
A unified view of gradient-based attribution methods for Deep Neural Networks

Marco Ancona

Department of Computer Science
ETH Zurich, Switzerland
anconam@inf.ethz.ch

Enea Ceolini

Institute of Neuroinformatics
University Zürich and ETH Zürich
enea.ceolini@ini.uzh.ch

Cengiz Öztireli*

ETH Zurich and
Disney Research Zurich
cengizo@inf.ethz.ch

Markus Gross

Department of Computer Science
ETH Zurich, Switzerland
grossm@inf.ethz.ch

Abstract

Understanding the flow of information in Deep Neural Networks (DNNs) is a challenging problem that has gain increasing attention over the last few years. While several methods have been proposed to explain network predictions, only a few attempts to analyze them from a theoretical perspective have been made in the past. In this work, we analyze various state-of-the-art attribution methods and prove unexplored connections between them. We also show how some methods can be reformulated and more conveniently implemented. Finally, we perform an empirical evaluation with six attribution methods on a variety of tasks and architectures and discuss their strengths and limitations.

1 Introduction and Motivation

While DNNs have had a large impact on several different tasks [8, 7, 11, 19, 25], explaining their predictions is still a challenge and a limitation for those scenarios where predictability is crucial. In this work, we study the problem of assigning an *attribution* value, sometimes also called "relevance" or "contribution", to each input feature of a network. More formally, consider a DNN that takes an input $x = [x_1, \dots, x_N] \in \mathbb{R}^N$ and produces an output $S(x) = [S_1(x), \dots, S_C(x)]$, where C is the total number of output neurons. Given a specific target neuron c , the goal of an attribution method is to determine the attribution $R^c = [R_1^c, \dots, R_N^c] \in \mathbb{R}^N$ of each input feature x_i to the output S_c . For a classification task, the target neuron of interest is usually the output neuron associated with the correct class for a given sample. Attributions are often visualized as heatmaps, called *attribution maps*, where red and blue colors indicate respectively features that contribute positively to the activation of the target output and features having a suppressing effect on it (Figures 1-2).

The problem of attribution for deep networks has been tackled in several previous works [20, 27, 22, 3, 17, 23, 12, 28]. Unfortunately, due to slightly different problem formulations, lack of compatibility with all different existing DNN architectures and no common benchmark, a comprehensive comparison is not available. In this work we prove that ϵ -LRP [3] and DeepLIFT (Rescale) [17] can be reformulated as computing backpropagation that contribute positively to the activation of the target output and features having a suppressing effect on it (Figures 1-2). This allows the construction of a unified framework that comprises several gradient-based attribution methods and facilitates their implementation. In the last section, we empirically compare several attribution methods and discuss practical directions for their application.

*The author contributed to this work while employed at ETH Zurich

2 Overview over existing attribution methods

2.1 Perturbation-based methods

Perturbation-based methods directly compute the attribution of an input feature (or set of features) by removing, masking or altering them, and running a forward pass on the new input, measuring the difference with the original output. This technique has been applied to Convolutional Neural Networks (CNNs) in the domain of image classification, visualizing the probability of the correct class as a function of the position of a grey patch occluding part of the image [27]. While perturbation-based methods allow a direct estimation of the marginal effect of a feature, they tend to be very slow as the number of features to test grows (ie. up to hours for a single image [28]). What is more, given the nonlinear nature of DNNs, the result is strongly influenced by the number of features that are removed altogether at each iteration (Figure 1).

In the remainder of the paper, we will consider the occluding method by [27] as a comparison benchmark for perturbation-based methods. We will use this method, referred to as *Occlusion-1*, replacing one feature x_i at the time with a zero baseline and measuring the effect of this perturbation on the target output, ie. $S_c(x) - S_c(x_{[x_i=0]})$ where we use $x_{[x_i=v]}$ to indicate a sample $x \in R^N$ whose i -th component has been replaced with v .




Figure 1: Attributions generated by occluding portions of the input image with squared grey patches of different sizes. Notice how the size of the patches influence the result, with focus on the main subject only when using bigger patches. Best seen in electronic form.

2.2 Backpropagation-based methods

Backpropagation-based methods compute the attributions for all input features in a single forward and backward pass through the network². While these methods are generally faster than perturbation-based methods, their outcome can hardly be related directly to a variation of the output. Except where explicitly reported, the mathematical formulation of the method below is reported in Table 1.

