

CONTROLLING EXPLANATORY HEATMAP RESOLUTION AND SEMANTICS VIA DECOMPOSITION DEPTH

Sebastian Bach[†], Alexander Binder^{*}, Klaus-Robert Müller[‡] and Wojciech Samek[†], Member, IEEE

[†] Fraunhofer Heinrich Hertz Institute, Einsteinufer 37, 10587 Berlin, Germany

[‡] Berlin Institute of Technology (TU Berlin), Straße des 17. Juni 135, 10623 Berlin, Germany

^{*} Singapore University of Technology (SUTD), 8 Sompah Road, Singapore 487372, Singapore

ABSTRACT

We present a comparison of the perception of input images for classification of the state of the art deep convolutional neural networks and high performing Fisher Vector predictors. Layer-wise Relevance Propagation (LRP) is a method to compute scores for the components of input points denoting their contribution to the prediction of the classifier for one particular test point. We demonstrate the impact of different choices of decomposition cut-off points during the LRP-process, controlling the degree of detail and resolution of the heatmap on test images from the PASCAL VOC 2007 test data set.

1. INTRODUCTION

Nonlinear models play an integral part in many well-predicting machine learning algorithms. They include, for example, kernel machines, artificial neural networks, and other nonlinear mapping functions and feature space transformations e.g. during preprocessing steps. Many high-performing predictors consist of multiple layers of such mappings resulting in powerful predictive capabilities, which comes at the cost of an obfuscation of the decision making. Oftentimes, however, knowledge about *how* a prediction comes to pass is as equally important as the confidence of the prediction, as it may e.g. help to identify the weaknesses of a classifier, the training data or reveal the policies followed by the predictor.

Quite recently, multiple endeavours have been made to gain insight into those black box classifiers, e.g. for neural network type classifiers by highlighting dominant filter activations [1], computing saliency maps visualizing local sensitivities [2] or the identification of support regions [3] critical to the prediction for SVM [4] classifiers with max-pooling feature mapping and the explanation of hard feature mappings with HIK kernels [5].

With Layer-wise Relevance Propagation [6] (LRP), a principled approach applicable to a wide range of classifier architectures and problem domains has been introduced. LRP is a method for explaining the output of a classifier wrt to the input data. Specifically, the method allows to generate an

explanation of how the individual components of the input in their given state cause the evaluating model to arrive at its decision. LRP assigns *relevance scores* to each input (or intermediate representation) component, which can then be visualized as a *heatmap*. In [7], the algorithm has been used to compare the perception of Deep Neural Network (DNN) classifiers and a state-of the art configuration of the improved Fisher Vector [8] (FV) classifier, demonstrating the method's general applicability to a wide range of classifier architectures. However, due to limits in the transparencies of the feature extraction process which in general apply to Bag of Feature (BoF) classifiers, heuristic steps have been incorporated into the decomposition process [7, 6, 3], leaving the decision-explaining heatmaps to appear much less sparse and at a lower resolution when compared to the output of LRP applied to DNNs. To alleviate this issue, we introduce the notion of a *mapping influence cut-off point* with allows to control the degree of detail and semantics of a heatmap. A qualitative analysis on heatmaps at different resolutions computed for images of the PASCAL VOC 2007 [17] and both predictors is performed.

2. LAYER-WISE RELEVANCE PROPAGATION

The aim of LRP is to attribute shares of upper layer relevances $R_j^{(l+1)}$ to all components i of the adjacent lower layer l , such that each component of l receives a relevance score $R_i^{(l)}$ proportionally to its contribution to the output values at layer $l+1$ when performing a forward pass. In its simplest and most general formulation, this is realized via the local decomposition rule

$$R_i^{(l)} = \sum_j \frac{z_{ij}}{z_j} R_j^{(l+1)} \quad (1)$$

with z_{ij} representing the outcome of a forward mapping operation from component i of layer l to component j in layer $l+1$ and z_j being the combined output at layer $l+1$. Note that the case $0/0$ is being treated as 0. Starting at the predictor output with $R^{(l+1)} = f(x)$, this decomposition is then performed iteratively layer-by-layer under consideration of the classifier

architecture until relevance scores at the input layer have been obtained. Several pertinent adaptations of Equation 1 specific to different forward mappings have been discussed and evaluated in [9, 6, 7], with select variants relevant to the classifiers used within the scope of this paper being remarked in below Sections 2.1 and 2.2.

