

Layer-wise Relevance Propagation for Neural Networks with Local Renormalization Layers

Alexander Binder¹, Grégoire Montavon², Sebastian Lapuschkin³,
Klaus-Robert Müller^{2,4}, and Wojciech Samek³

¹ ISTD Pillar, Singapore University of Technology and Design

² Machine Learning Group, Technische Universität Berlin

³ Machine Learning Group, Fraunhofer Heinrich Hertz Institute

⁴ Department of Brain and Cognitive Engineering, Korea University

Abstract. Layer-wise relevance propagation is a framework which allows to decompose the prediction of a deep neural network computed over a sample, e.g. an image, down to relevance scores for the single input dimensions of the sample such as subpixels of an image. While this approach can be applied directly to generalized linear mappings, product type non-linearities are not covered. This paper proposes an approach to extend layer-wise relevance propagation to neural networks with local renormalization layers, which is a very common product-type non-linearity in convolutional neural networks. We evaluate the proposed method for local renormalization layers on the CIFAR-10, Imagenet and MIT Places datasets.

Keywords: Neural Networks, Image Classification, Interpretability

1 Introduction

Artificial neural networks enjoy increasing popularity for image classification tasks. They have shown excellent performance in large scale competitions [4]. One reason is the ability to train neural networks with millions of training samples by parallelizing them on GPU hardware. This allows to use numbers of training samples which match the large number of parameters in deep neural networks. However, understanding what region of the image is important for a classification decision, is still an open question for neural networks, as well as for many other non-linear models. The work of [1] proposed Layer-wise Relevance Propagation (LRP) as a solution for explaining what pixels of an image are relevant for reaching a classification decision. This was done for neural networks, bag of word models [2, 11], and in a subsequent work [5], for Fisher vectors.

This paper proposes an approach to extend LRP to neural networks with non-linearities beyond the commonly used neural network formulation. One example of such nonlinearities are local renormalization layers which can not be handled by standard LRP [1]. The presented approach is based on first (or higher) order Taylor expansion. We consider a classification setup with real-valued outputs. A classifier f is a mapping of an input space $f : X \rightarrow \mathbb{R}$ such that $f(x) > 0$ denotes the presence of the class.




Fig. 1. Pixel-wise decompositions for classes wolf, frog and wolf using a neural network pretrained for the 1000 classes of the ILSVRC challenge.

2 Layer-wise Relevance Propagation for Neural Networks

In the following we consider neural networks consisting of layers of neurons. The output x_j of a neuron j is a non-linear activation function g as given by

$$x_j = g\left(\sum_i w_{ij}x_i + b\right) \quad (1)$$

Given an image x and a classifier f the aim of layer-wise relevance propagation is to assign each pixel p of x a pixel-wise relevance score $R_p^{(1)}$ such that

$$f(x) \approx \sum_p R_p^{(1)} \quad (2)$$

Pixels p with $R_p^{(1)} < 0$ contain evidence against the presence of a class, while $R_p^{(1)} > 0$ is considered as evidence for the presence of a class. These pixel-wise relevance scores can be visualized as an image called *heatmap* (see Fig. 1 for examples). Obviously, many possible such decompositions exist which satisfy equation 2. The work of [1] yield pixel-wise decompositions which are consistent with evaluation measures [9] and human intuition.

Assume that we know the relevance $R_j^{(l+1)}$ of a neuron j at network layer $l+1$ for the classification decision $f(x)$, then we like to decompose this relevance into messages $R_{i \leftarrow j}^{(l,l+1)}$ sent to those neurons i at the layer l which provide inputs to neuron j such that equation 3 holds.

$$R_j^{(l+1)} = \sum_{i \in (l)} R_{i \leftarrow j}^{(l,l+1)} \quad (3)$$

We can then define the relevance of a neuron i at layer l by summing all messages from neurons at layer $l+1$ as in equation 4

$$R_i^{(l)} = \sum_{j \in (l+1)} R_{i \leftarrow j}^{(l,l+1)} \quad (4)$$

Equations 3 and 4 define the propagation of relevance from layer $l+1$ to layer l . The relevance of the output neuron at layer M is $R_1^{(M)} = f(x)$. The pixel-wise scores are the resulting relevances of the input neurons $R_d^{(1)}$.