Saliency Maps [20] constructs attributions by taking the absolute value of the partial derivative of the target output S_c with respect to the input features x_i . Intuitively, the absolute value of the gradient indicates those input features (pixels, for image classification) that can be perturbed the least in order for the target output to change the most, with no regards for the direction of this change. Nevertheless, Saliency Maps are usually rather noisy [15, 12, 21] and taking the absolute value prevents the detection of positive and negative evidence that might be present in the input.

Gradient * Input [18] was at first proposed as a technique to improve the sharpness of attribution maps. The attribution is computed taking the (signed) partial derivatives of the output with respect to the input and multiplying them feature-wise with the input itself.

Integrated Gradients [23], similarly to Gradient * Input, computes the partial derivatives of the output with respect to each input feature. However, instead of evaluating the partial derivative at the provided input x only, it computes its *average* value while the input varies along a linear path from a baseline \bar{x} to x . The baseline is defined by the user and often chosen to be zero.

Layer-wise Relevance Propagation (LRP) [3] is computed with a backward pass on the network. Let us consider a quantity $r_i^{(l)}$, called "relevance" of unit i of layer l . The algorithm starts at the output layer L , assigning the relevance of the target neuron c equal to the activation of the neuron itself, and

²Sometimes several of these steps are necessary, but the number does not depend on the number of input feature and generally much smaller than for perturbation-based methods

the relevance of all other neurons to zero (Equation 1). Then it proceeds layer by layer, redistributing the prediction score S_i until the input layer is reached. One recursive rule for the redistribution of a layer's relevance to the following layer is the ϵ -rule described in Equation 2, where we defined $z_{ij} = x_i^{(l)} w_{ij}^{(l,l+1)}$ to be the weighted activation of a neuron i onto neuron j in the next layer and b_j the additive bias of unit j . A small quantity ϵ is added to the denominator to avoid numerical instabilities. Once reached the input layer, the final attributions are defined as $R_i^c(x) = r_i^{(1)}$.

$$r_i^{(L)} = \begin{cases} S_i(x) & \text{if unit } i \text{ is the target unit of interest} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

$$r_i^{(l)} = \sum_j \frac{z_{ij}}{\sum_{i'} (z_{i'j} + b_j) + \epsilon \cdot \text{sign}(\sum_{i'} (z_{i'j} + b_j))} r_j^{(l+1)} \quad (2)$$

LRP together with the propagation rule described in Equation 2 is called ϵ -LRP, analyzed in the remainder of this paper. Alternative stabilizing methods exists [3, 12] but are not considered here. Additionally, we assume a small and fixed ϵ , with the only purpose of avoiding divisions by zero.

DeepLIFT [17] proceeds in a backward fashion, similarly to LRP. Each unit i is assigned an attribution that represents the relative effect of the unit activated at the original network input x compared to the activation at some reference input \bar{x} (Equation 3). Reference values \bar{z}_{ij} for all hidden units are determined running a forward pass through the network, using the baseline \bar{x} as input, and recording the activation of each unit. The baseline is a user-defined parameter often chosen to be zero. Equation 4 describes the relevance propagation rule.

$$r_i^{(L)} = \begin{cases} S_i(x) - S_i(\bar{x}) & \text{if unit } i \text{ is the target unit of interest} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

$$r_i^{(l)} = \sum_j \frac{z_{ij} - \bar{z}_{ij}}{\sum_{i'} z_{i'j} - \sum_{i'} \bar{z}_{i'j}} r_j^{(l+1)} \quad (4)$$

In Equation 4, $\bar{z}_{ij} = \bar{x}_i^{(l)} w_{ij}^{(l,l+1)}$ is weighted activation of a neuron i onto neuron j when the baseline \bar{x} is fed into the network. The attributions at the input layer are defined as $R_i^c(x) = r_i^{(1)}$. The rule here described ("Rescale rule") is used in the original formulation of the method and it is the one we will analyze in the remainder of the paper. The "Reveal-Cancel" rule [17] is not considered here.




Figure 2: Backpropagation-based attribution methods applied to an Inception V3 network for natural image classification [24]. All signed methods produce attributions affected by higher local variance compared to perturbation-based methods (Figure 1). Best seen in electronic form.

