When saliency goes off on a tangent: Interpreting Deep Neural Networks with nonlinear saliency maps

Jan Rosenzweig (✉ jan.rosenzweig@kcl.ac.uk)

King's College London

Zoran Cvetkovic

King's College London

Ivana Rosenzweig

King's College London

Research Article

Keywords:

Posted Date: February 18th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1295925/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

When saliency goes off on a tangent: Interpreting Deep Neural Networks with nonlinear saliency maps

Jan Rosenzweig, Zoran Cvetković*, Ivana Rosenzweig†

Abstract

A fundamental bottleneck in utilising complex machine learning systems for critical applications has been not knowing why they do and what they do, thus preventing the development of any crucial safety protocols. To date, no method exist that can provide full insight into the granularity of the neural network’s decision process. In past, saliency maps were an early attempt at resolving this problem through sensitivity calculations, whereby dimensions of a data point are selected based on how sensitive the output of the system is to them. However, the success of saliency maps has been at best limited, mainly due to the fact that they interpret the underlying learning system through a linear approximation. We present a novel class of methods for generating nonlinear saliency maps which fully account for the nonlinearity of the underlying learning system. While agreeing with linear saliency maps on simple problems where linear saliency maps are correct, they clearly identify more specific drivers of classification on complex examples where nonlinearities are more pronounced.

This new class of methods significantly aids interpretability of deep neural networks and related machine learning systems. Crucially, they provide a starting point for their more broad use in serious applications, where ‘why’ is equally important as ‘what’.

Main

One of the key barriers limiting the mass adoption of machine learning in serious applications is the lack of interpretability of such methods. In layman’s terms, machine learning is quite good at providing the answers along the lines of “computer says no”, but it is quite difficult to find out “why not”.

This is especially important when any decisions made by such systems have serious consequences. In a well publicized case, Amazon was forced to abandon their AI recruitment tool when it turned out it had a clear bias against female candidates [1]. Similar anecdotes revolving around machine learning systems behaving in unintended ways are a daily staple in the world of AI and machine learning, and most researchers

*Department of Engineering, King’s College London, Strand, London WC2R 2LS

†Centre for Neuroimaging Sciences, King’s College London, De Crespigny Park, Camberwell, SE5 8AF

have encountered numerous examples, ranging from amusing to distasteful. Biases in AI and machine learning have even been addressed by the World Economic Forum [2].

In essence, deep learning systems are an example of black-box systems, providing answers but no insights into how these answers were reached.

Saliency maps are a decade-old tool for interpreting deep learning systems. In computer vision and related applications, they are used to highlight the areas of the image that are most salient to its classification, which has since been generalized to other, non-vision related deep neural networks (DNN) [3].

At their core, saliency maps rely on the simple calculation of gradients. The sensitivity of the classification (output) to the image (input) is measured by the gradient of the output with respect to the input. Hence traditional saliency maps work by calculating this gradient, and then selecting the regions of the input that contribute most to the gradient. This is in no small part aided by the fact that deep learning systems already calculate all gradients through backpropagation, and hence the resulting gradients are easily available at zero or minimal additional computational cost.

This strength of saliency maps as an interpretation tool is also their greatest weakness. A deep neural network implements a nonlinear mapping of the input to the output. The gradient, however, only calculates the linear tangent to this mapping, evaluated at the input. In essence, it calculates the linear component of the DNN as a linear regressor [4].

In numerous cases, the tangent approximation is sufficient, which has lead to the rapid adoption of saliency maps as an early DNN interpretation tool.

In many other cases, however, the nonlinearities of the DNN are such that the linear tangent is a poor approximation or, worse, misleading. This is especially noticeable in attempts to use DNNs to classify medical images[6]; while saliency maps for simple cases indeed highlight the correct region of the image, indicating that they are classified more or less linearly, more complex cases generate saliency maps that invariably pick out wrong parts of the image, indicating that their classification significantly depends on the nonlinearity of the DNN[7].

This puts a fundamental ceiling on the applicability of saliency maps to DNN interpretation. Consequently, saliency maps have fallen out of favour as a DNN interpretation tool, and their use these days is mostly limited to research purposes [4, 5].

In this work, we describe a method to overcome these difficulties, by moving away from saliency maps as a linear tangent approximation, towards a fully nonlinear mapping from the input to the output.

