

UNDERSTANDING AND MANIPULATING NEURAL NET FEATURES USING SPARSE OBLIQUE CLASSIFICATION TREES

Suryabhan Singh Hada Miguel Á. Carreira-Perpiñán Arman Zharmagambetov

Dept. of Computer Science and Engineering, University of California, Merced
{shada, mcarreira-perpinan, azharmagambetov}@ucmerced.edu

ABSTRACT

The widespread deployment of deep nets in practical applications has led to a growing desire to understand how and why such black-box methods perform prediction. Much work has focused on understanding what part of the input pattern (an image, say) is responsible for a particular class being predicted, and how the input may be manipulated to predict a different class. We focus instead on understanding what internal features computed by the neural net are responsible for a particular class. We achieve this by mimicking part of the net with a decision tree having sparse weight vectors at the nodes. We are able to learn trees that are both highly accurate and interpretable, so they can provide insights into the deep net black box. Further, we show we can easily manipulate the neural net features in order to make the net predict, or not predict, a given class, thus showing that it is possible to carry out adversarial attacks at the level of the features. We demonstrate this robustly in MNIST and ImageNet with LeNet5 and VGG networks.

Index Terms— interpretability, deep nets, decision trees

1. INTRODUCTION

Deep neural nets are accurate black-box models. They are highly successful in terms of predictive performance (say, classifying an input image) but remarkably difficult to understand in terms of how exactly they come up with a prediction for an input. Both of these issues have been known to researchers and practitioners for many years, but it is in the 2010s that deep learning has achieved a wild, unexpected success that has attracted widespread attention beyond computer science. Thus making it urgent to understand the behavior of these models in explanatory terms.

Much work in this regard seeks to understand what a specific neuron in a deep net does. This includes work on finding input patterns that invert the activation of a neuron [1, 2, 3] or maximally activate it [4, 5, 6, 7]; and work that finds input patterns, or parts of them (such as image regions) that have an important effect on the output class, essentially a sensitivity analysis via gradients or other measure of saliency [5, 8, 6, 9, 10, 11].

Other work seeks to replace a deep net with a simpler, interpretable model that can then be inspected, such as a decision tree, a set of rules or a (sparse) linear model. This can be done locally around an instance [12] or globally for all instances—which is much harder since an interpretable model cannot generally approach the accuracy of the deep net. The topic of extracting sets of rules from neural nets was actively researched in the 1990s [13, 14, 15]. Two basic approaches were used: in rule extraction as search [16, 17], a specialized heuristic search over possible rules was based on the neurons’ connectivity pattern, but this assumed binary activations and did not scale beyond small nets. In rule extraction as learning,

or teacher-student approach [18, 19, 20], one trains a decision tree to mimic the neural net by using the latter’s predictions on the training set. The fundamental problem with it that traditional decision tree learning algorithms such as CART [21] or C4.5 [22] are unable to learn small yet accurate enough trees to be useful mimics of a neural net except in very small problems.

Our paper has two contributions that can improve our ability to explain and manipulate trained deep nets. Firstly, we propose decision trees as a tool to understand deep nets. As mentioned above this is by itself not a new idea. What is new is the specific, novel type of tree we use, and how we apply it to a given deep net. Traditional tree learning algorithms typically construct trees where each decision node thresholds a single input feature. Although such trees are considered among the most interpretable models, this is only true if the tree is relatively small. Unfortunately, such trees often produce too low accuracy, and are wholly inadequate for high-dimensional complex inputs such as image pixels or neural net features. We capitalize on a recently proposed *Tree Alternating Optimization (TAO)* algorithm [23, 24] which can learn far more accurate trees that remain small and very interpretable because each decision node uses a small, learnable subset of features. It has been shown to outperform existing tree algorithms such as CART [21] or C4.5 [22] by a large margin [25], and to improve forests [26, 27, 28, 29]. Counterfactual explanations can also be solved exactly for oblique trees [30].

Second, we apply the tree to an internal layer of the deep net, hence mimicking its remaining (classifier) layers, rather than attempting to mimic the entire deep net. *This allows us to study the relation between deep net features (neuron activations) and output classes (unlike the work cited in the second paragraph, which studies the relation between input features and neuron activations)*. As a subproduct, inspection of the tree allows us to construct a new kind of adversarial attacks where we manipulate the deep net features via a mask to block a specific set of neurons. This gives us surprising control on what class the deep net will output. Among other possibilities, we can make it output the same, desired class for all dataset instances; or make it never output a given class; or make it misclassify certain pairs of classes.