Other methods that are designed only for specific architectures (ie. **Grad-CAM** [16] for CNNs) or activation functions (ie. **Deconvolutional Network** [27] and **Guided Backpropagation** [22] for Rectified Linear Units (ReLUs)) are out of the scope of this analysis.

3 A gradient-centric formulation

Saliency Maps, Gradient * Input and Integrated Gradients are, by definition, computed as a function of the partial derivatives of the target output with respect to all input features. In this section, we will show that ϵ -LRP and DeepLIFT can also be computed by applying the chain rule for gradients, if the instant gradient at each nonlinearity is replaced with a function that depends on the method.

In a DNN where each layer performs a linear transformation of the input $z_j = \sum_i x_i w_{ij} + b_j$ followed by a nonlinear mapping $x_j = f(z_j)$, a path connecting any two units consists of a sequence of such operations. As a consequence, the chain rule along a single path is the product of the partial derivatives of all linear and nonlinear transformations along the path. For two units i and j in *subsequent* layers we have $\partial x_j / \partial x_i = w_{ij} \cdot f'(z_j)$, whereas for any two generic units i and c connected by a set of paths P_{ic} the partial derivative is sum of the product of all weights and all derivatives of the nonlinearities that are found along the each path. We introduce a notation to indicate a modified chain-rule, where the derivative of the nonlinearities $f'()$ is replaced by a function $g()$:

$$\frac{\partial^g x_c}{\partial x_i} = \sum_{P_{ic}} \left(\prod w_{yj} \prod g(z_j) \right) \quad (5)$$

When $g() = f'()$ this is the definition of partial derivative of the output of unit c with respect to unit i , computed as the sum of contributions over all paths connecting the two units. Given that a zero weight can be used for non-existing path this is valid for any architecture that involves fully-connected, convolutional or recurrent layers without multiplicative units, as well as for pooling operations. Table 1 reports the definition of the six attribution methods discussed in Section 2, using Definition 5 to reformulate LRP and DeepLIFT. These formulations are based on the following propositions.

Method	Attribution $R_i^c(x)$	Example of attributions on MNIST			
		ReLU	Tanh	Sigmoid	Softplus
Saliency Maps	$\left \frac{\partial S_c(x)}{\partial x_i} \right $				
Gradient * Input	$x_i \cdot \frac{\partial S_c(x)}{\partial x_i}$				
<u>ϵ-LRP</u>	$x_i \cdot \frac{\partial^g S_c(x)}{\partial x_i}, \quad g = \frac{f(z)}{z}$				
<u>DeepLIFT</u>	$(x_i - \bar{x}_i) \cdot \frac{\partial^g S_c(x)}{\partial x_i}, \quad g = \frac{f(z) - f(\bar{z})}{z - \bar{z}}$				
Integrated Gradient	$(x_i - \bar{x}_i) \cdot \int_{\alpha=0}^1 \frac{\partial S_c(\tilde{x})}{\partial \tilde{x}_i} \Big _{\tilde{x}=\bar{x}+\alpha(x-\bar{x})} d\alpha$				
<u>Occlusion-1</u>	$x_i \cdot \int_{\alpha=0}^1 \frac{\partial S_c(\tilde{x})}{\partial \tilde{x}_i} \Big _{\tilde{x}=x[x_i=\alpha \cdot x_i]} d\alpha$				

Table 1: Definition of several attribution methods. Underlined methods have been reformulated to highlight the connection to the gradient chain-rule. On the right, examples of attributions on a sample from the MNIST dataset [9]. Details about the network architectures can be found in Appendix A.

It has been previously proven that, in the case of ReLUs, ϵ -LRP is equivalent to Gradient * Input [18, 4]. The connection of ϵ -LRP to the gradient chain-rule can be generalized to any activation functions as follows.

Proposition 1. *ϵ -LRP is equivalent the feature-wise product of the input and the modified partial derivative $\partial^g S_c(x) / \partial x_i$, with $g = g^{LRP} = f_i(z_i) / z_i$, i.e. the ratio between the output and the input at each nonlinearity.*

Proof. For the following proof we refer to the ϵ propagation rule defined as in Equation 56 of [3], according to which the bias terms can be assigned part of the relevance. We also assume the stabilizer term $\epsilon \cdot \text{sign}(\sum_{i'} (z_{i'j} + b_j))$ at the denominator of Equation 2 is small enough to be neglected.