Graphically, a saliency map can be interpreted as deselecting the pixels that do not contribute significantly to the classification, and that can therefore be removed from the image without affecting its classification. In practice, this is usually implemented by masking the non-highlighted pixels with a background colour. The masking is intended to filter out competing classifications, with the effect of increasing the probability of the target classification.

This, however, does not have the desired effect. The masking itself introduces new information to the image, and that new information competes with the desired classifica-

tion. Generally speaking, if more than 20% of the pixels have been masked, the image is invariably classified as a "piece of a puzzle" or equivalent, regardless of what is depicted in it.

We modify this approach two-fold. Firstly, instead of masking non-highlighted pixels, we subject them to a low-pass filter, effectively blurring the non-salient parts of the image. Secondly, we enhance the highlighted pixels in the direction of the gradient. By iteratively repeating this step, we are able to selectively amplify the features that are important to classifying the image into any chosen class. The corresponding saliency map is then obtained by masking any blurred pixels.

Neither of those steps is sufficient in its own right; blurring pixels without gradient enhancement only increases the entropy of the image, effectively blurring the image to the extent that it can not be reliably classified as anything. Gradient enhancement without blurring, on the other hand, reaches the boundaries of the RGB box and stops there, without being able to add significantly to the classification.

The key difference between our nonlinear saliency maps and linear saliency maps is the updating of the gradient at each step; modifying our method to re-use the same gradient, calculated at the original image, at each step, would revert to a linear saliency map. The difference between the linear and nonlinear saliency map at any given point is therefore directly linked to the local nonlinearity of the DNN at that point.

The specific optimization procedure described in this paper does not converge generally, due to irreversibility of the low-pass filtering operation. Once the local maximum is overshoot by blurring salient pixels, it is impossible for the optimization to un-blur them in order to improve the approximation.

This lack of convergence does not pose a serious problem. Even though the maximum is never reached, the step parameters can be chosen so that the trajectory passes sufficiently close to it, before possibly subsequently diverging. Hence the only remaining step is to re-trace the trajectory and return the points on the trajectory that maximised the target class probability, as various approximations of the relevant maximum.

We have performed this procedure over a range of well-known DNNs such as VGG16 [8], ResNet50 [9], MobileNet [10], InceptionV3 [11] and Xception [12], using mostly images from Imagenet ILSRVC 2017 (<https://image-net.org>), and others.

The results are striking. As shown in Figure 1, we can work with classes that score as low as 0.01% probability for a given image, and enhance the image to classify in such a class with probability higher than 95%. We thereby come as close to a pure class representation as is reasonably possible for a given image. By subsequently masking the non-highlighted pixels, we obtain the equivalent masked saliency map, but this time reflecting the full nonlinearity of the DNN, instead of only its tangent.

Looking at optimization trajectories in Figure 2, it is apparent that not all trajectories are equal. Some trajectories reach the pure class representation almost immediately from the original image, indicating that the initial gradient, calculated at the original image, is sufficient, and that nonlinearities of the DNN are relatively unimportant at that point. Other trajectories, on the other hand, only find a path to the pure class representation once 50% or 75% of the image has been blurred out. For these images,

the nonlinearities of the DNN play a more important part in their classification, and the linear approximation fails to provide a correct interpretation.

More specifically, is apparent from the optimization trajectories that there is a threshold in the proportion of the image that has to be blurred before the maximum is reached. The blurring operator serves to destroy competing classes, thus enabling the target class to dominate the classification. The threshold corresponds to the proportion of the image contributing to the competing classification. Only when the competing classifications have been sufficiently annihilated by the blurring operator, the target class is able to come through.

The threshold value is not directly related to the class probability in the original classification. This is apparent in the examples of, on one hand, (h) **goldfish** ($p = 0.27\%$; $\text{thr} = 5 - 10\%$) and (b) **whistle** ($p = 0.06\%$; $\text{thr} = 5 - 10\%$), and on the other (e) **Border_collie** ($p = 0.01\%$; $\text{thr} = 50 - 75\%$) and (d) **sombrero** ($p = 0.01\%$; $\text{thr} = 75 - 95\%$) – different images with class probabilities of the order of 0.1% or less, but with thresholds ranging from 0 to 95%. Instead, as expected, the threshold value appears to be directly linked to the convexity of the DNN at the relevant point.

To verify this, we looked at direct comparisons between linear and nonlinear saliency maps in more detail. The comparison between the linear and nonlinear vanilla gradient saliency maps for two tentatively nonlinear examples (b) and (e) from Figures 1 and 2, is shown in Figure 3.