The work in [1] established two formulas for computing the messages $R_{i \leftarrow j}^{(l,l+1)}$. The first formula called ϵ -rule is given by

$$R_{i \leftarrow j}^{(l,l+1)} = \frac{z_{ij}}{z_j + \epsilon \cdot \text{sign}(z_j)} R_j^{(l+1)} \quad (5)$$

with $z_{ij} = (w_{ij}x_i)^p$ and $z_j = \sum_{k:w_{kj} \neq 0} z_{kj}$. The variable ϵ is a ‘‘stabilizer’’ term whose purpose is to avoid numerical degenerations when z_j is close to zero, and which is chosen to be small. The second formula called β -rule is given by

$$R_{i \leftarrow j}^{(l,l+1)} = \left((1 + \beta) \frac{z_{ij}^+}{z_j^+} - \beta \frac{z_{ij}^-}{z_j^-} \right) R_j^{(l+1)} \quad (6)$$

where the positive and negative weighted activations are treated separately. The variable β controls how much inhibition is incorporated in the relevance redistribution. A fairly large value for β (e.g. $\beta = 1$) leads to sharper heatmaps. In both formulas the message $R_{i \leftarrow j}^{(l,l+1)}$ has the following structure

$$R_{i \leftarrow j}^{(l,l+1)} = v_{ij} R_j^{(l+1)} \quad \text{with} \quad \sum_i v_{ij} = 1 \quad (7)$$

The meaningfulness of the resulting pixel-wise decomposition for the input layer $R_d^{(1)}$ comes from the fact that the terms v_{ij} are derived from the weighted activations $w_{ij}x_i$ of the input neurons. Note that layer-wise relevance propagation does not use gradients in contrast to backpropagation during the training phase. For full details on layer-wise relevance propagation the reader is referred to [1]. An implementation of LRP can be found in [6] and downloaded from www.heatmapping.org.

3 Extending LRP to local renormalization layers

We consider a general neuron j whose pooling and activation does not fit into the structure given by equation 1, and consequently, intuition for a possible redistribution formula is lacking. In this paper we propose a strategy for such neurons, based on the Taylor expansion of its activation function. A Taylor-based approach was used in [7] for decomposing ReLU neurons by exploiting their local linearity. Here, we consider instead fully nonlinear neurons.

Suppose we can define for each neuron i input to neuron j a term v_{ij} which is derived from its activation x_i such that $\sum_i v_{ij} = 1$. Then we can define a message $R_{i \leftarrow j}^{(l,l+1)} = v_{ij} R_j^{(l+1)}$. Such messages were used in equations 5 and 6 where the weighting v_{ij} was chosen to depend on the weighted activations of neuron i : $v_{ij} = c(w_{ij}x_i)^p$ and $v_{ij} = c_1 z_{ij}^+ + c_2 z_{ij}^-$, respectively. For differentiable neurons, such weighting can be obtained by performing a first order Taylor expansion. Let $x_j = g(x_{h_1}, \dots, x_{h_n})$ be a nonlinear activation function. Then, by Taylor expansion at some reference point $(\tilde{x}_{h_1}, \dots, \tilde{x}_{h_n})$, we get

$$x_j \approx g(\tilde{x}_{h_1}, \dots, \tilde{x}_{h_n}) + \sum_{i \leftarrow j} \frac{\partial g}{\partial x_{h_i}}(\tilde{x}_{h_1}, \dots, \tilde{x}_{h_n})(x_{h_i} - \tilde{x}_{h_i}). \quad (8)$$