We proceed by induction. By definition, the ϵ -LRP relevance of the target neuron c on the top layer L is defined to be equal to the output of the neuron itself, S_c :

$$r_c^{(L)} = S_c(x) = f\left(\sum_j w_{cj}^{(L,L-1)} x_j^{(L-1)} + b_c\right) \quad (6)$$

The relevance of the parent layer is:

$$\begin{aligned} r_j^{(L-1)} &= r_c^{(L)} \frac{w_{cj}^{(L,L-1)} x_j^{(L-1)}}{\sum_{j'} w_{cj'}^{(L,L-1)} x_{j'}^{(L-1)} + b_c} && \blacktriangleright \text{LRP prop. rule (Eq. 2)} \\ &= f\left(\sum_{j'} w_{cj'}^{(L,L-1)} x_{j'}^{(L-1)} + b_c\right) \frac{w_{cj}^{(L,L-1)} x_j^{(L-1)}}{\sum_{j'} w_{cj'}^{(L,L-1)} x_{j'}^{(L-1)} + b_c} && \blacktriangleright \text{replacing Eq. 6} \\ &= g^{LRP} \left(\sum_{j'} w_{cj'}^{(L,L-1)} x_{j'}^{(L-1)} + b_c\right) w_{cj}^{(L,L-1)} x_j^{(L-1)} && \blacktriangleright \text{by definition of } g^{LRP} \\ &= \frac{\partial g^{LRP} S_c(x)}{\partial x_j^{(L-1)}} x_j^{(L-1)} && \blacktriangleright \text{by definition of } \partial^g \text{ (Eq. 5)} \end{aligned}$$

For the inductive step, we start from the hypothesis that on a generic layer l the LRP explanation is:

$$r_i^{(l)} = \frac{\partial g^{LRP} S_c(x)}{\partial x_i^{(l)}} x_i^{(l)} \quad (7)$$

then for layer $l-1$ it holds:

$$\begin{aligned} r_j^{(l-1)} &= \sum_i r_i^{(l)} \frac{w_{ij}^{(l,l-1)} x_j^{(l-1)}}{\sum_{j'} w_{ij'}^{(l,l-1)} x_{j'}^{(l-1)} + b_i} && \blacktriangleright \text{LRP propagation rule (Eq. 2)} \\ &= \sum_i \frac{\partial g^{LRP} S_c(x)}{\partial x_i^{(l)}} \frac{x_i^{(l)}}{\sum_{j'} w_{ij'}^{(l,l-1)} x_{j'}^{(l-1)} + b_i} w_{ij}^{(l,l-1)} x_j^{(l-1)} && \blacktriangleright \text{replacing Eq. 7} \\ &= \sum_i \frac{\partial g^{LRP} S_c(x)}{\partial x_i^{(l)}} g_i^{LRP} w_{ij}^{(l,l-1)} x_j^{(l-1)} && \blacktriangleright \text{by definition of } g^{LRP} \\ &= \frac{\partial g^{LRP} S_c(x)}{\partial x_j^{(l-1)}} x_j^{(l-1)} && \blacktriangleright \text{by definition of } \partial^g \text{ (Eq. 5)} \end{aligned}$$

□

Proposition 2. *DeepLIFT (Rescale) is equivalent to the feature-wise product of the $x - \bar{x}$ and the modified partial derivative $\partial^g S_c(x)/\partial x_i$, with $g = g^{DL} = (f_i(z_i) - f_i(\bar{z}_i))/(z_i - \bar{z}_i)$, i.e. the ratio between the difference in output and the difference in input at each nonlinearity, for a network provided with some input x and some baseline input \bar{x} defined by the user.*

Proof. Similarly to how a chain rule for gradients is constructed, DeepLIFT computes a multiplicative term, called "multiplier", for each operation in the network. These terms are chained to compute a

global multiplier between two given units by summing up all possible paths connecting them. The chaining rule, called by the authors "chain rule for multipliers" (Eq. 3 in [17]) is identical to the chain rule for gradients, therefore we only need to prove that the multipliers are equivalent to the terms used in the computation of our modified backpropagation.

Linear operations. For Linear and Convolutional layers implementing operations of the form $z_j = \sum_i (w_{ij} \cdot x_i) + b_j$, the DeepLIFT multiplier is defined to be $m = w_{ij}$ (Sec. 3.5.1 in [17]). In our formulation the gradient of linear operations is not modified, hence it is $\partial z_i / \partial x_i = w_{ij}$, equal to the original DeepLIFT multiplier.

