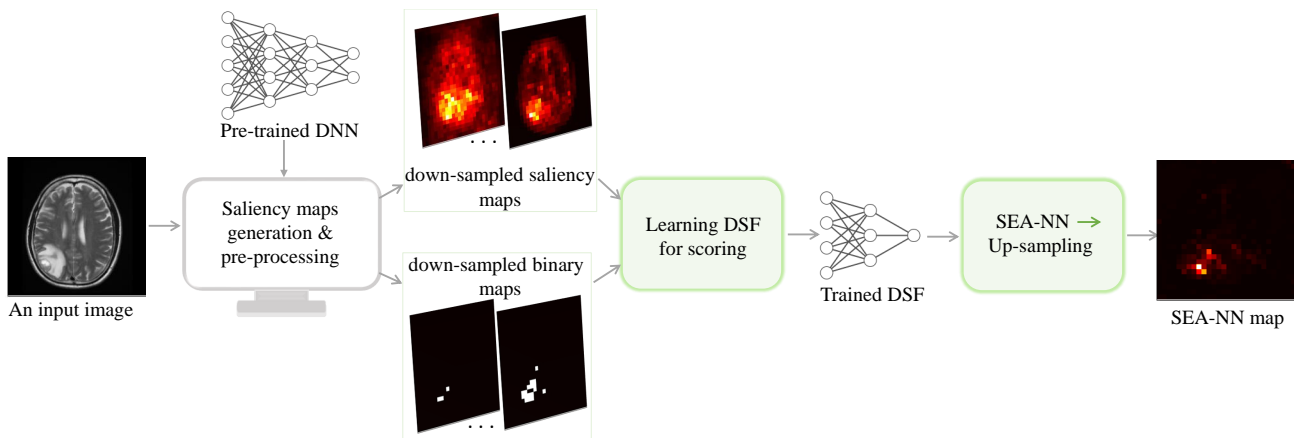


Figure 2: Overview of SEA-NN for the more general high-dimensional setting

calls to the DSF where in each call, we pass n inputs as a single batch. (We report quantitative measures of time complexity in the supplementary section.)

Scaling to High-Dimensional Images. For high-dimensional images, we learn a DSF using down-sampled versions of heatmaps and their binary (thresholded) counterparts. The proposed SEA-NN algorithm then computes attribution scores using the learned DSF. The attribution map thus obtained is up-sampled by interpolation (we use simple nearest-neighbor interpolation in this work) to match the original resolution of the image. Similar kind of upsampling techniques are also part of widely used attribution algorithms like CAM [Zhou et al., 2016] and Grad-CAM [Selvaraju et al., 2017]. For the downsampling itself, we divide a heatmap into grids and represent each grid by a single pixel in the downsampled heatmap where we assign average of attribution scores in the grid as the corresponding pixel’s value in the downsampled version. For downsampling the binary maps, we divide the binary map into grids, and each grid cell’s value is set to one if the sum of values inside it is greater than the average of the sums of values across all grid cells in the images. This provides a sense of relative importance for each grid location. An illustration is provided in the supplementary section for clarity. Empirically, however, we found that our method was not sensitive to the specifics of upsampling and downsampling.

An overall workflow of our approach for the more general high-dimensional setting is shown in Fig 2.

4 EXPERIMENTS AND RESULTS

In order to validate the proposed method, we performed a comprehensive suite of experiments on the Fashion-MNIST (or FMNIST) [Xiao et al., 2017], CUB [Wah et al., 2011], Tiny Imagenet [Le and Yang, 2015] and Brain Tumor

Detection ¹ datasets. Following [Lu et al., 2021], we used Vanilla Gradient (VG), Integrated Gradients (IG) and Smooth Integrated Gradient (SG) as the baseline attribution maps for our experiments in this section. These methods are widely used across DNN architectures, and were also found to be sufficient to learn a submodular scoring function that gave good results. We show more results with other methods in the supplementary section.

Baselines and Evaluation Metrics: We show qualitative and quantitative comparison against VG, IG, SG, as well as Agg-Mean [Rieger and Hansen, 2019] that does linear averaging of attribution maps. (Due to unavailability of source code, we are unable to compare against [Wagner et al., 2019], which however did not have the objective of ensembling attribution maps, but provided fine-grained explanations.) In addition to showing the qualitative results of the heatmaps, we also studied the methods quantitatively using the Area Under Perturbation Curve (AUPC) metric [Samek et al., 2017]. We perturb regions in the input image cumulatively, in the decreasing order of their relevance where relevance of a region depends on the attribution score assigned to it by an attribution method. The classifier’s score corresponding to the predicted label is recorded after each perturbation forming a perturbation curve. If an attribution method correctly highlights the relevant regions in an image, there will be a sharp decrease in the classifier’s score, making the area under the perturbation curve low. We report the AUPC metric averaged across all inputs. We follow the implementation of AUPC as suggested in [Goh et al., 2021] and [Petsiuk et al., 2018].

Implementation Details: We used PyTorch [Paszke et al., 2019] for all our implementations, and

¹<https://www.kaggle.com/navoneel/brain-mri-images-for-brain-tumor-detection>

Table 1: Area Under Perturbation Curve averaged over images of the datasets (lower is better)

Method	FMNIST	BTD	CUB	T.Img
SEA-NN	3.46	11.11	4.82	5.42
IG	3.62	16.78	4.9	5.49
VG	3.68	17.31	5.2	5.51
SG	3.57	17.51	5.21	5.79
Agg-Mean	3.51	17.24	4.84	5.47

used the Captum library [Kokhlikyan et al., 2020] to obtain the heatmaps for learning our DSF. In order to obtain an attribution map for a specific input (local explanation), we learn a DSF per input. The DSF architecture was chosen as Linear(784, 512)→Linear(512, 256)→Linear(256, 32)→Linear(32, 1) with square-root as the activation function after all but the last layer. The dataset for training a DSF comprised of heatmaps from VG, IG, SG along with their binary counterparts obtained using 10 threshold values equally spaced between 5 and 50. As our DSF is a small neural network with very few parameters, this size of training data was empirically found sufficient. Weights of the DSF NN were initialized to 1 and training was done for 50 epochs. For training, we used Adagrad optimizer with learning rate decay coefficient as 0.1. The hyperparameter δ was set to 10^{-5} . For the high-dimensional setting (CUB, TinyImagenet and BTD datasets), the downsampled version of SEA-NN map was chosen to be of resolution 28×28 . Other hyperparameters specific to a dataset are mentioned while describing the corresponding experiments.