Next, we describe how we use trees to understand and manipulate deep net features (section 2), and demonstrate this in MNIST and ImageNet [31] with LeNet5 and VGG16 [32] deep nets (section 4).

2. SPARSE OBLIQUE TREES AS A TOOL TO OBSERVE A DEEP NEURAL NET

Our overall approach is as follows. Assume we have a trained deep net classifier $\mathbf{y} = \mathbf{f}(\mathbf{x})$, where input $\mathbf{x} \in \mathbb{R}^D$ and $\mathbf{y} \in \mathbb{R}^K$. We can write \mathbf{f} as: $\mathbf{f}(\mathbf{x}) = \mathbf{g}(\mathbf{F}(\mathbf{x}))$, where \mathbf{F} represents the features-extraction part ($\mathbf{z} = \mathbf{F}(\mathbf{x}) \in \mathbb{R}^F$), and \mathbf{g} represents the classifier part ($\mathbf{y} = \mathbf{g}(\mathbf{z})$). Then:

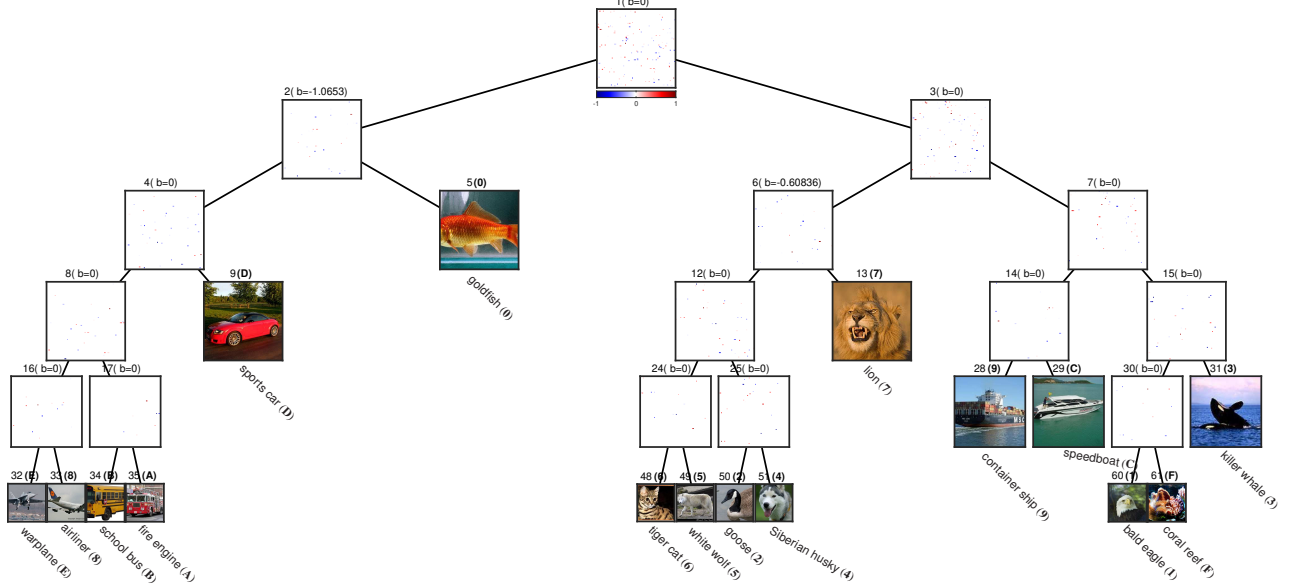


Fig. 1. Tree having one leaf per class ($\lambda = 33$). At each decision node we show its weight vector, node index and bias (always zero). At each leaf we show their index, class label, an image of from their class and class description in the format: class description (class label). We plot the weight vector, of dimension 8 192, as a 91×91 square (the last pixels are unused), with features in the original order in VGG16 (which is determined during training and arbitrary, hence the random aspect of the images), and colored according to their sign and magnitude (positive, negative and zero values are blue, red and white, respectively). You may need to zoom in the figure.

