

# Global Explanations of Neural Networks

## Mapping the Landscape of Predictions

Mark Ibrahim\*, Melissa Louie\*, Ceena Modarres\*, John Paisley\*\*

\*Center for Machine Learning, Capital One \*\*Columbia University

### Abstract

A barrier to the wider adoption of neural networks is their lack of interpretability. While local explanation methods exist for one prediction, most global attributions still reduce neural network decisions to a single set of features. In response, we present an approach for generating global attributions called GAM, which explains the landscape of neural network predictions across subpopulations. GAM augments global explanations with the proportion of samples that each attribution best explains and specifies which samples are described by each attribution. Global explanations also have tunable granularity to detect more or fewer subpopulations. We demonstrate that GAM’s global explanations 1) yield the known feature importances of simulated data, 2) match feature weights of interpretable statistical models on real data, and 3) are intuitive to practitioners through user studies. With more transparent predictions, GAM can help ensure neural network decisions are generated for the right reasons.

### Introduction

The present decade has seen the widespread adoption of neural networks across many areas from natural language translation to visual object recognition (LeCun, Bengio, and Hinton 2015). In domains where the stakes are high however, adoption has been limited due to a lack of interpretability. The hierarchical non-linear nature of neural networks limits the ability to explain predictions. This black-box nature exposes the risk of adverse societal consequences such as bias and prejudice.

Researchers have developed several approaches to explain neural network predictions through an observation’s saliency or unique identifying factors (Kindermans et al. 2018a)(Simonyan, Vedaldi, and Zisserman 2013)(Montavon et al. 2015). Another set of techniques produce explanations using attributions, which estimate each feature’s importance to a prediction. Some of these techniques are model agnostic while others attempt to take advantage of a neural network’s architecture to generate attributions (Lundberg and Lee 2017)(Kindermans et al. 2018b) (Binder et al. 2016). However, these techniques are local, meaning the explanations are limited to a single prediction.

Global attributions can be a powerful tool for interpretability, since they explain feature importance across a population. Global explanations can aid in model diagnostics, uncover opportunities for feature engineering, and expose possible bias. Existing global techniques rely on more interpretable surrogate models such as decision trees or manipulate the input space to assess global predictive power (Frosst and Hinton 2017) (Yang, Rangarajan, and Ranka 2018) (Lakkaraju, Bach, and Leskovec 2016). While these techniques produce an interpretable set of rules, they can fail to capture the non-linear feature interactions learned by neural networks. Submodular-Pick LIME (SP-LIME) is another technique based on summarizing local attributions to better reflect learned interactions. Although SP-LIME maximizes coverage of local explanations, it does not describe the proportion or subpopulations best explained by each global attribution (Ribeiro, Singh, and Guestrin 2016).

We propose a global attribution method called GAM (Global Attribution Mapping) that’s able to explain non-linear representations learned by neural networks across subpopulations. GAM groups similar local feature importances to form human-interpretable global attributions. Each global attribution best explains a particular subpopulation, augmenting global explanations with the proportion and specific samples they best explain. In addition, GAM provides tunable granularity to capture more or fewer subpopulations for global explanations.

### Background: Attribution Techniques

GAM builds on local attribution techniques. Here we consider three techniques that are rooted in a strong theoretical foundation, well cited, and relatively easy to implement: Locally Interpretable Model-Agnostic Explanations (LIME), Integrated Gradients, and Deep Learning Important Features (DeepLIFT) (Ribeiro, Singh, and Guestrin 2016) (Sundararajan, Taly, and Yan 2017) (Shrikumar, Greenside, and Kundaje 2017).

**Locally Interpretable Model-Agnostic Explanations (LIME)** LIME produces interpretable representations of complicated models by optimizing two metrics: interpretability and local fidelity. In practice, an input is perturbed by sampling in a local neighborhood to fit a simpler

linear model. The now explainable linear model’s weights can be used to interpret a particular model prediction.