Nonlinear operations. For a nonlinear operation with a single input of the form $x_i = f(z_i)$ (i.e. any nonlinear activation function), the DeepLIFT multiplier (Sec. 3.5.2 in Shrikumar et al. [17]) is:

$$m = \frac{\Delta x}{\Delta z} = \frac{f(z_i) - f(\bar{z}_i)}{z_i - \bar{z}_i} = g^{DL} \quad (8)$$

Nonlinear operations with multiple inputs (eg. 2D pooling) are not address in [17]. For these, we keep the original operations' gradient unmodified as in the DeepLIFT public implementation.³

□

This formulation has an important practical advantage. Since all modern frameworks for graph computation, like TensorFlow [1], implement backpropagation for efficient computation of the chain rule, it is possible to implement all methods above by just overriding the gradient of the graph nonlinearities, with no need to implement custom layers or operations. Moreover, thanks to automatic differentiation, all methods can be applied to any architecture trainable via backpropagation.

This formulation also enables a direct comparison between methods. We have already mentioned that ϵ -LRP is equivalent to Gradient * Input if only ReLUs are used as nonlinearities [18, 4]. We can also easily see that ϵ -LRP is equivalent to DeepLIFT computed with a zero baseline, if applied to a network with no additive biases and with nonlinearities f such that $f(0) = 0$ (eg. ReLU or Tanh). The proof comes directly from the observation that, in a network satisfying these conditions, the propagation of the baseline produces a zero reference value for *all* hidden units (ie. $\forall i : \bar{z}_i = f(\bar{z}_i) = 0$). Then $g^{LRP} = g^{DL}$. Notice that $g^{LRP}(z) = (f(z) - 0)/(z - 0)$ which, in the case of ReLU and Tanh, is the *average gradient* of the nonlinearity in $[0, z]$. It also easy to see that $\lim_{z \rightarrow 0} g^{LRP}(z) = f'(0)$, which explain why g can not assume arbitrarily large values as $z \rightarrow 0$, even without stabilizers. On the contrary, if the discussed condition on the nonlinearity is not satisfied, for example with Sigmoid or Softplus, we found empirically that ϵ -LRP fails to produce meaningful results (see Section 4). We speculate this is due to the fact $g^{LRP}(z)$ can become extremely large for small values of z , being its upper-bound only limited by the stabilizer. This causes attribution values to concentrate on a few features as shown in Table 1. Notice also that the interpretation of g^{LRP} as average gradient of the nonlinearity does not hold in this case, which explains why ϵ -LRP diverges from other methods⁴.

DeepLIFT and Integrated Gradients are related as well. While Integrated Gradients computes the average partial derivative of each feature as the input varies from a baseline to its final value, DeepLIFT approximates this quantity in a single step by replacing the gradient at each nonlinearity with its average gradient. Although the chain rule does not hold in general for average gradients, we show empirically in Section 4 that DeepLIFT is often a good approximation of Integrated Gradients.

4 Empirical comparison

As pointed out by others [23, 15], attributions methods are hard to evaluate empirically because it is difficult to distinguish errors of the model from errors of the attribution method explaining the model. For this reason, the final evaluation is often qualitative, based on the inspection of the produced

³DeepLIFT public code: <https://github.com/kundajelab/deeplift>. Retrieved on 25 Sept. 2017

⁴We are not claiming any general superiority of gradient-based methods but rather observing that ϵ -LRP can only be considered gradient-based for precise choices of the nonlinearities. In fact, there are backpropagation-based attribution methods, not directly interpretable as gradient methods, that exhibit other desirable properties. For a discussion about advantages and drawbacks of gradient-based methods we refer the reader to [17, 13, 23].

attribution maps. We argue, however, that this introduces a strong bias in the evaluation: as humans, one would judge more favorably methods that produce explanations closer to his own expectations, at the cost of penalizing those methods that might more closely reflect the network behavior. Instead, we perform a quantitative evaluation using a variation of the perturbation method by [15]. This method starts from the assumption that a good attribution method should identify features that are more relevant for the activation of the target neuron. As such, these relevant features should produce a sensible variation on the target value when removed or altered.