Elements of the sum can be assigned to incoming neurons, and the zero-order term can be redistributed equally between them, leading to the decomposition

$$\forall i \leftarrow j : z_{ij} = \frac{1}{n} g(\tilde{x}_{h_1}, \dots, \tilde{x}_{h_n}) + \frac{\partial g}{\partial x_{h_i}}(\tilde{x}_{h_1}, \dots, \tilde{x}_{h_n})(x_{h_i} - \tilde{x}_{h_i}) \quad (9)$$

of the neuron activation onto its input neurons. Local renormalization layers have been shown to improve the performance in deep neural networks [4]. Consider the local renormalization y_k of a neuron x_k by the set of its surrounding neurons $\{x_1, \dots, x_n\}$ as

$$y_k(x_1, \dots, x_n) = \frac{x_k}{(1 + b \sum_{i=1}^n x_i^2)^c} \quad (10)$$

This interaction can be modeled by a layer in the network that has an activation function as given in equation 10. Local renormalization layers represent a non-linearity which cannot be tackled exactly by LRP as introduced in [1], however the strategy proposed above can be applied.

One choice to be made is the point at which to perform the Taylor expansion. There are two apparent candidates, firstly the actual input to the renormalization layer $z_1 = (x_1, \dots, x_n)$ and, secondly, the input corresponding to the case when only the neuron k fires which is to be normalized $z_2 = (0, \dots, 0, x_k, 0, \dots, 0)$. The partial derivative of y at z_2 is zero for all variables x_i with $i \neq k$ due to

$$\frac{\partial y_k}{\partial x_j} = \frac{\delta_{kj}}{(1 + b \sum_{i=1}^n x_i^2)^c} - 2bc \frac{x_k x_j}{(1 + b \sum_{i=1}^n x_i^2)^{c+1}} \quad (11)$$

This implies that the Taylor approximation has no off-diagonal contribution.

$$y_k(z_1) \approx y_k(z_2) + 0 = \frac{x_k}{(1 + b x_k^2)^c} \quad (12)$$

Therefore we apply the Taylor series around the point z_1 :

$$y_k(z_2) \approx y_k(z_1) + \nabla y_k(z_1) \cdot (z_2 - z_1) \quad (13)$$

$$\Rightarrow y_k(z_1) \approx y_k(z_2) + \nabla y_k(z_1) \cdot (z_1 - z_2) \quad (14)$$

$$\Rightarrow y_k(z_1) \approx \frac{x_k}{(1 + b x_k^2)^c} - 2bc \sum_{j:j \neq k} \frac{x_k x_j^2}{(1 + b \sum_{i=1}^n x_i^2)^{c+1}} \quad (15)$$

This weighting satisfies the following qualitative properties: for the neuron input x_k which is to be normalized, the sign of the relevance is kept. For suppressing neighboring neurons x_i , $i \neq k$, the sign of the relevance can be flipped in line with their suppressing property. The absolute value of the relevance received by the suppressing neurons is proportional to the square of their input. In the limits $c \rightarrow 0$ and $b \rightarrow 0$, the local renormalization converges against the identity, and the approximation recovers the identity. A baseline to compare against is to treat the normalization as constant. In that case the weights v_{ij} for the relevance propagation in equation 3 become a zero one vector, the relevance is propagated only to that neuron which is to be normalized: $v_{ij} = 1$ if and only if i is the neuron which is to be normalized by neuron j .

4 Experiments

We need to define a measure for meaningfulness and quality of a pixel-wise decomposition in order to evaluate the various strategies to compute it. Here we




Fig. 2. Decrease of classification score as pixels are sequentially replaced by random noise on the CIFAR-10 dataset. Red curve: pixels with highest pixel-wise scores are flipped first. Blue curve: pixels are flipped in random order. Green curve: least relevant pixels are flipped first. A similar comparison for Imagenet is found in [9].

use an idea from [9]: A pixel p is considered highly relevant for the classification score $f(x)$ of the image x if modifying it by assigning it a random RGB value $\tilde{x}(p)$, and classifying the modified image $\tilde{x}_p = x \setminus \{x(p)\} \cup \{\tilde{x}(p)\}$ results in a strong decrease of the real-valued classification score $f(\tilde{x}_p)$. This idea can be extended by sequentially modifying pixels from the most relevant to the least relevant. The result is a graph of the prediction score $f(\tilde{x})$ as a function of the number of modified pixels. An example for some sequences which will be explained below is shown in Fig. 2. We can use these graphs to evaluate the meaningfulness of a pixel-wise decomposition.