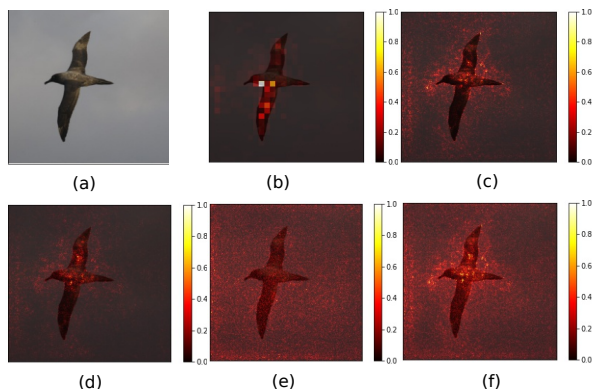


Figure 3: Example image and attribution maps on CUB dataset image: (a) Image of Sooty Albatross (b) SEA-NN (c) Integrated Gradient (d) Vanilla Gradient (e) Smooth Integrated Gradient (f) Agg-Mean

Results on CUB. We used a fine-tuned Resnet-18 classifier and experiment on the train split of CUB. The hyperparameters were set to: λ_1 as 0.1, λ_2 as 10 and Adagrad’s weight decay coefficient as 10^{-6} .

From Table 1, we observe that SEA-NN outperforms baseline methods on the average AUPC metric. Fig 3 shows qualitative results, demonstrating the higher specificity of SEA-NN over baseline methods. More qualitative results are provided in the supplementary section.

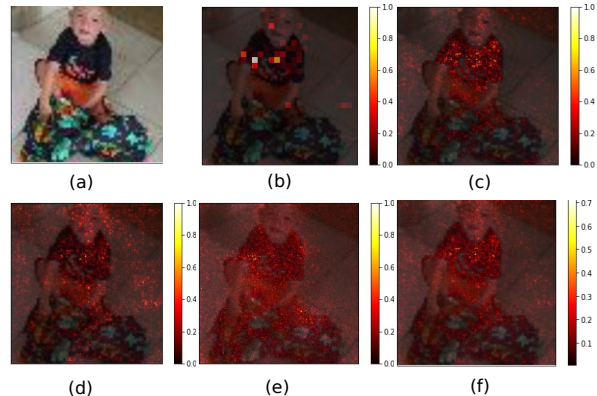


Figure 4: Example image and attribution maps on TinyImagenet dataset image: (a) Image of boy (b) SEA-NN (c) Integrated Gradient (d) Vanilla Gradient (e) Smooth Integrated Gradient (f) Agg-Mean

Results on TinyImagenet (T.Img). A fine-tuned Resnet-18 with top-1 accuracy of 57% was used as the classifier. The hyperparameters were set to: λ_1 as 0.1, λ_2 as 10 and Adagrad’s weight decay coefficient as 10^{-6} . Fig 4 shows an example where SEA-NN outperforms baseline methods on specificity of the heatmap. As shown in Table 1, the AUPC score averaged across the validation split of Tiny Imagenet for SEA-NN was the best among the considered methods. More qualitative results are in the supplementary section.

Correlations with human annotations. We conducted an experiment to see how well SEA-NN maps align with human annotated heatmaps. For this, we use human annotations provided by [Mohseni et al., 2020] for 98 images of the Imagenet dataset. We compute the Jaccard score between an attribution map and the annotated heatmap after hard-thresholding them to keep only the top- k pixels. We compute this for multiple thresholds (k values) and show the results in Fig 6. The figure demonstrates that the pixels picked by SEA are the most human-interpretable among all baselines method. We also show this through a qualitative example in Fig 5. Hyperparameters were: λ ’s as 0.1 and weight decay coefficient as 10^{-3} .

Results on Brain Tumor Detection (BTD). We used a fine-tuned VGG-11 classifier with top-1 accuracy of 85%. In the BTD dataset, white mass inside the skull is known to be indicative of brain tumor. The hyperparameters were set to: λ_1 as 1, λ_2 as 100 and weight decay coefficient for Adagrad as 10^{-3} . Fig

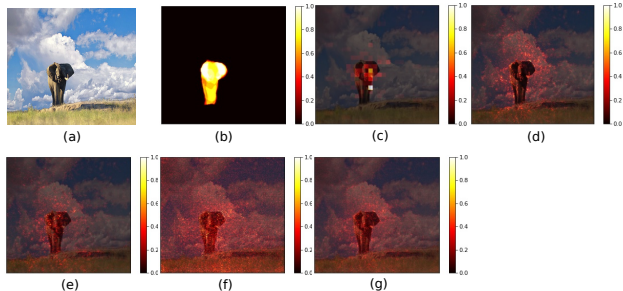


Figure 5: Correlations with human annotated heatmaps experiment (Imagenet image and attribution maps): (a) Image of an elephant (b) Human annotation (c) SEA-NN (d) Integrated Gradient (e) Vanilla Gradient (f) Smooth Integrated Gradient (g) Agg-Mean

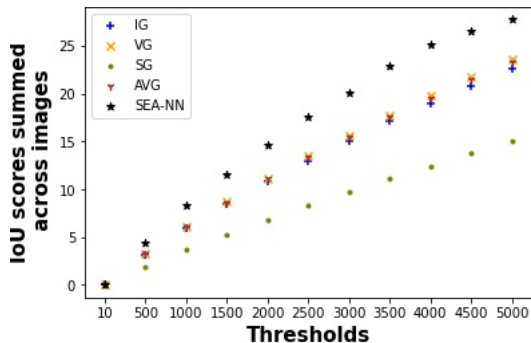


Figure 6: Results on Correlations with human annotations (higher is better)

1 shows that SEA-NN correctly identifies the affected regions with minimum noise compared to all baseline methods. Sparsity is highly beneficial in medical diagnosis where false positives could be fatal. Table 1 shows that SEA-NN obtains the best average AUPC score.

Experiment on Fashion-MNIST (FMNIST). We used a simple CNN classifier with accuracy of 90.5%. The hyperparameters were set to: λ_1 as 10, λ_2 as 10 and Adagrad’s weight decay coefficient as 10^{-6} . SEA-NN identified discriminatory pixels to achieve the best AUPC averaged across the test split of FMNIST (Table 1). SEA-NN map shows high specificity (Fig 4).