The image from example (e) is a fairly complex image containing two people and a dog in the front left of the image, another person and a dog to the back right, and several more people, parasols, cars and marquees in the background. The Xception DNN identified the image as being 99.98% **English_springer**, 0.01% **Border_collie**.

All the saliency maps correctly identified the two dogs in the image as being crucial to the classification. Linear saliency maps for either the **English_springer** or **Border_collie** classes select both dogs more or less equally at the masking ratio of 95%. According to the linear saliency maps, both dogs are roughly equally responsible for both classes, and the DNN is pretty clear that they are both English Springer Spaniels, or, significantly less likely, Border Collies.

Nonlinear saliency maps, however, show something quite different. While largely agreeing with linear saliency maps at low masking ratios, at 95% masking ratio they clearly show that the dog in the front left of the picture is identified as an English Springer Spaniel, while the dog in the back right is identified as a Border Collie. As it happens, the DNN is probably wrong on the latter point – to the untrained eye of the authors, the dog in the back right also appears to be an English Springer Spaniel. Whether the DNN is right or wrong is, however, in this particular context immaterial. The relevant point is that this classification can't be correctly interpreted without accounting for the local nonlinearity of the DNN through nonlinear saliency maps.

A similar pattern is seen in example (b). The image shows a hummingbird sat on a bird feeder. In this case, both linear and nonlinear saliency maps correctly identify the hummingbird, blurring out the bird feeder. On the second category, **whistle**, they, however, differ. The linear saliency map selects both the bird and the feeder as con-

tributing to **whistle**, while the nonlinear saliency map selects only the feeder, masking the bird.

We can therefore conclude that nonlinear saliency maps do generally provide a different interpretation from that of linear saliency maps. In our randomly chosen examples, the interpretation given by nonlinear saliency maps appears to be more sensible than that given by linear saliency maps.

Finally, we note that the optimization algorithm we used is a simple extension of gradient descent, which is a relatively unsophisticated optimizer. Different core optimization methods generally generate somewhat different trajectories. The direction of image enhancement is, however, only one part of the optimization procedure, with the other being the filtering operation. As noted above, the critical value governing the speed of convergence is the threshold in the proportion of image that is blurred, which is only weakly linked to the direction of image enhancement. Changing the optimizer therefore generally does not significantly affect the critical filtering threshold after which the maximum is reached, and it does not have a strong effect on the speed of convergence.

How important are nonlinear saliency maps for interpretability of machine learning models? That is, of course, at this stage still unclear. It is clear that fully accounting for nonlinearities in the model is of considerable importance for the task of interpretability. The class of methods that we describe is a step forward in this direction. It is able to do the same job as linear saliency maps where linear saliency maps are correct, and it provides an important, distinct change of direction where they are not.

The examples we presented come from computer vision and image classification, which is admittedly a task that often crosses the boundary into being to some extent subjective. Assigning complex images to simple, single classes is always somewhat arbitrary. The example image of two dogs could equally validly be classified as depicting people, a fair, or a lawn.

The true test of the methods presented here will come in real life applications which fall outside the scope of this study, such as medical image classification, material defect analysis and others.

Methods

All calculations were done in Tensorflow 1.18 (<http://www.tensorflow.org>) and NumPy (<https://numpy.org>) running on Python 3.6 (<http://python.org>) .

We used pretrained tensorflow models (<https://tfhub.dev/>) corresponding to the well documented DNNs such as VGG16, ResNet15, MobileNet, InceptionV3 and Xception.

Each image was classified by the DNN models. The nonlinear saliency maps corresponding to any class was calculated by:

1. selecting the appropriate column for the desired class from the Jacobian of the output vector with respect to the input tensor
2. for a given threshold, updating the image according to the logic

```

next_img = np.clip(
    blur_below_level(prev_img, sensitivity_norm, L) +
    littleH * sensitivity, 0., 1.)

```

3. obtaining the saliency map using the related logic

```

next_map = np.clip(
    mask_below_level(prev_img, sensitivity_norm, L) +
    littleH * sensitivity, 0., 1.)

```

Here, **sensitivity** refers to the relevant column of the Jacobian, **sensitivity_norm** to its pixel-wise norm (Euclidean norm over the RGB channel direction), **blur_below_level** performs the low pass filtering on pixels where **sensitivity_norm** falls below the selected level **L**, **mask_below_level** covers the same pixels in a background colour (in all our examples, white) and **littleH** is a small number controlling the size of the optimization step. The clipping operation ensures that all RGB values remain within the RGB range.