1. Train a sparse oblique tree $y = T(\mathbf{z})$ with TAO (see details in [23, 24]) on the training set $\{(\mathbf{F}(\mathbf{x}_n), y_n)\}_{n=1}^N \subset \mathbb{R}^F \times \{1, \dots, K\}$. Choose the sparsity hyperparameter $\lambda \in [0, \infty)$ such that, T have close to highest validation accuracy and is as sparse as possible.
2. Inspect the tree to find interesting patterns about the deep net.

Our goal is to achieve a tree that both mimics well the deep net and is as simple as possible.

Step 2 is purposely vague. There is probably a wealth of information in the tree regarding the features’ meaning and effect on the classification, both at the level of a specific input instance or more globally. Here, we focus on one specific pattern described next.

3. MANIPULATING THE FEATURES OF A DEEP NET TO ALTER ITS CLASSIFICATION BEHAVIOR

Our overall objective is to control the network prediction by manipulating the value of the deep net features $\mathbf{z} \in \mathbb{R}^F$. We do not alter the network weights, i.e., \mathbf{F} and \mathbf{g} remain the same. We just alter \mathbf{z} into a masked $\bar{\mathbf{z}} = \boldsymbol{\mu}(\mathbf{z}) = \boldsymbol{\mu}^\times \odot \mathbf{z} + \boldsymbol{\mu}^+$ via a *multiplicative and an additive mask* $\boldsymbol{\mu}^\times, \boldsymbol{\mu}^+ \in \mathbb{R}^F$, respectively (where “ \odot ” means elementwise multiplication).

$$\text{Original net: } \mathbf{y} = \mathbf{f}(\mathbf{x}) = \mathbf{g}(\mathbf{F}(\mathbf{x})) \quad (1)$$

$$\text{Original features: } \mathbf{z} = \mathbf{F}(\mathbf{x}) \quad (2)$$

$$\text{Masked net: } \bar{\mathbf{y}} = \bar{\mathbf{f}}(\mathbf{x}) = \mathbf{g}(\boldsymbol{\mu}(\mathbf{F}(\mathbf{x}))) \quad (3)$$

$$\text{Masked features: } \bar{\mathbf{z}} = \boldsymbol{\mu}(\mathbf{F}(\mathbf{x})) = \boldsymbol{\mu}(\mathbf{z}) \quad (4)$$

In the simplest, most intuitive version of the mask, we just need a binary multiplicative mask $\bar{\mathbf{z}} = \boldsymbol{\mu}^\times \odot \mathbf{z}$ where $\boldsymbol{\mu}^\times \in \{0, 1\}^F$. Using an additive mask and real-valued masks makes the manipulation’s effect more robust and harder to detect. We will construct a mask by inspecting the tree, specifically by observing the weight of each feature in each decision node. By selectively zeroing some features we can guarantee that any instance will follow a specific child in a given node and hence direct instances towards a target leaf.

3.1. All instances to one child

We define decision rule at a decision node i as: “if $\mathbf{w}_i^T \mathbf{z} + b_i \geq 0$ then go to right child, else go to left child”, where $\mathbf{w}_i \in \mathbb{R}^F$ is the weight vector and $b_i \in \mathbb{R}$ is the bias¹. We also describe a mask for node i , that divert all instances to one child, we call it NODE-MASK = $\{\boldsymbol{\mu}^\times, \boldsymbol{\mu}^+\}$.