Visualizing Top-K Pixels. In order to show the effect of feature-interaction in attribution, we present binary heatmaps showing the top- k pixels identified by various attribution methods on the FMNIST dataset. As shown in Fig 4, the top- k pixels identified by SEA-NN are less redundant. We note that for the same budget, SEA-NN selects a more diverse set of pixels which is especially visible in the rightmost column when k is chosen as 10.

More results – qualitative and quantitative – are in-

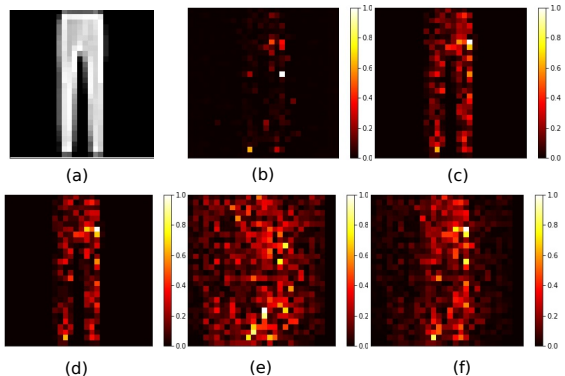


Figure 7: Example image and attribution maps on FMNIST dataset: (a) Image of Trouser (b) SEA-NN (c) Integrated Gradient (d) Vanilla Gradient (e) Smooth Integrated Gradient (f) Agg-Mean

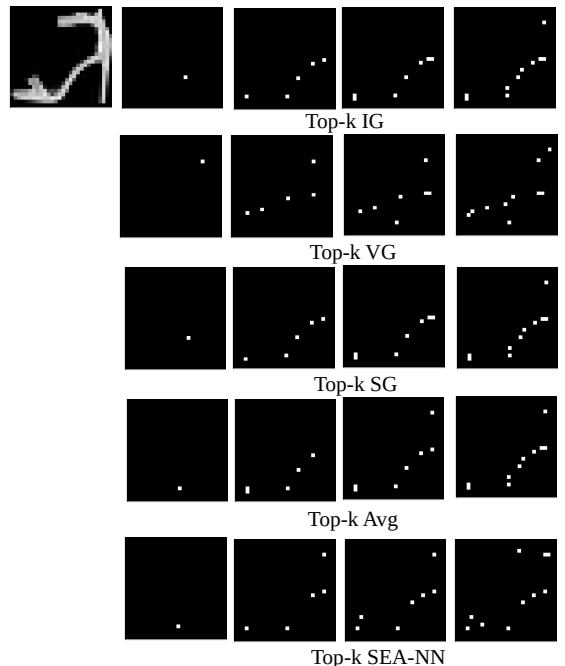


Figure 8: Top- k pixels selected by attribution methods on image from Sandal class in FMNIST. Columns from left to right represent k as 1, 5, 7, 10 respectively

cluded in the supplementary section owing to space constraints.

5 CONCLUSIONS AND FUTURE WORK

In this work, we propose a framework, SEA-NN, to improve existing attribution methods through submodular ensembling that provides more specificity to the attribution maps, and also context-dependent attribution scores. Ensembling attribution maps of different methods is expected to reduce epistemic uncertainty and biases associated with each attribution method. The special properties satisfied by our learned submodular scoring function and the ease of adding ad-

ditional constraints during training help making our solution suitable for needs of different users. We validated SEA-NN against well-known gradient-based attribution methods as well as existing ensembling methods and showed that our attribution maps are superior quantitatively and qualitatively. We also showed that SEA-NN aligns well with human annotated heatmaps available for a subset of Imagenet images.

Future Work and Social Impact. Attribution methods help enhancing the interpretability of DNNs which will encourage their usage in critical domains such as healthcare and criminal justice. The proposed method can be especially beneficial for domains like medical diagnosis and astronomy where specificity in attribution is highly desired. As a future work, we would like to explore a transfer learning approach such that DSF’s trained on just a subset of images, could work well on the entire dataset. This will also help further mitigate the computational effort incurred by our method.

Acknowledgements

We are grateful to the Ministry of Human Resource Development, India; Department of Science and Technology, India; Honeywell, India for the financial support of this project. We also thank the Japan International Cooperation Agency and IIT-Hyderabad for the provision of GPU servers used for this work. We thank the anonymous reviewers for their valuable feedback. We would also like to thank the AISTATS Mentorship program and our submission mentor, Dr Bilal Alsallakh. PM thanks Google for the PhD Fellowship.

References

- [Ancona et al., 2018] Ancona, M., Ceolini, E., Öztireli, C., and Gross, M. (2018). Towards better understanding of gradient-based attribution methods for deep neural networks. In *Proceedings of the 6th International Conference on Learning Representations*.
- [Bach et al., 2015] Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., and Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. volume 10. Public Library of Science.
- [Ben-Tal and Nemirovski, 2021] Ben-Tal, A. and Nemirovski, A. (2021). Lectures On Modern Convex Optimization. Available freely at <https://www2.isye.gatech.edu/~nemirovs/LMCOLN2021WithSol.pdf>.
- [Bhatt et al., 2020] Bhatt, U., Weller, A., and Moura, J. M. F. (2020). Evaluating and aggregating feature-based model explanations. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence, IJCAI-20*, volume abs/2005.00631.
- [Chattopadhyay et al., 2018] Chattopadhyay, A., Sarkar, A., Howlader, P., and Balasubramanian, V. N. (2018). Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *Proceedings of the 2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*.
- [Chattopadhyay et al., 2019] Chattopadhyay, A., Manupriya, P., Sarkar, A., and Balasubramanian, V. N. (2019). Neural network attributions: A causal perspective. In *Proceedings of the 36th International Conference on Machine Learning*.
- [Cheraghchi et al., 2011] Cheraghchi, M., Klivans, A., Kothari, P., and Lee, H. (2011). Submodular functions are noise stable. volume 18.
- [Dolhansky and Bilmes, 2016] Dolhansky, B. W. and Bilmes, J. A. (2016). Deep submodular functions: Definitions and learning. In *Proceedings of the 30th Neural Information Processing Systems*.
- [Elenberg et al., 2017] Elenberg, E., Dimakis, A. G., Feldman, M., and Karbasi, A. (2017). Streaming weak submodularity: Interpreting neural networks on the fly. In *Proceedings of the 31st Conference on Neural Information Processing Systems*.
- [Goh et al., 2021] Goh, G. S. W., Lapuschkin, S., Weber, L., Samek, W., and Binder, A. (2021). Understanding integrated gradients with smooth-taylor for deep neural network attribution. In *2020 25th International Conference on Pattern Recognition, (ICPR)*, pages 4949–4956, Virtual Event / Milan, Italy. IEEE.
- [Hartley et al., 2021] Hartley, T., Sidorov, K., Willis, C., and Marshall, D. (2021). Swag: Superpixels weighted by average gradients for explanations of cnns. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- [Iyer and Bilmes, 2020] Iyer, R. K. and Bilmes, J. A. (2020). Concave aspects of submodular functions. In *Proceedings of the IEEE International Symposium on Information Theory ISIT*.
- [Kapishnikov et al., 2019] Kapishnikov, A., Bolukbasi, T., Viegas, F., and Terry, M. (2019). Xrai: Better attributions through regions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