2.1. LRP for Deep Neural Networks

Neural network type classifiers typically consist of a sequence of mapping layers

$$z_{ij} = x_i w_{ij} ; z_j = \sum_j z_{ij} + b_j ; x_j = g(z_j)$$

where x_i is the input, w_{ij} and b_j are the learned weights and bias terms and $g(\cdot)$ incorporates a non-linear activation or/and pooling function. This formulation of a network is in general enough to encompass a wide range of architectures, including convolution operations. Equation 1 and variants thereof are directly applicable to DNN classifiers, resulting in (sub)pixel-accurate¹ relevance scores $R_p^{(1)}$ per input pixel p . For the ConvNet models with ReLu activation layers considered throughout this paper, we apply the α/β -weighted decomposition formula with $\beta = 1$, which has been identified to result in heatmaps best representing the classifier decision in [10].

2.2. LRP for Bag of Feature Classifiers

The computations performed in the context of Bag of Feature methods fit well into the framework described by LRP. Similar to DNNs, BoF classifiers operate by executing consecutive layers of feature extraction, mappings and transformations relative to a visual vocabulary and pooling steps. The resulting vector description is then being fed into a support vector classifier for training and classification. LRP decomposes the prediction of such a pipeline in inverse direction, starting with $R^{(4)} = f(x)$. [6] points out solutions for decomposing any kernel-based classification function to compute relevances $R_d^{(3)}$ for the dimensions d of the vector representation of the input image. Also, relevance scores $R_l^{(2)}$ for all local descriptors $l \in L$ as extracted from an input image can easily be computed for sum- and max-pooled (and anything in between) mappings. In previous work, the notation $m(l)_d$ has been used to describe the output of such a mapping function m with input l to output dimension d , which corresponds to z_{ij} (as z_{ld}) in Equation 1.

Since feature extraction algorithms encode and compress the image within a the receptive field of a local descriptor, the relationship between single pixels and local descriptor dimensions – or even visual prototypes – is not inherently clear

¹DNNs receiving rgb color images as input, each color channel per pixel receives a relevance score. We use their sum as $R_p^{(1)}$ for visualizations.

in general. This is especially true for quantile-based feature descriptors. For that reason [6, 7] uniformly distribute a local feature’s relevance score across all pixels within its receptive field to compute $R_p^{(1)}$ for each input pixel. This leads to heatmaps appearing more coarse when compared to heatmaps produced from DNN classifiers (see Fig. 1. and [6, 7]), with heatmap granularity being limited by descriptor size.

Within the scope of this paper we use a FV classifier to represent the class of BoF models, in a setup as described in [11]. The mappings of features l onto a FV representation may be of both positive and negative sign, potentially destabilizing the relevance decomposition. We therefore employ the following ϵ -stabilized decomposition formula with $\epsilon = 100$. This redistribution rule has been found to compute explanatory heatmaps matching the classifier decision well [7].

3. CONTROLLING HEATMAP DETAIL AND SEMANTIC

We introduce the notion of an *mapping influence cut-off point*, describing the step from which on the forward mapping function of the classifier does not influence relevance propagation anymore and only the receptive field of the classifier does. For the DNN architecture as described in Section 2.1, there is no such cut-off point being used in previous work, whereas for the BoF architecture as of Section 2.2 the cut-off is located at the decomposition layer resulting in $R_p^{(1)}$ (for brevity, we say the cut-off is “at $R_p^{(1)}$ ”) and has been chosen out of necessity in past work. In this Section our aim is to explain how the choice of such a cut-off point may be voluntary for both considered predictor architectures by potentially increasing or decreasing the heatmap resolution for BoF models and DNNs respectively.