NODE-MASK works as follows. Write \mathbf{w} and \mathbf{z} as $\mathbf{w} = (\mathbf{w}_0 \ \mathbf{w}_- \ \mathbf{w}_+)$ and $\mathbf{z} = (\mathbf{z}_0 \ \mathbf{z}_- \ \mathbf{z}_+)$, where $\mathbf{w}_0 = \mathbf{0}$, $\mathbf{w}_- < \mathbf{0}$ and $\mathbf{w}_+ > \mathbf{0}$ contain the zero, negative and positive weights in \mathbf{w} , and $\mathbf{z} \geq \mathbf{0}$ ² is arranged according to that. Call $\mathcal{S}_0, \mathcal{S}_-$ and \mathcal{S}_+ the corresponding sets of indices in \mathbf{w} . Then $\mathbf{w}^T \mathbf{z} + b = \mathbf{w}_-^T \mathbf{z}_- + \mathbf{w}_+^T \mathbf{z}_+$ with $\mathbf{w}_-^T \mathbf{z}_- \leq 0$ and $\mathbf{w}_+^T \mathbf{z}_+ \geq 0$. So if $\mathbf{z}_- = \mathbf{0}$ then $\mathbf{w}^T \mathbf{z} + b \geq 0$ and \mathbf{z} would go to the right child and if $\mathbf{z}_+ = \mathbf{0}$ then $\mathbf{w}^T \mathbf{z} + b < 0$ and \mathbf{z} would go to the left child. Hence, NODE-MASK defined as follows: to go left, $\boldsymbol{\mu}^\times \in \{0, 1\}^F$ is a binary vector containing ones at \mathcal{S}_- , zeros at \mathcal{S}_+ and * (meaning any value) at \mathcal{S}_0 ; and $\boldsymbol{\mu}^+ \geq \mathbf{0}$ is a vector containing small positive values at \mathcal{S}_- and zero elsewhere. To go right, exchange “-” and “+” in the procedure.

3.2. Masks

We now show how to construct masks that effect a certain class outcome. For each case, we state the desired goal and the corresponding mask. In the manipulations below we may use NODE-MASK repeatedly over several nodes to construct the mask (which is applied to the feature vector and hence applies globally to each node). In that case, we will only use the multiplicative mask produced by NODE-MASK at each node, and create the additive mask at the end given the final multiplicative mask.

ALL CLASS k_1 TO CLASS k_2 : let $k_1 \neq k_2 \in \{1, \dots, K\}$. For any instance originally classified as k_1 , classify it as k_2 . For any other instance, do not alter its classification. This case only works if

¹The bias (b_i) at each decision node i of the tree is zero. This holds very well in the trees we trained, specifically $|b_i| \ll \|\mathbf{w}_i\|$ at each decision node.

²We assume the deep net features are nonnegative: $\mathbf{z} = \mathbf{F}(\mathbf{x}) \geq \mathbf{0}$. This is true for ReLUs, which are used in most deep nets at present.

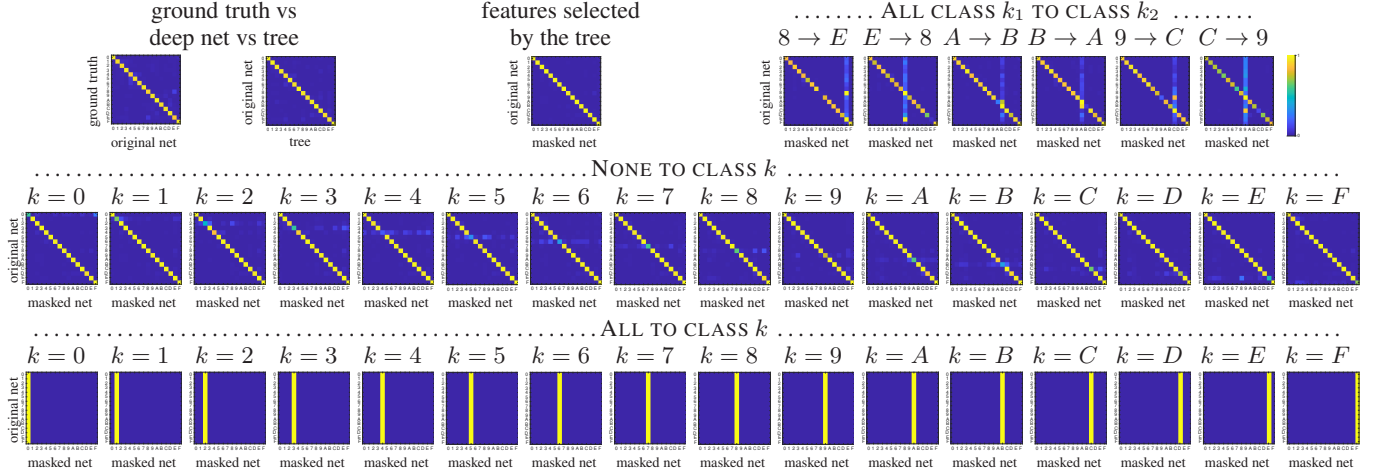


Fig. 2. Confusion matrices for VGG (test set). *Top left:* ground-truth vs deep net, and deep net vs tree. *Top middle:* deep net vs deep net with only the features selected by the tree. *Top right:* ALL CLASS k_1 TO CLASS k_2 (selected examples). *Middle:* NONE TO CLASS k . *Bottom:* ALL TO CLASS k . The confusion matrices for the training set (not shown) are very similar.

