

ANOMALOUS BEHAVIOUR IN LOSS-GRADIENT BASED INTERPRETABILITY METHODS

Vinod Subramanian

Queen Mary University of London
v.subramanian@qmul.ac.uk

Siddharth Gururani

Electronic Arts
sgururani@ea.com

Emmanouil Benetos

Queen Mary University of London
emmanouil.benetos@qmul.ac.uk

Mark Sandler

Queen Mary University of London
mark.sandler@qmul.ac.uk

ABSTRACT

Loss-gradients are used to interpret the decision making process of deep learning models. In this work, we evaluate loss-gradient based attribution methods by occluding parts of the input and comparing the performance of the occluded input to the original input. We observe that the occluded input has better performance than the original across the test dataset under certain conditions. Similar behaviour is observed in sound and image recognition tasks. We explore different loss-gradient attribution methods, occlusion levels and replacement values to explain the phenomenon of performance improvement under occlusion.

1 INTRODUCTION

Quantitative evaluation of interpretability methods usually involves the ranking of input features using an interpretability algorithm, occluding parts of the input based on the ranking and measuring the change in the output as a result. Different techniques can be found in Samek et al. (2017); Fong & Vedaldi (2017); Petsiuk et al. (2018); Kindermans et al.; Feng et al. (2018); Carter et al. (2021). Hooker et al. (2019) suggests that re-training a model from random initialization is required to have a more accurate evaluation of interpretability methods because the training and test data need to have similar distributions.

Most research focuses on the highest ranked input features and the effects of occluding them. If the interpretability algorithm is good at identifying important features for the deep learning model then occluding the highest ranked features should cause a larger decrease in performance compared to an inferior interpretability algorithm. Hooker et al. (2019) show that removing the lowest ranked features for saliency map based attribution methods causes the perform to degrade slower. Kim et al. (2019) perform experiments with the lowest ranked features of loss-gradient based methods but they do not report any results about performance improvements. Ancona et al. show an example where occluding the lowest ranked features increases the pre-softmax activation, however they do not investigate it further.

Our work evaluates loss-gradient based attribution methods, focusing on how different occlusion levels and replacement values impact the test accuracy for a given model and attribution method. We focus on the highest and lowest ranked features. Our code can be found at ¹ and the results can be summarised as follows:

1. Removing the lowest ranked inputs can cause the performance to improve over the unchanged input;
2. The sign of the gradients is important to the ranking process which seems counter-intuitive to the idea that the sign of the gradient just indicates the direction, not the importance of the input;

¹<https://github.com/VinodS7/investigate-gradients>

3. Replacing the input with different values changes the performance and appears related to the sign of the gradients.

2 LOSS-GRADIENT ATTRIBUTION METHODS

Kim et al. (2019) show that loss-gradients are perpendicular to the decision boundary. Additionally, in adversarial attack literature the input is perturbed in the direction of the loss gradient in order to change the prediction (Szegedy et al.; Goodfellow et al., 2014). Given that our eventual goal is to bridge the gap between adversarial robustness and interpretability we decided to use loss-gradient based methods over saliency based methods.

The loss-gradient is the gradient of the loss with respect to the input to the model. For a model D which has input $x \in \mathbb{R}^{m \times n}$ with ground truth label t and loss function L , the loss-gradient $g \in \mathbb{R}^{m \times n}$ can be computed as:

$$g = \frac{dL(D(x), t)}{dx} \quad (1)$$

Based on the loss-gradient computation we use three attribution methods and we compare them to random removal of inputs. The first method is the unprocessed loss-gradients (`grad_orig`) calculated in equation 1. The second method is the absolute value of the gradients (`abs_grad`) motivated by Hooker et al. (2019) who observed slight performance improvements with the absolute value of the gradients over the raw values. The third method is multiplying the gradient with the input (`grad_inp`) motivated by Shrikumar et al. (2017). The equations for the last two methods are:

$$\begin{aligned} \text{Absolute gradient (abs_grad)} &= \left| \frac{dL(D(x), t)}{dx} \right| \\ \text{Gradient} \times \text{input (grad_inp)} &= \frac{dL(D(x), t)}{dx} \times x \end{aligned}$$

Once the attributes are obtained they need to be ranked in order to determine what input features to occlude. We want the highest ranked attributes to correspond to the most important input features and the lowest ranked attributes to correspond to the least important input features. For the absolute gradient method the attributes are arranged in descending order so the highest valued attribute corresponds to the highest ranked attribute. The other two methods contain positive and negative valued attributes and depending on how the loss function is implemented the signs could be inverted. We use the loss functions implemented in pytorch² and our experiments show that the negative gradients rank highest and the positive gradients rank lowest.