The level parameter threshold **L**, and the distance parameter **littleH** control the step size in the optimization process.

The authors consent for publication of identifying information/images in an online open-access publication.

References

References

- [1] Dastin, J (2018) Amazon scraps secret AI recruiting tool that showed bias against women. Reuters, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- [2] Omowole, A (2021) Research shows AI is often biased. Here's how to make algorithms work for all of us. World Economic Forum 2021, <https://www.weforum.org/agenda/2021/07/ai-machine-learning-bias-discrimination/>
- [3] Simonyan, K, Vedaldi, A and Zisserman A. (2013) Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv preprint arXiv:1312.6034
- [4] Borji, A (2019) Saliency Prediction in the Deep Learning Era: Successes and Limitations. IEEE Transactions on Software Engineering PP(99)
- [5] Mudhenk, N, Chen, BY, Friedland, G (2020) ICLR 2020, <https://arxiv.org/pdf/1911.11293.pdf>

- [6] Saporta, A, Gui, X, Agrawal, A, Pareek, A, Truong, SQH, Nguyen, CDT, Ngo, VD, Seekins, J, Blankenberg, FG, Ng, AY, Lungren, MP, Rajpurkar, P (2021) Deep learning saliency maps do not accurately highlight diagnostically relevant regions for medical image interpretation. medRxiv 2021.02.28.21252634; doi: <https://doi.org/10.1101/2021.02.28.21252634>
- [7] Arun, N Gaw, N, Singh, P, Chang, K, Aggarwal, M, Chen, B, Hoebel, K, Gupta, S, Patel, J, Gidwani, M, Adebayo, J, Li, MD, Kalpathy-Cramer, J (2020) Assessing the (Un)Trustworthiness of Saliency Maps for Localizing Abnormalities in Medical Imaging. medRxiv preprint <https://doi.org/10.1101/2020.07.28.20163899>
- [8] Simonyan, K, Zissman, A. (2014) Very Deep Convolutional Networks for Large-Scale Image Recognition. ICLR 2015 <https://arxiv.org/abs/1409.1556>
- [9] He, K, Zhang, X, Sun, J (2015) Deep Residual Learning for Image Recognition. CVPR 2015 <https://arxiv.org/abs/1512.03385>
- [10] Howard, AG, Zhu, M, Chen, B, Kalanichenko, D, Wang, W, Weyand, T, Andreeto, M and Hartwig, A (2017) MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications arXiv preprint arXiv 1704.04861 . <https://arxiv.org/abs/1704.04861>
- [11] Szegedy, C, Vanhoucke, V, Ioffe, S, Shlens, J, wowjna, (2015). Rethinking the Inception Architecture for Computer Vision. CVPR 2016 <https://arxiv.org/abs/1512.00567v>
- [12] Chollet, F (2017) Xception: Deep Learning with Depthwise Separable Convolutions. CVPR 2017 <https://arxiv.org/abs/1610.02357>

Figure 1: Image modifications on eight examples, using the Xception DNN for the highest and second highest ranked categories, with masked saliency maps below. The snapshots are, from left to right, with 0%, 5%, 10%, 20%, 30%, 50%, 75% and 95% of pixels blurred / masked.

(a) 90.10% **swimming_trunks**, 8.03% **snorkel**

(b) 99.85% **hummingbird**, 0.06% **whistle**

(c) 32.85% **soccer_ball**, 18.45% **volleyball**

(d) 99.99% **cowboy_hat**, 0.01% **sombrero**

(e) 99.98% **English_springer**, 0.01% **Border_collie**

(f) 40.89% **shoe_shop**, 22.80% **turnstile**

(g) 92.29% **ice_lolly**, 6.14% **swimming_trunks**

(h) 99.70% **tench**, 0.27% **goldfish**

Figure 2: Optimization trajectories for image modifications from Figure 1. Graph titles are the names of classes being optimized, x-axes are % of pixels blurred, and y-axes are resulting class probabilities.

(a) hummingbird (top row), whistle (bottom row).

(b) English_springer (top row), Border_collie (bottom row).

Figure 3: Linear and nonlinear, 95% masked, vanilla gradient saliency maps for Figure 1(b) and (e). Columns (left to right): original image, linear map, nonlinear map.