the classes k_1 and k_2 are leaf siblings (have the same parent). Class k_2 may be represented by multiple leaves since we only need to deal with one of them (the sibling of k_1). *Mask:* simply apply NODE-MASK to the parent of the leaves of k_1 and k_2 . For instance, if class k_1 is left child, then final multiplicative mask μ^\times will contain ones at \mathcal{S}_+ , zeros at \mathcal{S}_- and $*$ (meaning any value) at \mathcal{S}_0 .

NONE TO CLASS k : let $k \in \{1, \dots, K\}$. For any instance originally classified as k , classify it as any other class. For any other instance, do not alter its classification. *Mask:* simply apply NODE-MASK to the parent of each leaf of k and combine the resulting multiplicative masks as extended-AND (defined below). Finally, add the additive mask.

Strictly speaking, we can guarantee that class- k instances are classified as some other class, but note that we do not alter the classification of other instances. This is because the features that are masked out may appear in other nodes and possibly affect the path of an instance. However, with our deep nets the number of features masked out is very small and the mask works well. If the features selected in a node only appear in that node, their effect is purely local, of course.

ALL TO CLASS k : let $k \in \{1, \dots, K\}$. Classify all instances \mathbf{x} as class k . *Mask:* find the path from the root to the leaf of class k . At each node i in the path, apply NODE-MASK (to divert instances along the path) and keep the multiplicative mask only. The final multiplicative mask, elementwise, has a 0 where any of the node masks has a 0, a 1 where all node masks have no 0s but at least one 1, and $*$ elsewhere. This masks out all the “undesired” features that might divert us from the path. Equivalently, this is the logical extended-AND of all the multiplicative masks along the path (where we extend AND to mean $\text{AND}(*, 0) = 0$, $\text{AND}(*, 1) = 1$ and $\text{AND}(*, *) = *$).

This would not work if the multiplicative mask is zero at all features, but this is unlikely if the nodes have sparse weight vectors, as happens in our experiments. This also works if class k is represented as multiple leaves. We take the union of the masks over each leaf.

4. EXPERIMENTS

We have evaluated our masks thoroughly on two deep nets. 1) VGG16 [32] in a subset of 16 classes of ImageNet [31], for which we select the $F = 8192$ neurons from its last convolutional layer. 2)

LeNet5 in MNIST on 10 digit classes [33], for which we select the $F = 800$ neurons at layer conv2 as features. For both of them, we can train trees that accurately mimic the deep net classifier g . The trees give remarkable insight in the relation of deep net features to classes and allow us to construct masks that indeed work as intended in the deep net for most instances. Here, we focus on VGG16. More detailed results appear in [34].

Our VGG16 net achieves an error of 0.2% (training) and 6.79% (test). To train the tree, we use as initial tree a deep enough, complete binary tree with random parameters, and run TAO for a range of increasing λ values. From there, we pick a tree with accuracy close to that of the deep net but as sparse as possible, which we will use as mimic. This tree ($\lambda = 1$) has an error of 0% (training) and 7.90% (test); it has 39 nodes and uses just 1366 features (17% of the total 8192). We normalize the final tree so each node weight vector has norm 1. We also discuss a tree of somewhat lower accuracy but which has exactly one leaf per class (fig. 1). This tree ($\lambda = 33$) has an error of 1.79% (training) and 9.56% (test); it has 31 nodes and uses just 408 features (5% of the total 8192).

4.1. Manipulating the deep net features via masks

We derive masks using the mimic tree ($\lambda = 1$). Fig. 2 shows confusion matrices for VGG16, over test instances. As shown in the confusion matrix for deep net vs tree prediction (second matrix in the top left), both models have the same prediction for almost all instances, showing the tree mimic the network really well. This was expected as the tree has training and test errors close to those of VGG16. The interesting confusion matrix is the original network vs network with only the feature selected by the tree (top middle). Here, even after using only 17% of the features, the network has the same prediction as to the original one. It suggests that 83% of the features and hence neurons and weights of the net are practically redundant, or perhaps code for properties that are useful for only a few specific instances. This is not surprising if one notes that deep nets (at least, as presently designed) seem to be vastly overparameterized and can be significantly compressed.