Our evaluation procedure is the following. We sort the input features according to their attributions and, based on this ranking, we proceed iteratively, altering one feature at each iteration and replacing it with a zero baseline until the input is left with none of the original features. At each iteration we run a forward pass with the modified input and measure the target output. The only exception to this procedure is on Inception V3, where we remove 100 features at each step to keep the evaluation time reasonable. Our method differs from [15] by two aspects: 1) we always remove features based on their ranking and we do not assume our inputs are natural images, where local pixel coherence can be exploited to alter local patches; 2) while [15] performs this test removing most relevant features first (MoRF), we also test the inverse ranking, removing least relevant features first (LeRF), as done by [2] in the context of text classification. This allows us to test whether the sign information provided by some attribution method is meaningful. We expect removing features with positive (negative) attribution to cause a drop (rise) in the target output. The resulting perturbation curves are evaluated with two criteria: firstly, the *initial steepness* of the curve is expected to be stronger for methods that better capture the true underlying attributions. Secondly, better methods are expected to produce a *larger variation* in the target output when several features are removed. Results are compared towards a baseline constructed by removing features in random order.

Figure 3 shows the results for all methods of Table 1. We use the well-known MNIST dataset [9] to test how the methods behave with two different architectures (a Multilayer Perceptron (MLP) and a CNN) and four different activation functions. We also test a simple CNN for image classification on CIFAR10 [6] and the more complex Inception V3 architecture [24] on ImageNet-like [14] samples. Finally, we test a model for sentiment classification from text data. For this task we use the IMDB dataset [10], applying both a shallow MLP and a LSTM model. Details about the architectures can be found in Appendix A. Notice that it was not our goal, nor a requirement, to reach the state-of-the-art in these tasks since attribution methods should be applicable to any model that produces some meaningful prediction. From the results we can formulate a few considerations:

- **Sign matters.** In most of our tests, the perturbation of features with negative attribution causes the target output to increase. This suggests that input samples can contain negative evidence and that these can be detected by signed attribution methods. On the other hand, on complex models like Inception V3, all gradient-based methods show low accuracy in predicting the attribution sign, leading to heatmaps affected by high-frequency noise (Figure 2) and poor performance in the LeRF test. The importance of sign information is particularly evident on simple models (eg. IMDB-MLP).
- **Saliency Maps detect neutral regions.** While Saliency Maps (MoRF) performs poorly compared to other methods (likely because sign information is discarded) the LeRF test shows that features with lowest (absolute) attribution are in fact neutral to the output. On some models, like MNIST w/ Relu, Saliency Maps identifies large portions of the image that can be removed without significantly affecting the target output. On Inception V3, Saliency Maps gains in performance since other methods fail to correctly identify sign information.
- **Occlusion-1 better identifies the few most important features.** This is a natural consequence of the fact that attributions for this methods are computed using the same procedure used for testing. Notice, however, that some other methods perform better when a relevant percentage of features has been removed, leading to a larger maximum variation of the output. This suggests that methods like Integrated Gradient, DeepLIFT and LRP are able to capture cross-interactions between different features that Occlusion-1, by considering a single feature at the time, cannot capture.
- **Integrated Gradients and DeepLIFT have very high correlation,** suggesting that the latter is a good (and faster) approximation of the former in practice. This does not hold with multiplicative interactions between features (eg. IMDB (LSTM)), where DeepLIFT performs sensibly worse.
- **ϵ -LRP is equivalent to Gradient * Input when all nonlinearities are ReLUs, while it fails when these are Sigmoid or Softplus.** When the nonlinearities are such that $f(0) \neq 0$, ϵ -LRP diverges from other methods, cannot be seen as a discrete gradient approximator and may lead to