3 EXPERIMENTAL SETUP

We conduct experiments on two tasks, singing voice detection (Humphrey et al., 2019) as the audio recognition task and image recognition on MNIST dataset (LeCun et al., 2010). We use the model from Schlüter & Grill (2015) for singing voice detection, the output of this model has a single head and is passed through a sigmoid function. If the output is closer to 1 then the audio contains singing voice and if the output is closer to 0 then the audio does not contain singing voice. We use the Area Under the Receiver Operating Characteristic curve (AUROC) to evaluate the model as it bypasses the need to set a threshold for classification. We train 4 versions of this model and the average AUROC across the models is 0.961. The loss function for this model is the binary cross entropy.

We simplify MNIST to a binary class problem between the "0" and "1" digits to keep it comparable to the singing voice detection task and to simplify the analysis. So all digits except the "0" and "1" digits are removed from the training and testing dataset. We use the example models from pytorch³ for the experiments. The output of the model has two heads that are passed through logsoftmax. The performance is evaluated using test accuracy. We train 5 versions of this model and the average accuracy is 99.99%. The loss function used is the negative log likelihood.

²pytorch.org

³<https://github.com/pytorch/examples/tree/master/mnist>

4 EXPERIMENTS AND RESULTS

4.1 EVALUATING DIFFERENT ATTRIBUTION METHODS

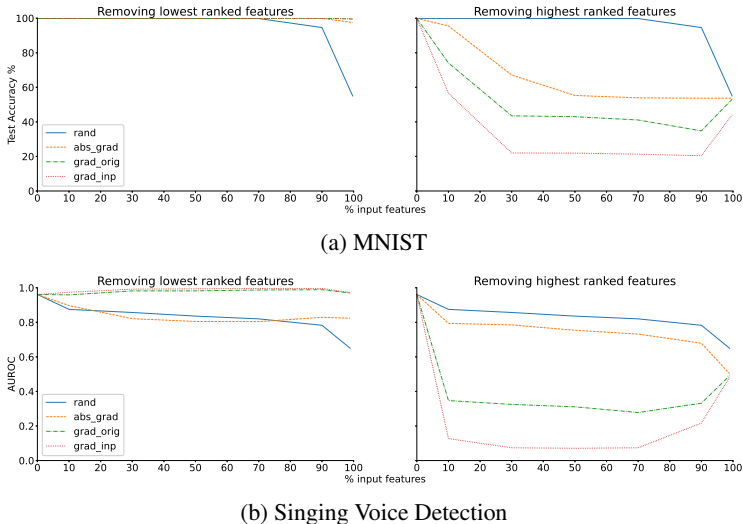


Figure 1: Performance of attribution methods at different occlusion levels

We evaluate the model at 6 different occlusion levels, 10%, 30%, 50%, 70%, 90% and a special case, for MNIST this is at 99.71% which corresponds to 782 out of 784 pixels and for singing voice it is 98.91% which corresponds to 9100 out of 9200 bins. The results are reported for the case of removing the highest ranked input features and the lowest ranked input features and the values are averaged across 5 models trained for each task. Figure 1 shows the results of this experiment.

The behavior of the model after absolute gradient occlusion is as expected. Removing the highest ranked gradients causes the performance to degrade. Removing the lowest ranked features causes negligible performance change in the MNIST dataset until 99.71% pixels are occluded, for singing voice detection the performance appears to decrease and then increase while overall performing much worse than without occlusion. The grad_orig method performs better than the abs_grad, where we observe that in both MNIST and singing voice detection the performance improves over normal evaluation for certain occlusion levels when removing the least important gradients. Finally, the grad_inp degrades the performance more than the grad_orig while removing the highest ranked features. For grad_orig and grad_inp the performance is worse than random chance while removing the highest ranked features which suggests that the occluded input consistently fools the classifier into predicting the wrong label, the increase in AUROC for singing voice detection indicates that the model is moving closer to a random classifier because informative features are occluded.

We visualize the scenario where 99.71% of the lowest ranked features of the input is occluded on the MNIST dataset in Figure 2. We observe a clear pattern where the central white pixels remain for the number 1 and the off centre white pixels remain for the number 0. Irrespective of the attribution method the two pixels remaining are the same. These examples suggest that there is an over emphasis on a small cluster of pixels to make a prediction.

4.2 EVALUATING DIFFERENT REPLACEMENT VALUES

In the previous experiment the occluded features in the input were replaced by the average value of the dataset. In this experiment we change the values that we are occluding the input with. We evaluate the performance at different occlusion levels for the grad_orig method by replacing the features with the input minimum, input maximum and dataset mean. The results are shown in figure 4. Replacing the occluded parts with the input minimum and mean has roughly similar behaviour for both tasks. Replacing by the input maximum inverts the behaviour of the model towards occluding by the highest and lowest ranked features.