Generally, the masks affect the deep net classification in the same way as the tree. This is to be expected since the tree has a very similar error and confusion matrix as the net, but it is still surprising in how well it works in most cases, like for classifying all

instances as class k mask (bottom row). This also indicates that certain deep net neurons (those critically involved in the masks) play a well-defined role in the classification. The number of features that a mask critically needs to perform its job is very small, around 200 (out of 8 192); for MNIST³ it is much smaller, around 40 (out of 800). Misclassifying class k_1 as k_2 (where k_1 must have a single leaf which is a sibling of k_2) works well too (top right), although a few instances from other classes are sometimes classified as k_2 . Not classifying any instance as class k (middle row) works also well but fails with some instances, which remain as class k . The confusion matrices for MNIST (not shown) are very similar.

Fig. 3 illustrates the mask behavior for VGG16 in an image not in the dataset. Column 2 shows how masking the features selected by the tree drastically alters in a controlled way the softmax output (hence the class prediction; row 2 and 5). We also show how the mask correlates with superpixels in the image, either manually cropped (row 3) or optimized to invert the desired deep net features (row 4).

4.2. Inspecting the sparse oblique trees

Fig. 1 shows a very interesting tree, obtained for a larger λ value so that there is exactly one leaf per class (the smallest number of leaves possibly unless we ignore classes). This tree has very few nonzero weights yet its test error is reasonable, so it probably extracts features that robustly classify most images. Also, its structure remains unchanged for a wide range of λ . *Inspecting it shows an intuitive hierarchy of classes that seem primarily related to the background or surroundings of the main object in the image.* Its leftmost subtree {warplane, airliner, school bus, fire engine, sports car} consists of man-made objects often found on roads. However, {container ship, speedboat} (man-made objects found on the sea) appears in the rightmost subtree, together with {killer whale, bald eagle, coral reef}, all of which are also typically found on the sea or on the air. Yet {goldfish} appears in a single subtree quite separate from all other classes: indeed, this fish is found on fishbowls (not the sea) in the training images. A subtree in the middle contains animals in land natural environments (forest, snow, grass, etc.): {tiger cat, white wolf, goose, Siberian husky, lion}. And so on. This is consistent with previous works that have found that, in some specific cases, the reason why a deep net classifies an object as a certain class is caused by the background or more generally by some confounding variables [12, 35]. It points to a possible vulnerability of the net, in that it may misclassify an object that happens to appear in an unusual background (say, a bald eagle standing on a road).

5. CONCLUSION

Our paper demonstrates the use of sparse oblique decision trees as a powerful “microscope” to investigate the behavior of deep nets, by learning interpretable yet accurate trees that mimic the classifier part of a deep net. Using the TAO algorithm is critical for this to succeed. The resulting tree gives insights about the relation between neurons and classes, and enables the design of simple manipulations of the neuron activations that can, for any training or test instance, change the class predicted in various, controllable ways (thus making adversarial attacks possible at the level of the deep net features).

This approach to interpreting or manipulating features applies to other types of deep nets and data, such as audio or language. It may also prove helpful in other areas where deep nets are being applied,

³For MNIST, our LeNet5 architecture achieves an error of 0.00545% (training) and 0.61% (test). We selected as mimic the tree for $\lambda = 20$, with depth 5 and only 27 nodes. It has an error of 1.28% (training) and 1.67% (test), which is very close to that of LeNet5.