For neural network type classifiers, decreasing the resolution of LRP-computed heatmaps follows the procedure to compute scores $R_p^{(1)}$ for BoF classifiers. For ConvNets, instead of applying the local decomposition rule from Equation 1 to the bottom-most convolution layer, relevance scores computed for the succeeding pooling layer are to be distributed uniformly across all of its inputs. This is equivalent to applying Equation 1 and then averaging lower layer relevances for each convolution operator – or in general, choosing a cut-off point at a higher layer is equal to substituting any mappings from that layer on with flat weights. Note that for simple fully connected layers, this approach will render the resulting lower layer relevances meaningless, since spatial structure – as it is the case with convolution layers – is not present and the resulting heatmap will be uniform. Here, the w^2 -rule proposed in [9] might yield satisfactory results by distributing relevance values according to mapping weights only

To bridge the gap from local feature relevance scores to pixel relevance scores with BoF classifiers, two obstacles need to be outmanoeuvred, namely (i) one has to be able

to compute relevance scores for each dimension of a local descriptor \mathbf{l} and (ii) the local feature dimensions need to be relatable to a grid of pixel coordinates. In this work, we concentrate on the FV classification model as described in [11], where the requirements to solve (i) are given fully and (ii) partially, as explained below.

(i) Computing relevances $R_{\mathbf{l}_i}^{(2)}$ for local feature all local feature dimensions i : To compute relevance scores at such a fine granularity, the forward mapping contribution of each feature dimension needs to be known (or at least its *influence*. This, however, will not be subject of this work). For the FV model considered here, this is the case, such that an exact relevance composition can be performed for each \mathbf{l}_i . For one, [11] compute projections

$$\Psi_{\mu_k}(\mathbf{l}) = \frac{1}{\sqrt{\pi_k}} \gamma_k(\mathbf{l}) \left(\frac{\mathbf{l} - \mu_k}{\sigma_k} \right)$$

$$\Psi_{\sigma_k}(\mathbf{l}) = \frac{1}{\sqrt{2\pi_k}} \gamma_k(\mathbf{l}) \left(\frac{(\mathbf{l} - \mu_k)^2}{\sigma_k^2} - 1 \right)$$

relating \mathbf{l} to all K components of a GMM $\lambda = \{(\pi_k, \mu_k, \Sigma_k)\}_{1..K}$ fit during training wrt to its 1st and 2nd moments, with $\mathbf{l} \in \mathbb{R}^D$ and also $\Psi_{\mu_k}(\mathbf{l})$ and $\Psi_{\sigma_k}(\mathbf{l}) \in \mathbb{R}^D$. Further, the covariance matrices of the trained GMM have been constrained to be diagonal (e.g. $\forall k : \Sigma_k = \text{diag}(\sigma_k)$). We therefore know that each \mathbf{l}_i corresponds to exactly one dimension in the mapping output space of $\Psi_{\mu_k}(\mathbf{l})$ and $\Psi_{\sigma_k}(\mathbf{l})$ for all k , which are concatenated to form the full FV representation $\Psi_\lambda(\mathbf{l})$. For simplicity, suppose a function $d = \delta(i, \Psi_{\{\mu, \sigma\}_k})$, which computes for a local feature dimension i and mapping of choice the output dimension d of the FV representation. We compute relevance scores for each \mathbf{l}_i as

$$R_{\mathbf{l}_i}^{(2)} = \sum_{k=1}^K \left(\frac{\Psi_{\mu_k}(\mathbf{l})_i}{\sum_{\mathbf{l}' \in L} \Psi_{\mu_k}(\mathbf{l}')_i} R_{\delta(i, \Psi_{\mu_k})}^{(3)} + \frac{\Psi_{\sigma_k}(\mathbf{l})_i}{\sum_{\mathbf{l}' \in L} \Psi_{\sigma_k}(\mathbf{l}')_i} R_{\delta(i, \Psi_{\sigma_k})}^{(3)} \right).$$

Note that in practice we do still apply the ϵ -stabilized decomposition variant. Above decomposition describes the basic approach. Before mapping the local descriptors into the FV space, the reference model from [11] projects the initial 128-dimensional SIFT [12] features onto a 80-dimensional subspace via mapping components computed with PCA [13] during training. The now also 80-dimensional $R_{\mathbf{l}_i}^{(2)}$ can easily be projected into the original SIFT space by applying Equation 1, since the appropriate operation in the forward pass realizes a linear projection. A numerical stabilization of the denominator was necessary for good results, and $\epsilon = 100$ has been chosen as the best suited parameter after visual inspection.