- [Kokhlikyan et al., 2020] Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., Melnikov, A., Kliushkina, N., Araya, C., Yan, S., and Reblitz-Richardson, O. (2020). Captum: A unified and generic model interpretability library for pytorch.
- [Krause and Golovin, 2014] Krause, A. and Golovin, D. (2014). Submodular function maximization.
- [Le and Yang, 2015] Le, Y. and Yang, X. (2015). Tiny imagenet visual recognition challenge.
- [Lecue et al., 2021] Lecue, F., Guidotti, R., Minervini, P., and Giannotti, F. (2021). On Explainable AI: From Theory to Motivation, Industrial Applications and Coding Practices. <https://xaitutorial2021.github.io/>. [Online; accessed 11-October-2021].
- [Lu et al., 2021] Lu, Y. Y., Guo, W., Xing, X., and Noble, W. S. (2021). Dance: Enhancing saliency maps using decoys. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*. PMLR.
- [Lundberg and Lee, 2017] Lundberg, S. M. and Lee, S. (2017). A unified approach to interpreting model predictions. In *Proceedings of the Neural Information Processing Systems (NeurIPS)*.
- [Mohseni et al., 2020] Mohseni, S., Block, J. E., and Ragan, E. D. (2020). Quantitative evaluation of machine learning explanations: A human-grounded benchmark. *arXiv preprint arXiv:1801.05075*.
- [Molnar, 2019] Molnar, C. (2019). *Interpretable Machine Learning*.
- [Nemhauser et al., 1978] Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1).
- [Paszke et al., 2019] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Proceedings of the 33rd Advances in Neural Information Processing Systems*. Curran Associates, Inc.
- [Petsiuk et al., 2018] Petsiuk, V., Das, A., and Saenko, K. (2018). RISE: randomized input sampling for explanation of black-box models. In *Proceedings of the 29th British Machine Vision Conference 2018, BMVC 2018*.
- [Rieger and Hansen, 2019] Rieger, L. and Hansen, L. K. (2019). Aggregating explainability methods for neural networks stabilizes explanations. *CoRR*, abs/1903.00519.
- [Rieger and Hansen, 2020] Rieger, L. and Hansen, L. K. (2020). A simple defense against adversarial attacks on heatmap explanations. In *Proceedings of the 5th Annual Workshop on Human Interpretability at the International Conference on Machine Learning (ICML)*, volume abs/2007.06381.
- [Samek et al., 2017] Samek, W., Binder, A., Montavon, G., Lapuschkin, S., and Müller, K.-R. (2017). Evaluating the visualization of what a deep neural network has learned. volume 28.
- [Selvaraju et al., 2017] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [Shrikumar et al., 2017] Shrikumar, A., Greenside, P., and Kundaje, A. (2017). Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*. PMLR.
- [Simonyan et al., 2014] Simonyan, K., Vedaldi, A., and Zisserman, A. (2014). Deep inside convolutional networks: Visualising image classification models and saliency maps. In *Proceedings of the 2nd International Conference on Learning Representations, ICLR Workshop*.
- [Smilkov et al., 2017] Smilkov, D., Thorat, N., Kim, B., Viégas, F. B., and Wattenberg, M. (2017). Smoothgrad: removing noise by adding noise. In *Proceedings of the Visualization for Deep Learning Workshop at the International Conference on Machine Learning (ICML)*.
- [Springenberg et al., 2015] Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. A. (2015). Striving for simplicity: The all convolutional net. In *Proceedings of the 3rd International Conference on Learning Representations, ICLR Workshop*.
- [Sundararajan et al., 2017] Sundararajan, M., Taly, A., and Yan, Q. (2017). Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*. PMLR.

- [Tjoa and Guan, 2020] Tjoa, E. and Guan, C. (2020). A survey on explainable artificial intelligence (xai): Toward medical xai. *Proceedings of the IEEE Transactions on Neural Networks and Learning Systems*.
- [Tsang et al., 2020] Tsang, M., Rambhatla, S., and Liu, Y. (2020). How does this interaction affect me? interpretable attribution for feature interactions. In *Proceedings of the 34th Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc.
- [Wagner et al., 2019] Wagner, J., Kohler, J. M., Gindele, T., Hetzel, L., Wiedemer, J. T., and Behnke, S. (2019). Interpretable and fine-grained visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [Wah et al., 2011] Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. (2011). The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology.
- [Wang et al., 2020] Wang, Z., Wang, H., Ramkumar, S., Mardziel, P., Fredrikson, M., and Datta, A. (2020). Smoothed geometry for robust attribution. In *Advances in Neural Information Processing Systems*, volume 33.
- [Xiao et al., 2017] Xiao, H., Rasul, K., and Vollgraf, R. (2017). Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms.
- [Yang and Kim, 2019] Yang, M. and Kim, B. (2019). Benchmarking Attribution Methods with Relative Feature Importance. *CoRR*, abs/1907.09701.
- [Zeiler and Fergus, 2014] Zeiler, M. D. and Fergus, R. (2014). Visualizing and understanding convolutional networks. In *Proceedings of the 13th European Conference on Computer Vision (ECCV)*.
- [Zhou et al., 2016] Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., and Torralba, A. (2016). Learning deep features for discriminative localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Supplementary Material: Improving Attribution Methods by Learning Submodular Functions

In this supplementary material, we include additional information which could not be included in the main paper due to space constraints. In particular, we include the following:

- Description of baseline attribution maps that we use for learning our scoring function
- More details on Deep Submodular Functions (DSF)
- More details on Lemma 1
- Toy illustration for down-sampling binary heatmaps
- More results (in continuation to Section 4)
- Quantitative comparison of computation time.