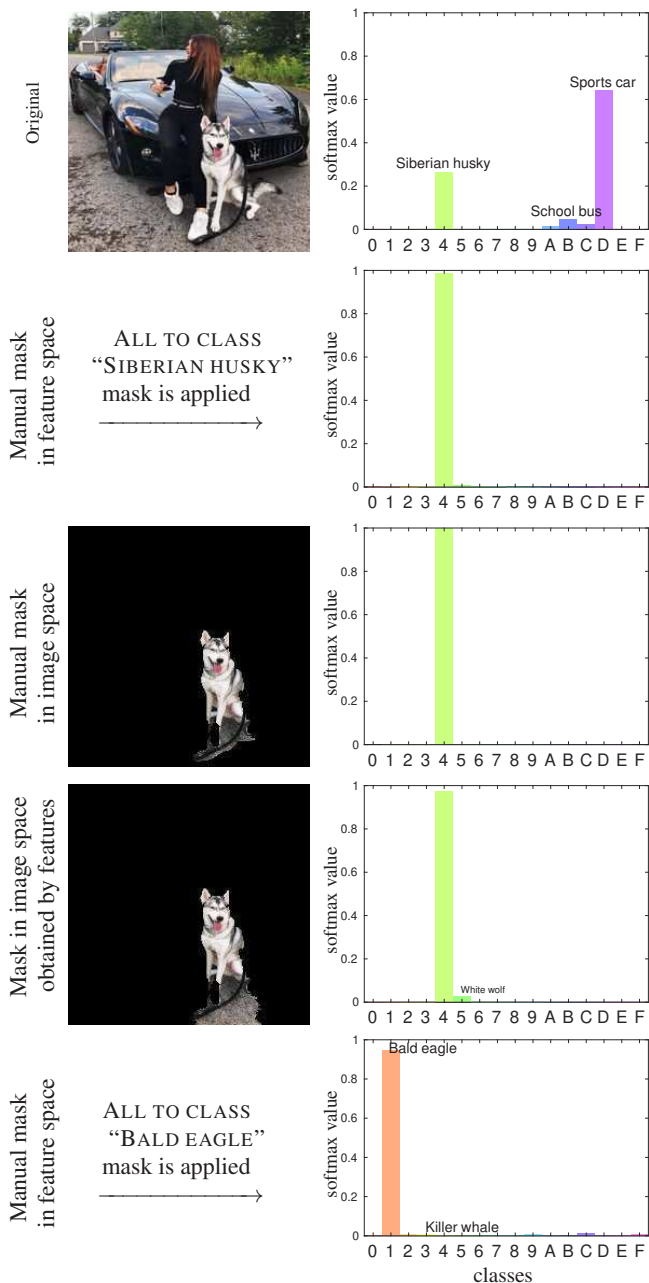


Fig. 3. Illustration of masks for a particular image in VGG16. Column 1 shows the image masks (when available). Column 2 shows the histogram of corresponding softmax values. Row 1 shows the original image. Row 2 shows a mask in feature space to classify it as “Siberian husky”. Row 3 shows a mask manually cropped in the image, whose features resemble those of row 2. Row 4 shows a mask in feature space obtained by finding the top-3 superpixels whose features most resemble those of the masked features of row 2. Row 5 shows a mask in feature space to classify the image as “bald eagle”.

such as in biology, where we may be able to relate neurons to genes or diseases, and observe the effect of “knocking out” such genes, which is essentially what our proposed masks do.

Acknowledgments. Work partially supported by NSF award IIS-2007147.