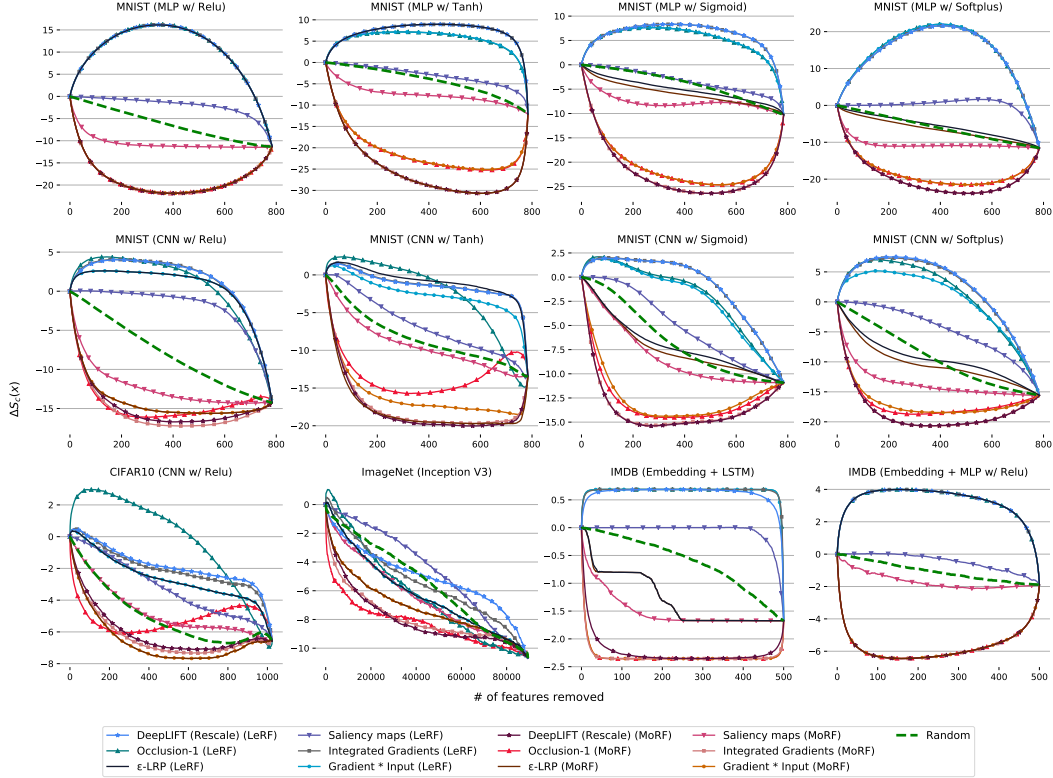


Figure 3: Results of the perturbation test for several attribution methods on various datasets and architectures. The variation of the target activation with respect to the original input is reported against the number of features removed. Best seen in electronic form.

numerical instabilities for small values of the stabilizer (Section 3). It has been shown, however, that adjusting the propagation rule for multiplicative interactions and avoiding critical nonlinearities, ϵ -LRP can be applied to LSTM networks, obtaining interesting results [2]. Unfortunately, these changes obstacle the formulation as modified chain-rule and make *ad-hoc* implementation necessary.

- **All methods are equivalent when the model behaves linearly.** On IMDB (MLP), where we used a very shallow network, all signed attribution methods perform equally well. It is easy to show that all signed methods in Table 1 produce the exact same attributions if applied to a linear model. While the opposite implication does not hold in general, we believe it is unlikely to observe this behavior with a highly nonlinear model.

5 Conclusions

We strongly believe a better understanding of existing attribution methods is necessary to highlight research direction for DNN interpretability. In this work, we have shown that Saliency Maps, Gradient * Input, ϵ -LRP, Integrated Gradients and DeepLIFT (Rescale) are strongly related under a unified framework for gradient-based attribution methods. We have also shown that ϵ -LRP and DeepLIFT, in their most simple formulation, can be easily implemented by overriding some operations of the gradient chain-rule, which is particularly convenient with modern libraries for automatic differentiation. We have proven some theoretical conditions of equivalence or approximation between these methods and conducted a quantitative evaluation to highlight their strengths and limitations.

Source code and acknowledgments A framework including the analyzed methods can be downloaded at <https://goo.gl/2R5YGS>. This work was partially funded by the Swiss Commission for Technology and Innovation (CTI, Grant No. 19005.1 PFES-ES).