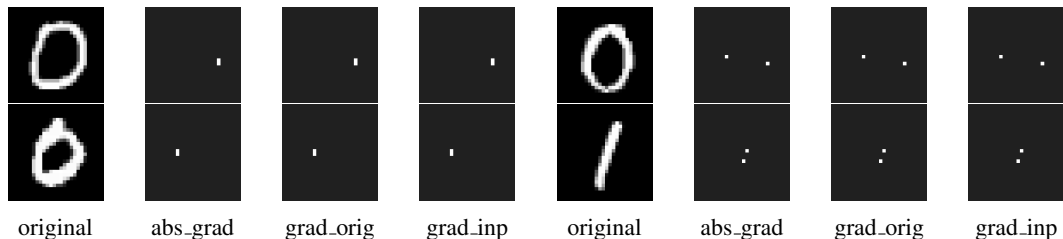


Figure 2: Comparison when 782 highest ranked pixel values are occluded by the dataset mean.

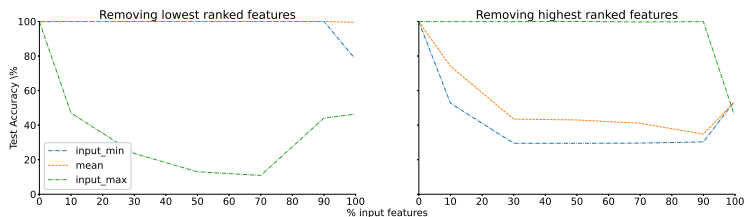


Figure 3: MNIST

Figure 4: Performance of different replacement values at different occlusion levels

In the two pixel scenario for MNIST we observe that replacing by the input minimum instead of the mean the accuracy drops from 99.5% to 78.7% which suggests that the replacement values are important. In figure 5 we compare the behaviour of replacing the highest ranked and lowest ranked features with the input maximum and the mean. In image 5b the lowest ranked features are replaced by the mean the output is predicted correctly as "1", we could infer that those central white pixels are important for prediction. Replacing by the input maximum in image 5c changes the prediction to "0" even though the same central white pixels are present. In image 5d we replace the highest ranked pixels with the mean and the prediction is "0" however replacing with the input maximum in image 5e is correctly predicted as "1". The similarity between image 5b and 5e is that the edge between the white and dark pixels is preserved.

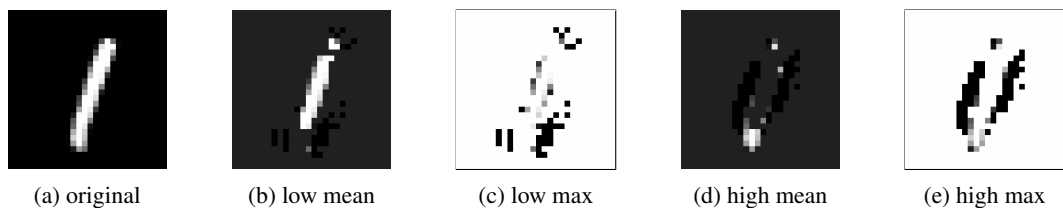


Figure 5: Replacing with mean and input maximum (max) of 90% of lowest ranked (low) and highest ranked (high) features

5 CONCLUSION

In this work we evaluate loss-gradient based attribution methods. We observe that the sign of the attribution method is important in the ranking process and the replacement value determines how the removed input features affect model placement. Our results suggest that the most important factor is the preservation of edges of the object. Future work will continue to explore the relationship of the replacement values with the occluded input features and also try to explain why the gradient \times input method outperforms the raw gradients as an attribution method.

ACKNOWLEDGEMENTS

This work has received funding from the EU’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 765068.