(ii) Relating local feature dimensions to a spatial pixel grid: Many local feature types aggregate information extracted from an image area, such that the relation between

	aer	bic	bir	boat	bot	bus	car
F	79.1	66.4	45.9	70.9	27.6	69.7	81.0
N	88.1	79.7	80.8	77.2	35.5	72.7	86.3
	cat	cha	cow	din	dog	hor	mot
F	59.9	51.9	47.6	58.1	42.3	80.5	69.3
N	81.1	51.0	61.1	64.6	76.2	81.6	79.3
	per	pot	she	sof	tra	tvm	mAP
F	85.1	28.6	49.6	49.3	82.7	54.3	60.0
N	92.4	50.0	74.0	49.5	87.1	67.1	72.1

Table 1: Prediction performance of the **FV** and **DNN** classifiers used in this paper, in average precision (AP) per class and in percent.

feature dimension and pixel coordinate is lost in the process on the interaction between groups pixel values. An appropriate example are quantile-based local descriptors, for the computation of which all scanned pixel values are of importance, yet only a select few are grouped into the final descriptor, maybe even in an interpolated manner. However, knowledge about spatial bin placement and local feature geometry can be used whenever possible in order to intelligently merge $R_{\mathbf{l}_i}^{(2)}$ and assign relevance scores to pixels at a higher resolution. The reference FV predictor uses SIFT descriptors at different sizes with 4×4 spatial bins, with each capturing a histogram of gradient magnitudes in 8 directions. We use knowledge about the feature geometry and indexing [14, 15] to evenly distribute the summed relevances corresponding to each spatial bin evenly over the covered pixels, thus increasing the resolution of each SIFT feature’s relevance feedback 16-fold.

4. RESULTS

We compute heatmap explanations for the FV model configured after [11] and trained for [7] and the BVLC reference model from the Caffe package [16] which has been retrained for the 20 classes of PASCAL VOC 2007 [17]. Classification results for both models in average precision (AP) are listed in Table 1. For both models, we compute heatmaps with cut-off points at $R^{(1)}$ and no cut-off (as far as possible for FV. See Section 3) and show results in Figure 1. We observe that heatmaps without cut-off are still sparser for DNNs due to the limits set to the FV classifier, despite images being fed into the network at a lower resolution. Nonetheless, the higher resolution heatmaps for the FV classifier allow for a better understanding of the classifier decision when compared to the heatmaps with cut-off at $R^{(1)}$. In the example for class “chair” the high resolution relevances demonstrate that the FV classifier mostly uses the object structure itself for classification, which was difficult to read for the lower resolution heatmaps. Both models seem to follow similar higher level strategies for most object classes, e.g. with the bottom half of cars, (also wheels of the bicycles, faces and clothing of peo-

ple) or the faces of dogs capturing the models’ focus of attention. Both classifiers strongly react to defining aspects of the object, while still focussing on most of it as can be reasoned from the input images with the lower resolution heatmap controlling the alpha channel, visualizing the respective model’s focus of attention. We also observe that both classifiers seem to prefer the use of hard edges (e.g. cutlery for class table) for detection as the result of training (DNN) or due to design (FV with SIFT), rather than texture. The more deep and complex DNN classifier seems to be in general more adept at abstracting object appearances and is therefore less prone to misleading noise information while using more sophisticated rules for prediction. For the image showing a dog, for example, the FV model reacts strongly to eye-like black pebbles in the snow, next to the dog’s eyes themselves, whereas the network model is reacting to a structure resembling a dog’s face (nose below eyes). However, this increased concentration of the DNN on higher structural information includes regularly co-appearing image features into the detection process, whereas the FV model tends to concentrate on simpler rules. Note the chandelier for class “chair” in living room scenes.