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, A. Vijayaraj, V. Vijay Panchanathan, D. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] L. Arras, G. Montavon, K.-R. Müller, and W. Samek. Explaining recurrent neural network predictions in sentiment analysis. *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 2017.
- [3] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one*, 10(7):e0130140, 2015.
- [4] P. Kindermans, K. Schütt, K. Müller, and S. Dähne. Investigating the influence of noise and distractors on the interpretation of neural networks. *CoRR*, abs/1611.07270, 2016.
- [5] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [6] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.
- [7] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proc. of NIPS*, pages 1097–1105, 2012.
- [8] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [9] Y. LeCun, C. Cortes, and C. J. Burges. The mnist database of handwritten digits, 1998.
- [10] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 142–150. Association for Computational Linguistics, 2011.
- [11] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [12] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller. Explaining nonlinear classification decisions with deep Taylor decomposition. *Pattern Recognition*, 65:211–222, 2017.
- [13] G. Montavon, W. Samek, and K.-R. Müller. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73(Supplement C):1 – 15, 2018.
- [14] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [15] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, and K.-R. Müller. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems*, 2016.
- [16] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra. Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization. *CoRR*, abs/1610.02391, 2016.
- [17] A. Shrikumar, P. Greenside, and A. Kundaje. Learning important features through propagating activation differences. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3145–3153. International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [18] A. Shrikumar, P. Greenside, A. Shcherbina, and A. Kundaje. Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713*, 2016.
- [19] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [20] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *ICLR Workshop*, 2014.
- [21] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- [22] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. *ICLR 2015 Workshop*, 2014.

- [23] M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3319–3328, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [24] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [25] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- [26] M. D. Zeiler. Adadelata: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.
- [27] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
- [28] L. M. Zintgraf, T. S. Cohen, T. Adel, and M. Welling. Visualizing deep neural network decisions: Prediction difference analysis. 2017.

A Experiments setup

A.1 MNIST

The MNIST dataset [9] was pre-processed to normalize the input images between -1 (background) and 1 (digit stroke). We trained both a DNN and a CNN, using four activation functions in order to test how attribution methods generalize to different architectures. The layers of the two architectures are listed below. The activations functions are defined as $ReLU(x) = \max(0, x)$, $Tanh(x) = \sinh(x)/\cosh(x)$, $Sigmoid(x) = 1/(1+e^{-x})$ and $Softplus(x) = \ln(1 + e^x)$ and have been applied to the output of the layers marked with † in the tables below. The networks were trained using Adadelata [26] and early stopping. We also report the final test accuracy.

MNIST MLP	MNIST CNN		Test set accuracy (%)		
				MLP	CNN
Dense† (512)	Conv 2D† (3x3, 32 kernels)	Conv 2D† (3x3, 64 kernels)	ReLU	97.9	99.1
Dense† (512)	Max-pooling (2x2)		Tanh	98.1	98.8
Dense (10)	Dense† (128)		Sigmoid	98.1	98.6
	Dense (10)		Softplus	98.1	98.8

A.2 CIFAR-10

The CIFAR-10 dataset [6] was pre-processed to normalized the input images in range [-1; 1]. As for MNIST, we trained a CNN architecture using Adadelata and early stopping. For this dataset we only used the $ReLU$ nonlinearity, reaching a final test accuracy of 80.5%. For gradient-based methods, the attribution of each pixel was computed summing up the attribution of the 3 color channels and Occlusion-1 was performed setting all color channels at zero at the same time for each pixel being tested. Similarly, for the final evaluation, each pixel was considered an atomic feature and all color channels replaced at the same time.

CIFAR-10 CNN
Conv 2D† (3x3, 32 kernels)
Conv 2D† (3x3, 32 kernels)
Max-pooling (2x2)
Dropout (0.25)
Conv 2D† (3x3, 64 kernels)
Conv 2D† (3x3, 64 kernels)
Max-pooling (2x2)
Dropout (0.25)
Dense† (256)
Dropout (0.5)
Dense (10)

A.3 Inception V3

We used a pre-trained Inception V3 network. The details of this architecture can be found in [24]. We used a test dataset of 1000 ImageNet-like images, normalized in $[-1; 1]$ that was classified with 95.9% accuracy. When computing attributions, the color channels were handled as for CIFAR-10.

A.4 IMDB

We trained both a shallow MLP and a LSTM network on the IMDB dataset [10] for sentiment analysis. For both architectures we trained a small embedding layer considering only the 5000 most frequent words in the dataset. We also limited the maximum length of each review to 500 words, padding shorter ones when necessary. We used *ReLU* nonlinearities for the hidden layers and trained using Adam [5] and early stopping. The final test accuracy is 87.3% on both architectures. For gradient-based methods, the attribution of each word was computed summing up the attributions over the embedding vector components corresponding to the word. Similarly, Occlusion-1 and the final evaluation were performed setting all components of the embedding vector at zero for each word to be tested.

IMDB MLP	IMDB LSTM
Embedding (5000x32)	Embedding (5000x32)
Dense (250)	LSTM (64)
Dense (1)	Dense (1)