A BASELINE ATTRIBUTION ALGORITHMS

In this section, we briefly describe the attribution methods using which we learn our submodular scoring function. We use x to denote the input image and F^c to denote neural network’s output corresponding to class c .

Vanilla Gradient (VG) VG assigns attribution score for the i^{th} input dimension as

$$\text{VG}_i(x; F^c) := \frac{\partial F^c}{\partial x_i}$$

Integrated Gradients (IG) IG tries to mitigate the saturating gradients problem in VG by aggregating the attribution maps corresponding to inputs lying on a path from a baseline x' to the original input x . IG assigns attribution score along i^{th} dimension for input x as

$$\text{IG}_i(x; F^c) := (x_i - x'_i) \int_{\alpha=0}^1 \frac{\partial F^c(x' + \alpha(x - x'))}{\partial x_i} \partial \alpha$$

Smooth Integrated Gradients (SG) Smooth Grad variants are designed to remove noise from attribution maps by adding noise to the input. Attribution score for Smooth Integrated Gradients along input dimension i is computed as

$$\text{SG}_i(x; F^c) := \frac{1}{N} \sum_{j=1}^N \text{IG}_i(x + n_j; F^c)$$

where $n_j \sim \mathcal{N}(0, \sigma^2)$ for $j = 1 \dots N$.

B MORE DETAILS ON DEEP SUBMODULAR FUNCTIONS

A Deep Neural Network (DNN) whose weights are restricted to be non-negative and the activation functions used are monotone non-decreasing concave for non-negative reals, constitutes a non-negative monotone non-decreasing submodular function when given Boolean input vectors. This is referred to as Deep Submodular Function (DSF)

that can be trained in the similar way as DNNs. The key result that forms the basis of DSF is that sums of concave functions, composed with modular functions (SCMMs), can be easily shown to be submodular. We urge the interested readers to refer [Dolhansky and Bilmes, 2016] for more details.

For choosing the NN architecture for DSF in our work, we started with a two-layer NN and kept on adding a layer until our objective over epoch curve started looking saturated at the end of 50 epochs. As described in section 4 of the main paper, the NN architecture chosen for DSF was $\text{Linear}(784, 512) \rightarrow \text{Linear}(512, 256) \rightarrow \text{Linear}(256, 32) \rightarrow \text{Linear}(32, 1)$. As the choice of concave activation function, we used $g(x) = \sqrt{x}$ after all but the last layer. We also tried $g(x) = \log(1 + x)$ as the activation function but we observed that using square root as the activation function resulted in smoother attribution maps.

C MORE DETAILS ON LEMMA 1

For the DSF architecture used in the paper, the Lipschitz continuity constant (L) for f_{ext} (concave extension of DSF) comes out to be $\frac{\prod_{i=1}^4 (\mathbf{w}_i^T \mathbf{1}) 2^{5^{i-1}}}{7}$ where \mathbf{w}_i denotes vectorized weights of layer i of the DSF. For deriving this L , we applied Theorem C.4.1 in [Ben-Tal and Nemirovski, 2021] and the property that f_{ext} is a monotonically increasing function defined over $[0, 1]^{784}$.

We also computed $\sup_{\mathbf{x}, \mathbf{y}} \frac{|f_{ext}(\mathbf{x}) - f_{ext}(\mathbf{y})|}{\|\mathbf{x} - \mathbf{y}\|}$ over 10^6 pairs of random inputs $\mathbf{x}, \mathbf{y} \sim [0, 1]^{784}$ for a DSF learnt on a random image from the Imagenet dataset and obtained L estimate as 6.54, which is not very high.

D DOWN-SAMPLING BINARY HEATMAPS

As detailed in section 3 of the main paper, for the case of high dimensional images, we learn our scoring function using down-sampled versions of baseline heatmaps and their binary (thresholded) counterparts in order to have better computational efficiency. In Fig 9, we present a toy illustration of the down-sampling procedure for binary heatmaps.

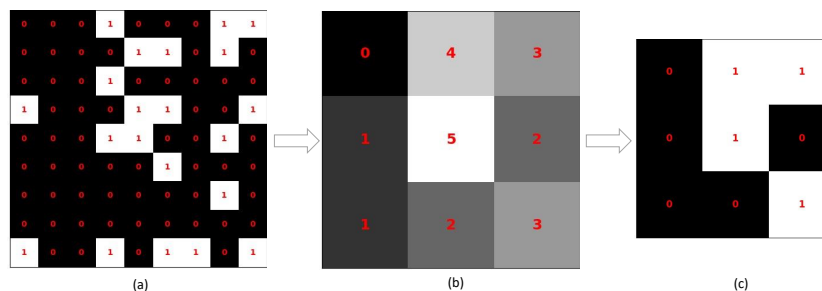


Figure 9: Toy illustration of down-sampling binary heatmaps (a) Binary attribution map (b) Intermediate representation considering 3×3 grids (c) Down-sampled binary map

E MORE RESULTS

SEA-NN is designed to enhance specificity in attribution. However, in the absence of standard metrics that simultaneously quantify sensitivity and specificity, we evaluate the discriminative power of our method using the AUPC metric and show qualitative results demonstrating good specificity.

We note that SEA-NN encompasses multiple benefits. It not only improves specificity in attribution but is a novel method to aggregate feature attribution and provides attribution scores that respect feature interactions. Moreover, we propose a new formulation for learning a submodular set function using real-valued inputs. To the best of our knowledge, there is no other method that provides all these benefits.

Following we demonstrate some more results showing the efficacy of the proposed approach.

E.1 More Qualitative Results

In continuation to section 4 of the main paper, we present more qualitative results in this section. Figures 10, 11, 12 and 13 show qualitative comparison of attribution maps for CUB, Brain Tumor Detection, Imagenet and Tiny Imagenet datasets respectively. Attribution maps of SEA-NN exhibit higher specificity than the attribution maps of other attribution methods.