There is a global extension of LIME called Submodule pick LIME (SP-LIME). SP-LIME reduces many local attributions into a smaller global set by selecting the least redundant set of local attributions. Formally, SP-LIME maximizes coverage defined as the total importance of the features in the selected global set.

**Integrated Gradients** Sundararajan et. al developed an approach that satisfies two proposed axioms for attributions. First, sensitivity states that for any two inputs that differ in only a single feature and have different predictions, the attribution should be non-zero. Second, implementation invariance states that two networks that have equal outputs for all possible inputs should also have the same attributions. The authors develop an approach that takes the path integral of the gradient for a neutral reference input to determine the feature dimension with the greatest activation. An attribution vector for a particular observation (local) is produced as a result.

**Deep Learning Important Features (DeepLIFT)** Much like Integrated Gradients, DeepLIFT seeks to explain the difference in output between an input and a neutral reference. An attribution score is calculated by constructing a function analogous to a partial derivative that’s used along with the chain rule to trace the change in the output layer relative to the input.

### GAM: Global Attribution Mapping

We treat each local attribution as a ranking of features. In contrast to domains such as speech or image recognition where individual inputs are not human-interpretable, here we focus on situations where features have well defined semantics (Datta, Sen, and Zick 2016). We then group similar attributions to form global human-interpretable explanations across subpopulations.

### Attributions as Rankings

Attributions contain more than the content encoded by treating each as a vector in  $\mathbb{R}^n$ : an attribution ranks and weighs each feature’s association with a particular prediction. Each element in an attribution vector answers the question: “how important is this feature for a particular prediction?” Furthermore, attributions are conjoined rankings, because every feature appears in every attribution along the same dimension. Therefore to incorporate both the rank and weight encoded in an attribution, we treat attributions as *weighted conjoined rankings*. With this treatment, we summarize the content of many local attributions by comparing their similarities to form global attributions.

### Rank Distances

We represent each attribution as a rank vector  $\sigma$ . The un-weighted rank of feature  $i$  is  $\sigma(i)$  and the weighted rank is  $\sigma_w(i)$ .

First, to ensure distances appropriately reflect similarity among vectors—as opposed to anomalies in the scale of the original inputs—we normalize each local attribution vector  $\sigma_w$  by

$$|\sigma_w| \circ \frac{1}{\sum_i |\sigma_w(i)|},$$

where  $\sigma_w$  is the weighted attribution vector and  $\circ$  is the Hadamard product.

By transforming the attribution into normalized percentages, we consider only feature importance, not whether a feature contributes positively or negatively to a particular target. This enables us to measure similarity among normalized attributions as weighted rankings of feature importance.

Next we consider two options for comparing attributions: Kendall’s Tau and Spearman’s Rho squared rank distances. The first option for comparing attributions is weighted Kendall’s Tau rank distance, an extension of Kendall’s Tau rank correlation as defined by Lee and Yu 2010 (Lee and Yu 2010).

For a second attribution vector  $\pi$ , we can define the weighted Kendall’s Tau distance between  $\pi$  and  $\sigma$  as

$$\sum_{i < j} w_i w_j I\{(\pi(i) - \pi(j))(\sigma(i) - \sigma(j)) < 0\}, \quad (1)$$

where  $w_i = \pi_w(i)\sigma_w(i)$  and  $I$  is the indicator function.

A feature’s weight,  $w_i$ , in this context is the product of the weights in each ranking. Two rankings containing all elements in the same order (possibly different weights) have zero distance to imply the same feature importance. Two rankings with at least one element in a different order have a distance proportional to the out-of-order elements’ weights. When a feature is ranked above or below in one ranking compared to the other, it is penalized by the product of its weights in the rankings. The total distance between two rankings is then the sum of the penalty for feature appearing in a different order. The more features that appear in a different order, the greater the distance between the rankings. Furthermore, the weighted Kendall’s Tau distance described satisfies symmetry, positive definiteness, and subadditivity (see Section 2 of (Lee and Yu 2012)) meaning it defines a proper distance metric that we can use for clustering rankings.