5. CONCLUSION

In this work we introduce the notion of a *mapping influence cut-off point*, which allows to control the resolution of the computed heatmap and simultaneously its semantics. We have compared heatmap explanations for state-of-the-art DNN and FV models at different degrees of detail. While (very) high resolution relevance maps provide information about the kind of visual cues a classifier has been conditioned on a very local level, lower resolution heatmaps shed light to the classifier’s focus of attention by directly mapping the relevance scores assigned to its receptive fields onto pixel level. By voluntarily choosing a mapping influence cut-off point, we are able to explain decisions wrt a desired aspect of the perception of the model. In combination with image-wise predictions, heatmaps computed at different degrees of detail open up the possibility of automated assistance in different problem settings, e.g. in localizing target content upon detection, and then highlighting characteristic features. We see possible applications for instance in the medical domain, where the screening of large images of stained tissue by a medical expert might be assisted with low resolution heatmaps for malignant tissue localization and a high resolution heatmap verifying the decision of the classifier.

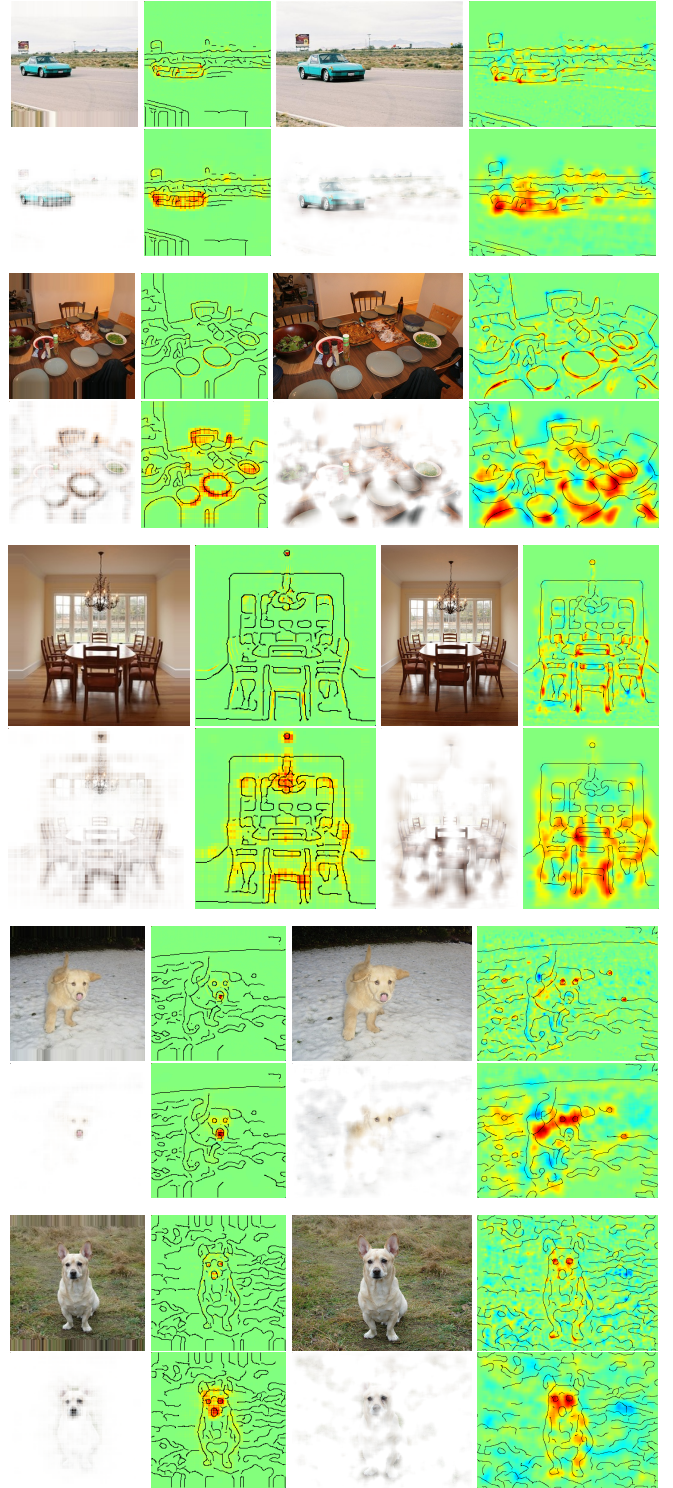


Fig. 1: Each two rows of images shows images corresponding to the CNN (left four images) and FV (right four) models. In clock-wise order: (\nwarrow) Input image, (\nearrow) high resolution heatmap without cut-off point, (\swarrow) low resolution heatmap with cut-off at $R_p^{(1)}$ and (\searrow) the same heatmap added as α -channel to the input. Green heatmap areas are rated neutral to the classifier prediction, yellow to dark red hues indicate a positive contribution to the target class with image areas marked with blue color receiving negative relevance ratings.

6. REFERENCES

- [1] Matthew D. Zeiler and Rob Fergus, “Visualizing and understanding convolutional networks,” in *ECCV*, 2014, pp. 818–833.
- [2] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman, “Deep inside convolutional networks: Visualising image classification models and saliency maps,” in *ICLR Workshop*, 2014.
- [3] Lingqiao Liu and Lei Wang, “What has my classifier learned? visualizing the classification rules of bag-of-feature model by support region detection,” in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 2012, pp. 3586–3593.