REFERENCES

- Marco Ancona, Enea Ceolini, Cengiz Öztireli, and Markus Gross. Towards better understanding of gradient-based attribution methods for deep neural networks. In *6th International Conference on Learning Representations, ICLR, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net. URL <https://openreview.net/forum?id=Sy21R9JAW>.
- Brandon Carter, Siddhartha Jain, Jonas Mueller, and David Gifford. Overinterpretation reveals image classification model pathologies, 2021. URL <https://openreview.net/forum?id=cP2fJWhYZe0>.
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. Pathologies of neural models make interpretations difficult. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3719–3728, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1407. URL <https://www.aclweb.org/anthology/D18-1407>.
- R. C. Fong and A. Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *IEEE International Conference on Computer Vision (ICCV)*, pp. 3449–3457, 2017. doi: 10.1109/ICCV.2017.371.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A benchmark for interpretability methods in deep neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/fe4b8556000d0f0cae99daa5c5c5a410-Paper.pdf>.
- E. J. Humphrey, S. Reddy, P. Seetharaman, A. Kumar, R. M. Bittner, A. Demetriou, S. Gulati, A. Jansson, T. Jehan, B. Lehner, A. Krupse, and L. Yang. An introduction to signal processing for singing-voice analysis: High notes in the effort to automate the understanding of vocals in music. *IEEE Signal Processing Magazine*, 36(1):82–94, 2019. doi: 10.1109/MSP.2018.2875133.
- Beomsu Kim, Junghoon Seo, and Taegyun Jeon. Bridging adversarial robustness and gradient interpretability. *Safe Machine Learning workshop at ICLR*, 2019.
- Pieter-Jan Kindermans, Kristof T. Schütt, Maximilian Alber, Klaus-Robert Müller, Dumitru Erhan, Been Kim, and Sven Dähne. Learning how to explain neural networks: Patternnet and patternattribution. In *6th International Conference on Learning Representations, ICLR, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net. URL <https://openreview.net/forum?id=Hkn7CBaTW>.
- Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2010.
- Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. In *British Machine Vision Conference (BMVC)*, 2018. URL <http://bmvc2018.org/contents/papers/1064.pdf>.
- W. Samek, A. Binder, G. Montavon, S. Lapuschkin, and K. Müller. Evaluating the visualization of what a deep neural network has learned. *IEEE Transactions on Neural Networks and Learning Systems*, 28(11):2660–2673, 2017. doi: 10.1109/TNNLS.2016.2599820.

Jan Schlüter and Thomas Grill. Exploring data augmentation for improved singing voice detection with neural networks. In *International Society of Music Information retrieval, ISMIR*, pp. 121–126, 2015.

Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3145–3153. PMLR, 06–11 Aug 2017. URL <http://proceedings.mlr.press/v70/shrikumar17a.html>.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*. URL <http://arxiv.org/abs/1312.6199>.

A APPENDIX

A.1 MNIST SIGMOID MODEL AND NORMAL MNIST

To bridge the gap between the logsoftmax MNIST model and the sigmoid singing voice detection model we create a sigmoid MNIST model with a single output where the output 1 is assigned to the label "1" and the output 0 is assigned to the label "0". We set a threshold of 0.5 for classification. We also redo the experiments for the 10-class MNIST dataset to show that this is not a phenomenon that is unique to binary classification problem. The average accuracy of the 10 class logsoftmax model is 99.03%, of the 2 class logsoftmax model is 99.95% and of the sigmoid model is 99.95%. The results are shown in table 1, in all these results the occluded inputs are replaced with the dataset mean.

Table 1: Performance for different occlusion levels. For the last 3 columns, the value on the left corresponds to the accuracy when removing the lowest ranked features and the value on the right is the accuracy when removing highest ranked features

(a) MNIST LogSoftmax model

Occlusion %	random	abs_grad	grad_orig	grad_inp
10	0.999	0.999 / 0.955	1.0 / 0.741	1.0 / 0.566
30	0.999	0.999 / 0.671	1.0 / 0.435	1.0 / 0.220
50	0.998	0.999 / 0.553	1.0 / 0.430	1.0 / 0.219
70	0.998	0.999 / 0.539	1.0 / 0.411	1.0 / 0.213
90	0.946	0.999 / 0.537	1.0 / 0.348	1.0 / 0.204
99.71	0.550	0.975 / 0.537	0.995 / 0.531	0.999 / 0.441

(b) MNIST sigmoid model

Occlusion %	random	abs_grad	grad_orig	grad_inp
10	0.999	0.999 / 0.959	1.0 / 0.710	1.0 / 0.567
30	0.999	0.999 / 0.665	1.0 / 0.393	1.0 / 0.276
50	0.999	0.999 / 0.556	1.0 / 0.390	1.0 / 0.274
70	0.998	0.999 / 0.539	1.0 / 0.384	1.0 / 0.275
90	0.970	0.999 / 0.537	0.999 / 0.350	0.999 / 0.277
99.71	0.571	0.982 / 0.537	0.999 / 0.517	1.0 / 0.397

(c) MNIST LogSoftmax 10

Occlusion %	random	abs_grad	grad_orig	grad_inp
10	0.984	0.990 / 0.693	0.996 / 0.585	0.995 / 0.567
30	0.928	0.987 / 0.368	0.985 / 0.459	0.994 / 0.276
50	0.730	0.979 / 0.252	0.978 / 0.416	0.994 / 0.274
70	0.427	0.952 / 0.188	0.965 / 0.362	0.993 / 0.275
90	0.205	0.772 / 0.140	0.892 / 0.224	0.984 / 0.277
99.71	0.119	0.150 / 0.115	0.220 / 0.114	0.265 / 0.397