6. REFERENCES

- [1] Matthew D. Zeiler and Rob Fergus, “Visualizing and understanding convolutional networks,” in *ECCV* 2014.
- [2] Alexey Dosovitskiy and Thomas Brox, “Inverting visual representations with convolutional networks,” in *CVPR* 2016.
- [3] Aravindh Mahendran and Andrea Vedaldi, “Visualizing deep convolutional neural networks using natural pre-images,” *Int. J. Computer Vision*, Dec. 2016.
- [4] Dumitru Erhan, Yoshua Bengio, Aaron Courville, and Pascal Vincent, “Visualizing higher-layer features of a deep network,” Tech. Rep. 1341, Université de Montréal, June 2009.
- [5] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman, “Deep inside convolutional networks: Visualising image classification models and saliency maps,” in *ICLR* 2014.
- [6] Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune, “Synthesizing the preferred inputs for neurons in neural networks via deep generator networks,” in *NIPS* 2016.
- [7] Suryabhan Singh Hada and Miguel Á. Carreira-Perpiñán, “Sampling the “inverse set” of a neuron: An approach to understanding neural nets,” arXiv:1910.04857, Sept. 27 2019.
- [8] Grégoire Montavon, Sebastian Lapuschkin, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller, “Explaining non-linear classification decisions with deep Taylor decomposition,” *Pattern Recognition*, May 2016.
- [9] Ruth C. Fong and Andrea Vedaldi, “Net2Vec: Quantifying and explaining how concepts are encoded by filters in deep neural networks,” in *CVPR* 2018.
- [10] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *ICCV* 2017.
- [11] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje, “Learning important features through propagating activation differences,” in *ICML* 2017.
- [12] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, ““Why should I trust you?”: Explaining the predictions of any classifier,” in *Proc. of the 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (SIGKDD’16)*.
- [13] Robert Andrews, Joachim Diederich, and Alan B. Tickle, “Survey and critique of techniques for extracting rules from trained artificial neural networks,” *Knowledge-Based Systems*, Dec. 1995.
- [14] Keith McCormick, Dean Abbott, Meta S. Brown, Tom Khabaza, and Scott R. Mutchler, *IBM SPSS Modeler Cookbook*, Packt Publishing, 2013.
- [15] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi, “A survey of methods for explaining black box models,” *ACM Computing Surveys*, May 2018.
- [16] LiMin Fu, “Rule generation from neural networks,” *IEEE Trans. Systems, Man, and Cybernetics*, Aug. 1994.
- [17] Geoffrey G. Towell and Jude W. Shavlik, “Extracting refined rules from knowledge-based neural networks,” *Machine Learning*, Oct. 1993.
- [18] Mark Craven and Jude W. Shavlik, “Using sampling and queries to extract rules from trained neural networks,” in *ICML* 1994.
- [19] Mark Craven and Jude W. Shavlik, “Extracting tree-structured representations of trained networks,” in *NIPS* 1996.
- [20] Pedro Domingos, “Knowledge discovery via multiple models,” *Intelligent Data Analysis*, 1998.
- [21] Leo J. Breiman, Jerome H. Friedman, R. A. Olshen, and Charles J. Stone, *Classification and Regression Trees*, Wadsworth, Belmont, Calif., 1984.
- [22] J. Ross Quinlan, *C4.5: Programs for Machine Learning*, Morgan Kaufmann, 1993.
- [23] Miguel Á. Carreira-Perpiñán and Pooya Tavallali, “Alternating optimization of decision trees, with application to learning sparse oblique trees,” in *NEURIPS* 2018.
- [24] Miguel Á. Carreira-Perpiñán, “The Tree Alternating Optimization (TAO) algorithm: A new way to learn decision trees and tree-based models,” arXiv, 2021.
- [25] Arman Zharmagambetov, Suryabhan Singh Hada, Magzhan Gabidolla, and Miguel Á. Carreira-Perpiñán, “Non-greedy algorithms for decision tree optimization: An experimental comparison,” in *IJCNN* 2021.
- [26] Miguel Á. Carreira-Perpiñán and Arman Zharmagambetov, “Ensembles of bagged TAO trees consistently improve over random forests, AdaBoost and gradient boosting,” in *FODS* 2020.
- [27] Arman Zharmagambetov and Miguel Á. Carreira-Perpiñán, “Smaller, more accurate regression forests using tree alternating optimization,” in *ICML* 2020.
- [28] Arman Zharmagambetov, Magzhan Gabidolla, and Miguel Á. Carreira-Perpiñán, “Improved boosted regression forests through non-greedy tree optimization,” in *IJCNN* 2021.
- [29] Arman Zharmagambetov, Magzhan Gabidolla, and Miguel Á. Carreira-Perpiñán, “Improved multiclass AdaBoost for image classification: The role of tree optimization,” in *ICIP* 2021.
- [30] Miguel Á. Carreira-Perpiñán and Suryabhan Singh Hada, “Counterfactual explanations for oblique decision trees: Exact, efficient algorithms,” in *Proc. of the 35th AAAI Conference on Artificial Intelligence (AAAI’21)*.
- [31] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in *CVPR* 2009.
- [32] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *ICLR* 2015.
- [33] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner, “Gradient-based learning applied to document recognition,” *Proc. IEEE*, Nov. 1998.
- [34] Suryabhan Singh Hada, Miguel Á. Carreira-Perpiñán, and Arman Zharmagambetov, “Sparse oblique decision trees: A tool to understand and manipulate neural net features,” arXiv:2104.02922, Apr. 7 2021.
- [35] John R. Zech, Marcus A. Badgeley, Manway Liu, Anthony B. Costa, Joseph J. Titano, and Eric Karl Oermann, “Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study,” *PLoS Medicine*, Nov. 2018.