In the first experiment we compare the measure when flipping highest-scoring pixels first, against flipping pixels in random order, and against flipping lowest scoring pixels first. If the classifier is able to identify pixels that are important for classification, then flipping highest scoring pixels first should result in the fastest decaying curve, while flipping lowest scoring pixels first should result in the slowest decrease. Fig. 2 tests this property on the CIFAR-10 dataset [3] which consists of 50000 images of size 32×32 drawn from 10 object classes. Scores are averaged over the 5000 images of the test set of CIFAR-10 for a classifier in which local renormalization layers are treated as the identity during computation of pixel-wise scores. Experiments corroborate that flipping highest scoring pixels first results in the fastest decrease of the prediction score on average over the test set. The decrease is sharper compared to random flipping, or flipping lowest scoring pixels first.

In a second experiment we compare which treatment of the local renormalization layer is best to identify those pixels that are most relevant for classifying an image. The two tested approaches for treating the local renormalization are (1) like it would be the identity, (2) by first order Taylor expansion as given by equation 15. These approaches are furthermore tested when used in conjunction with the two methods proposed by [1], namely, the ϵ -rule in equation 5 with a fixed value of the numerical stabilizer ϵ , and the β -rule shown in equation 6, with fixed β .

rule for basic layers	rule for normalization layers	AUC score
eq. 4,5, $\epsilon = 0.01$	identity	37.10
eq. 4,5, $\epsilon = 0.01$	first-order Taylor	35.47
eq. 4,6, $\beta = 1$	identity	56.13
eq. 4,6, $\beta = 1$	first-order Taylor	53.82

Table 1. Comparison of different types of LRN layer treatments for two approaches of computing pixel-wise scores for CIFAR-10. Lower scores are better.

dataset	methods	$\Delta_{\epsilon=1}^{\epsilon=0.01}$	$\Delta_{\epsilon=0.01}^{\epsilon=100}$	$\Delta_{\epsilon=100}^{\beta=1}$	$\Delta_{\beta=1}^{\beta=0}$
Imagenet	identity	-21.29	2.75	-42.61	-49.07
	Taylor	-12.29	-41.75	-34.44	-50.76
MIT Places	identity	-20.19	12.91	-14.55	-49.37
	Taylor	-11.65	-22.55	-8.82	-48.7

Table 2. Comparison of different types of heatmap computations for Imagenet and MIT Places. We use the shortcut notation Δ_a^b for expressing $AUC_a - AUC_b$. Thus, a negative value indicates that the method produces better heatmaps with parameter a than with parameter b . Note that ϵ refers to equations 4 and 5; β refers to eq. 4 and 6.

We measure the quality of heatmaps by perturbing highest pixels first and computing the area under the curve (AUC). Lower AUC averaged over a large number of images indicates a better identification of pixel relevance by the heatmap. Results on CIFAR-10 are shown in Table 1. We observe that in all cases using first order Taylor in normalization layers improves the heatmap AUC score. This shows its effectiveness for dealing with non-linear neuron layers.

We perform the same experiments also with Imagenet [8] and MIT Places [13] datasets, each time evaluating results for 5000 images from their respective unlabeled test sets. Note that computing a heatmap requires only a predicted class label, not a ground truth. We evaluated results for the parameter settings $\beta = 0, \beta = 1$ in equation 6 and $\epsilon = 0.01, \epsilon = 1, \epsilon = 100$ in equation 5. Table 2 shows the difference of AUC between variants of LRP, when using either the identity or the Taylor expansion for local renormalization layers. We observe the following ordering starting with the lowest (best) AUC: $\epsilon = 1, \epsilon = 0.01, \epsilon = 100, \beta = 1, \beta = 0$. This order holds independent of whether we consider Imagenet or MIT places, when using Taylor for local renormalization layers. When using identity instead of Taylor, the order remain the same, except for $\epsilon = 100$ and $\epsilon = 0.01$ that are swapped. This is by itself an interesting result demonstrating that use of Taylor in the normalization layer does not disrupt the overall properties of relevance propagation techniques. For a comparison to other approaches such as heatmaps based on deconvolutions [12], or backpropagated gradients [10] we refer to [9].