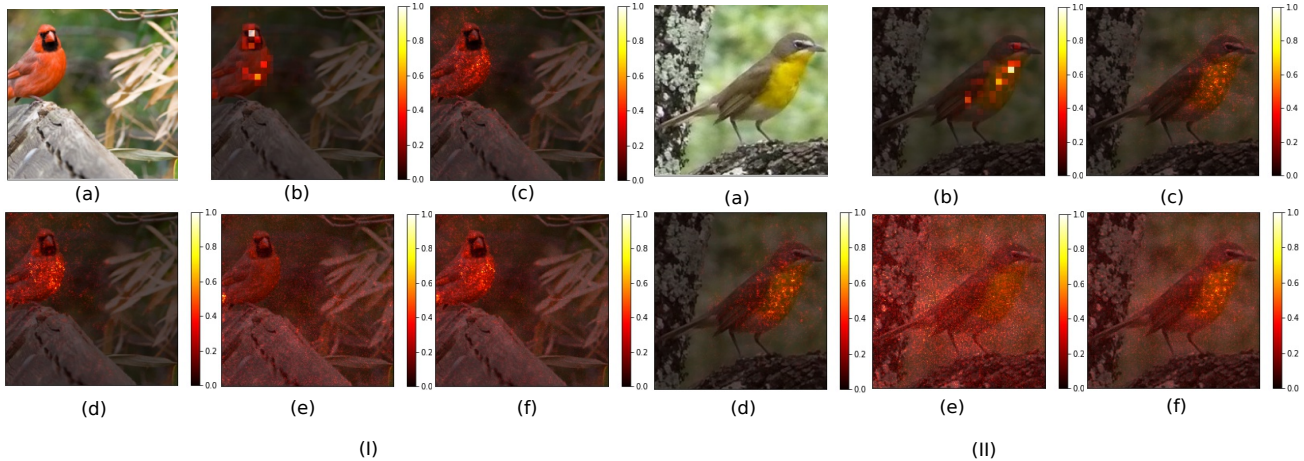


Figure 10: (I) and (II) are images of Golden Retriever birds shown with attribution maps overlaid on the respective images: (a) Image; (b) SEA-NN map; (c) Integrated Gradient map; (d) Vanilla Gradient map; (e) Smooth Integrated Gradient map; (f) Agg-Mean map of (c)-(e) maps.

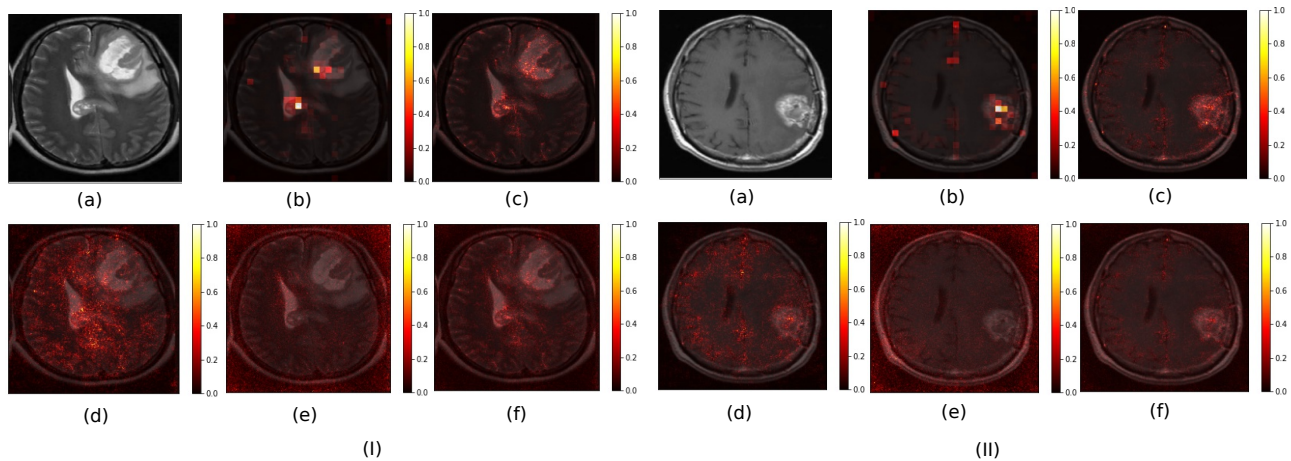


Figure 11: (I) and (II) are images from Brain Tumor dataset shown with attribution maps overlaid on the respective images: (a) Image with brain tumor; (b) SEA-NN map; (c) Integrated Gradient map; (d) Vanilla Gradient map; (e) Smooth Integrated Gradient map; (f) Agg-Mean map of (c)-(e) maps.

Global Attribution using SEA-NN In addition to providing local explanation for a specific input, our proposed attribution algorithm can also provide global explanation for an entire class in a dataset. Global explanations provide us with a more holistic understanding of the classifier NN. To extend SEA-NN as a global attribution algorithm, we only need to expand the training dataset for DSF to include heatmaps and their binary counterparts of all inputs belonging to a given class in the dataset. In Fig 14, we show global attribution maps of SEA-NN corresponding to three different FMNIST classes - Trouser, Coat and Sneaker; overlaid on randomly selected images belonging to these classes. The hyper-parameters for this experiment were the same as that of the FMNIST experiment included in section 4 of the main paper.