- [4] Corinna Cortes and Vladimir Vapnik, “Support-vector networks,” *Machine learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [5] Jasper RR Uijlings, Arnold WM Smeulders, and Remko JH Scha, “The visual extent of an object,” *International journal of computer vision*, vol. 96, no. 1, pp. 46–63, 2012.
- [6] 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*, vol. 10, no. 7, pp. e0130140, 2015.
- [7] Sebastian Bach, Alexander Binder, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek, “Analyzing classifiers: Fisher vectors and deep neural networks,” *arXiv preprint arXiv:1512.00172*, 2015.
- [8] Florent Perronnin, Jorge Sánchez, and Thomas Mensink, “Improving the fisher kernel for large-scale image classification,” in *Computer Vision–ECCV 2010*, pp. 143–156. Springer, 2010.
- [9] Grégoire Montavon, Sebastian Bach, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller, “Explaining nonlinear classification decisions with deep taylor decomposition,” *arXiv preprint arXiv:1512.02479*, 2015.
- [10] Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Bach, and Klaus-Robert Müller, “Evaluating the visualization of what a deep neural network has learned,” *arXiv preprint arXiv:1509.06321*, 2015.
- [11] Ken Chatfield, Victor S Lempitsky, Andrea Vedaldi, and Andrew Zisserman, “The devil is in the details: an evaluation of recent feature encoding methods,” in *BMVC*, 2011, p. 8.
- [12] David G Lowe, “Distinctive image features from scale-invariant keypoints,” *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [13] Karl Pearson, “Liii. on lines and planes of closest fit to systems of points in space,” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
- [14] Andrea Vedaldi and Brian Fulkerson, “Vlfeat: An open and portable library of computer vision algorithms,” in *Proceedings of the international conference on Multimedia*. ACM, 2010, pp. 1469–1472.
- [15] “SIFT - vl_phow,” http://www.vlfeat.org/matlab/vl_phow.html, last accessed 2016-01-31.
- [16] Yangqing Jia, “Caffe: An open source convolutional architecture for fast feature embedding,” <http://caffe.berkeleyvision.org/>, 2013.
- [17] M Everingham, L Van Gool, CKI Williams, J Winn, and A Zisserman, “The pascal visual object classes challenge 2007 (voc 2007) results (2007),” 2008.