Table 3 shows the difference of AUC between Taylor and identity for local renormalization layers, for various choices of datasets and LRP parameters. We observe that for the parameters with best AUC ($\epsilon = 1$ and $\epsilon = 0.01$), using Tay-

dataset	methods	$\epsilon = 1$	$\epsilon = 0.01$	$\epsilon = 100$	$\beta = 1$	$\beta = 0$
Imagenet	$AUC_{\text{Taylor}} - AUC_{\text{identity}}$	-35.84	-26.84	8.47	0.29	1.98
MIT Places	$AUC_{\text{Taylor}} - AUC_{\text{identity}}$	-33.13	-24.59	5.34	-0.39	-1.06

Table 3. Impact of using the Taylor method in various settings. Negative value indicates that using the Taylor expansion for the local renormalization is better in AUC terms (i.e. heatmaps are more representative of the importance of each pixel).

lor expansion for representing local renormalization layers further improves the AUC scores. For the remaining choices the results are on par or slightly worse. This is consistent with the interpretation of large values of ϵ as smoothing out small contributions. It is also consistent with the observation that $\beta = 1$ and $\beta = 0$ yield both smooth heatmaps in general. Heatmaps for some parameters of interest are shown in Figure 3. Taylor with $\epsilon = 1$ has both high pixel selectivity and low noise, which in agreement with its measured superiority in the quantitative experiments.

5 Conclusion

We have presented an extension of layer-wise relevance propagation (LRP) based on first-order Taylor expansions for product-type nonlinearities. Such nonlinearities occur in the local renormalization layers of deep convolutional neural networks. The proposed extension is evaluated on three popular datasets and it is shown to clearly outperform the original LRP method. In future work we will investigate the potential gain of using higher order Taylor expansions, and apply the method to a larger class of neural network layers.

References

1. Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLOS ONE*, 10(7):e0130140, 2015.
2. Gabriela Csurka, Christopher R. Dance, Lixin Fan, Jutta Willamowski, and Cdric Bray. Visual categorization with bags of keypoints. In *In Workshop on Statistical Learning in Computer Vision, ECCV*, pages 1–22, 2004.
3. Alex Krizhevsky. Learning multiple layers of features from tiny images. <http://www.cs.toronto.edu/~kriz/cifar.html>, 2009.
4. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1106–1114, 2012.
5. Sebastian Lapuschkin, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. Analyzing classifiers: Fisher vectors and deep neural networks. In *Proc. IEEE CVPR*, 2016.
6. Sebastian Lapuschkin, Alexander Binder, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. The layer-wise relevance propagation toolbox for artificial neural networks. *Journal of Machine Learning Research*, 2016. in press.



Fig. 3. Top row shows original unwarped image. Remaining rows show heatmaps produced by various parameters of the LRP method.

7. Grégoire Montavon, Sebastian Bach, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller. Explaining nonlinear classification decisions with deep Taylor decomposition. *arXiv:1512.02479*, 2015.
8. Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, pages 1–42, April 2015.
9. Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Bach, and Klaus-Robert Müller. Evaluating the visualization of what a deep neural network has learned. *arXiv:1509.06321*, 2015.
10. Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR*, abs/1312.6034, 2013.
11. Koen E. A. van de Sande, Theo Gevers, and Cees G. M. Snoek. Evaluating color descriptors for object and scene recognition. *IEEE TPAMI*, 32(9):1582–1596, 2010.
12. Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *ECCV*, pages 818–833, 2014.
13. Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. In *Adv. in NIPS*, pages 487–495, 2014.