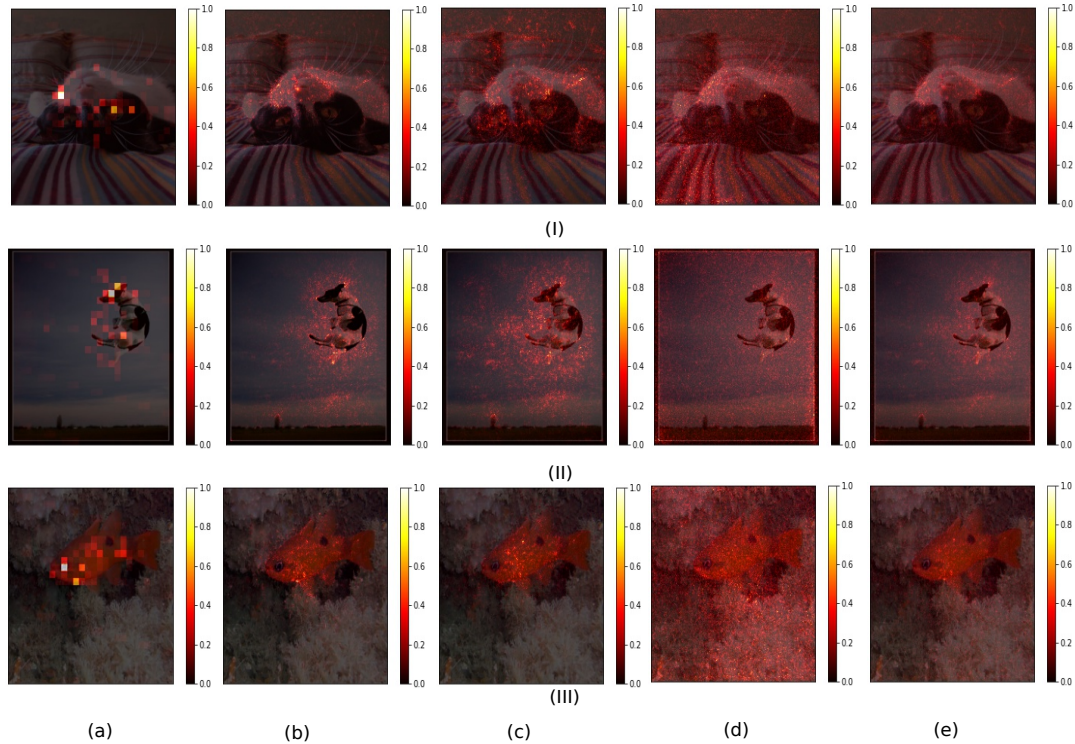


Figure 12: (I), (II), (III) are attribution maps overlaid on respective images of a Cat, a Dog and a Shark respectively: (a) SEA-NN map; (b) Integrated Gradient map; (c) Vanilla Gradient map; (d) Smooth Integrated Gradient map; (e) Agg-Mean map of (b)-(d) maps.

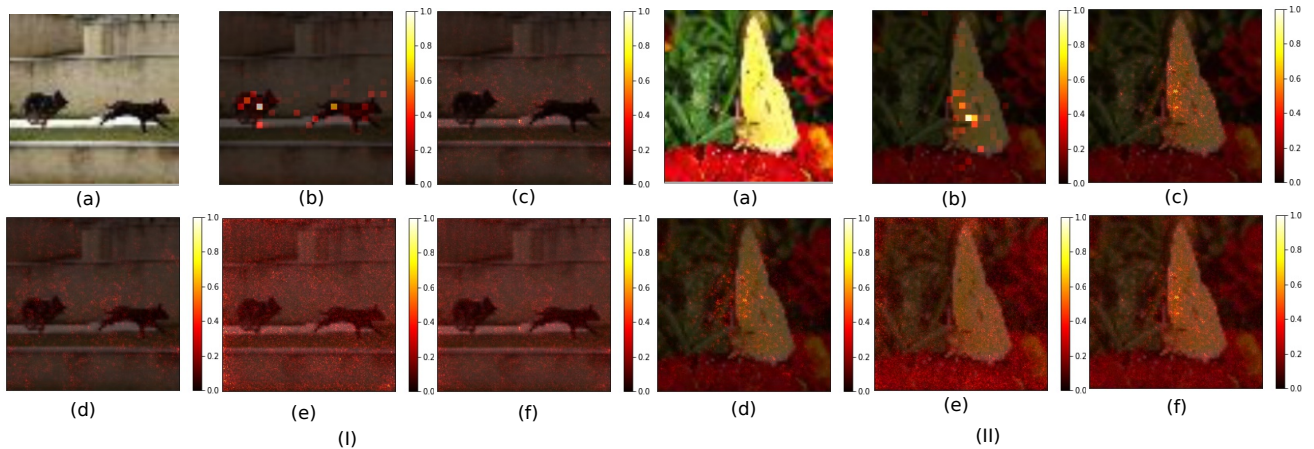


Figure 13: (I), (II) show images of German Shepherd dogs and a Sulphur Butterfly respectively with attribution maps overlaid on them: (a) Image; (b) SEA-NN map; (c) Integrated Gradient map; (d) Vanilla Gradient map; (e) Smooth Integrated Gradient map; (f) Agg-Mean map of (c)-(e) maps.

Visualizing Top- k pixels In continuation to section 4 in the main paper, we present another result visualizing the top- k pixels from FMNIST dataset. Fig 15 shows comparison of top- k pixels for image of a Bag, as identified by different attribution maps. While most of the attribution maps only identified the handle of the bag, SEA-NN identifies a more diverse set of discriminatory pixels especially when k is 10. The effect of feature-interaction and the benefits of submodularity help SEA-NN identify a more diverse set.

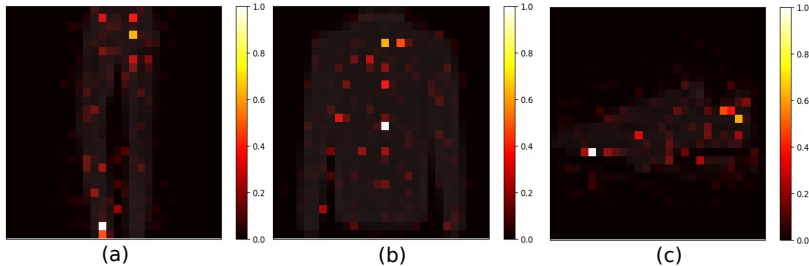


Figure 14: Global attribution maps of SEA-NN overlaid on images of corresponding classes (a) Trouser class (b) Coat class (c) Sneaker class

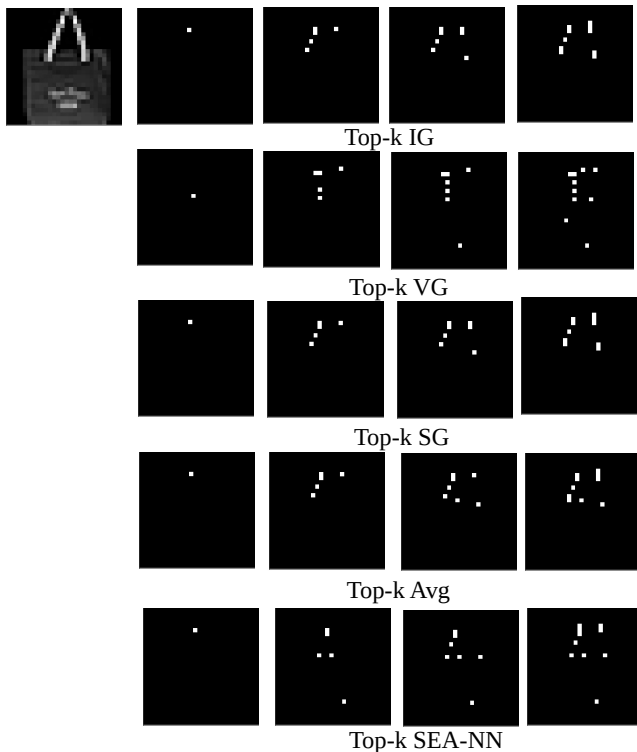


Figure 15: Top- k pixels selected by attribution methods on image from Bag class in FMNIST. Columns from left to right represent k as 1, 5, 7, 10 respectively

E.2 More Quantitative Results

We present some more details on the quantitative evaluation, some ablation studies and additional quantitative results in this section.

We tuned hyper-parameters of our method for a given dataset using a random subset of the same dataset. For computing average AUPC, we only considered the correctly classified images in case of FMNIST, CUB and Tiny Imagenet datasets. We considered all images for computing average AUPC on BTM as it was a relatively small dataset. In order to perturb pixels for computing AUPC, a popular approach is to choose top- k pixels with the choice of k as $\frac{1}{28}$ th of the total number of pixels [Petsiuk et al., 2018]. We followed this while computing AUPC for BTM. For datasets FMNIST, CUB, Tiny Imagenet, we perturb input regions of size $s \times s$ following [Goh et al., 2021] and consider total 8 such perturbations. The hyper-parameter s was chosen as 16 for CUB and Tiny Imagenet. For FMNIST, s was chosen as 8.

Effect of Regularization Hyper-parameters We perform ablation studies on BTM to gauge the effect of regularization hyper-parameters λ_1 and λ_2 used for learning the scoring function (Eq. 2 in the main paper) and show results in Table 3.

Comparing against more methods In Table 2, we show comparisons against more attribution methods in addition to the ones we compared against in Table 1.

Table 2: Area Under Perturbation Curve averaged over images of the datasets (lower is better)

Method	FMNIST	BTD	CUB	T.Img
SEA-NN	3.46	11.11	4.82	5.42
<i>input \circ gradient</i>	3.69	16.67	6.21	5.79
Deep-Lift	4.39	16.8	4.85	5.57
LIME	4.48	18.24	5.17	5.73

Robustness We conducted an experiment to test the robustness of attribution methods to noise in the input. We perturbed each input image of CUB by adding noise sampled from the uniform distribution, $\mathcal{U}(-0.02, 0.02)$ and re-computed the attribution maps on these noisy inputs. For an attribution method that shows robustness to noise, we would expect these re-computed attribution maps to exhibit high similarity to the original attribution maps. To quantify this similarity, we computed Intersection over Union (IoU) scores between the thresholded versions of the re-computed attribution maps and the original attribution maps. Top 5000 pixels identified by the respective attribution maps were retained after thresholding these maps. For this experiment, we compare against Agg-Mean and Smooth Integrated Gradients (SG) as these methods have been observed to be relatively robust to noise [Rieger and Hansen, 2020, Wang et al., 2020]. In Table 4, we report the IoU scores averaged over the entire dataset. SEA-NN achieves the best average IoU score.

(λ_1, λ_2)	AUPC
(1, 1)	11.91
(1, 10)	10.75
(1, 100)	11.47
(10, 1)	11.55
(10, 10)	11.43
(10, 100)	11.23
(100, 1)	11.46
(100, 10)	11.39
(100, 100)	11.15

Table 3: Ablation experiment on BTD

Table 4: Robustness comparison on CUB (higher average IoU is better)

Method	Average IoU
SEA-NN	2.94×10^{-1}
Agg-Mean	2.69×10^{-1}
SG	1.79×10^{-1}

Additional Experiment We also got promising performance on evaluating our method on randomly chosen 100 images of the Benchmarking Attribution Methods (BAM) dataset [Yang and Kim, 2019] using the Input Dependence Rate (IDR) metric proposed by them.

Metric	SEA-NN	Agg-Mean	IG	SG
IDR \uparrow	0.87	0.84	0.84	0.79

F DETAILS OF COMPUTATION TIME

In Table 5, we present a comparison of average computation time taken by different attribution methods on twenty randomly selected inputs from the datasets. These computation times were recorded after running experiments on an NVIDIA GTX 1080 Ti. As shown in Table 5, the computational time for SEA-NN becomes comparable to a popular attribution algorithm LIME for deeper classifiers. Furthermore, if we already have the down-sampled

Table 5: Computation time (in sec) averaged across twenty randomly selected inputs.

Method	FMNIST (5-layer CNN)	BTD (VGG-11)	CUB (ResNet-18)	T.Img (ResNet-18)
Proposed	$1.45 \times 10^{+1}$	$1.48 \times 10^{+1}$	$1.55 \times 10^{+1}$	$1.56 \times 10^{+1}$
LIME	3.89×10^{-1}	$1.27 \times 10^{+1}$	$1.26 \times 10^{+1}$	$1.44 \times 10^{+1}$
SG	2.49×10^{-2}	6.50×10^{-1}	7.20×10^{-1}	7.28×10^{-1}
IG	1.11×10^{-2}	2.12×10^{-1}	9.64×10^{-2}	9.43×10^{-2}
VG	2.08×10^{-3}	4.03×10^{-3}	8.07×10^{-3}	8.35×10^{-3}
<i>inp</i> \circ <i>grad</i>	2.19×10^{-3}	3.99×10^{-3}	9.84×10^{-3}	9.61×10^{-3}
GBP	2.17×10^{-3}	4.78×10^{-3}	9.39×10^{-3}	9.79×10^{-3}
Deconv	2.44×10^{-3}	5.05×10^{-3}	1.15×10^{-2}	1.17×10^{-2}
DL	7.11×10^{-3}	1.84×10^{-2}	2.03×10^{-2}	2.43×10^{-2}

baseline attribution maps and their binary counterparts then the computation time for the proposed algorithm does not scale-up as the depth of the classifier NN increases. However, the computation time for other attribution methods increases as the classifier NN becomes deeper.

We note that attribution methods like XRAI [Kapishnikov et al., 2019], RISE [Petsiuk et al., 2018] and SWAG [Hartley et al., 2021] also require similar computational effort as the proposed algorithm. Our future work will include exploring a transfer learning approach so that a DSF trained on a subset of inputs can be used for the entire dataset. This will further reduce the computation time for us.