A second option for comparing attributions is weighted Spearman’s Rho squared rank distance, as defined by Shieh et al. (Shieh, Bai, and Tsai 2000),

$$\sum_{i=1}^k \pi_w(i)\sigma_w(i)(\pi(i) - \sigma(i))^2. \quad (2)$$

Weighted Spearman’s Rho squared rank distance as defined compares attributions as weighted conjoined rankings much in the same way as weighted Kendall’s Tau rank distance, but with a runtime of  $\mathcal{O}(n \log n)$ . Weighted Spearman’s Rho squared distance also defines a distance metric (Lee and Yu 2012).

Kendall’s Tau rank distance does have computational and mathematical limitations. First, Kendall’s Tau has a

quadratic runtime, which can be computationally expensive as the number of feature and/or samples grows (Knight 1966). Second, Kendall’s Tau tends to produce smaller values compared to Spearman’s Rho, making differences between attributions less pronounced (Fredricks and Nelsen 2007).

Therefore, in situations where the number of features and/or samples is large, Spearman’s Rho squared is a more effective alternative from a computational perspective and due to its ability to make differences among rankings more pronounced. Despite the advantages of Spearman’s Rho squared rank distance, Kendall’s Tau remains an appropriate choice when computational efficiency is less important due to its desirable statistical properties. Kendall’s Tau as a correlation metric has faster asymptotic convergence with respect to bias and variance (Xu et al. 2013). Having defined metrics with which to compare attributions, we now turn our attention to identifying similar attributions.

### Grouping Similar Local Attributions

Clustering algorithms can detect global patterns in a dataset by grouping similar data points into a cluster. We apply this idea in the context of local attribution vectors by transforming many local attributions into a few global attributions. We use weighted Kendall’s Tau or Spearman’s Rho squared rank distances to measure similarity among attributions. We adapt K-medoids to use rank distance rather than Manhattan distance to group similar attributions (Park and Jun 2009). We then use the resulting medoids to succinctly summarize global attribution patterns.

The clustering algorithm identifies the  $K$  local attributions (medoids) minimizing the pairwise dissimilarity within a cluster relative to the medoid. We then use the medoid to summarize the pattern detected for each cluster. The algorithm works by initializing  $K$  local attributions (medoids)  $m_1, \dots, m_K$  selected randomly from the full set of local attributions. The local attributions are reassigned until convergence or for a sufficiently large number of iterations, each taking  $\mathcal{O}(n(n - K)^2)$ . The choice of medoids at each iteration is updated by

1. determining cluster membership for each local attribution  $\sigma$  via the nearest medoid,
 
$$\sigma \in C_a \text{ if } \text{rankDist}(\sigma, m_a) \leq \text{rankDist}(\sigma, m_j) \forall j \in [1, K],$$
2. updating each cluster’s medoid by minimizing pairwise dissimilarity within a cluster,

$$m_a = \sigma_l = \underset{\sigma_l \in C_a}{\text{argmin}} \sum_{\sigma_j \in C_a} \text{rankDist}(\sigma_j, \sigma_l).$$

The advantage this clustering approach over other algorithms such as K-means is that it allows us to compare attributions as conjoined weighted rankings rather than simply vectors in  $\mathbb{R}^N$ . K-means, for example, treats attributions as vectors by minimizing inter-cluster variance in Euclidean space, ignoring the content encoded in the rank and weights of each feature. Our approach, on the other hand, takes advantage of both the rank and weight information in local attributions during clustering. Each cluster specifies a subpopulation best explained by its medoid.

---

### Algorithm 1: Generating Global Attributions (GAM)

---

```

Input: local attributions
Output: medoids and corresponding members
/* 1. Normalize the set of local
   attributions */
foreach local attribution do
   normalized = abs(attribution) /
   sum(abs(attribution))
end

/* 2. Compute pair-wise rank
   distance matrix */
distances = []
foreach attribution1 in normalizedAttributions do
  foreach attribution2 in normalizedAttributions
  do
    distances += rankDistance(attribution1,
    attribution2)
  end
end

/* 3. Cluster Attributions */
initialMedoids = random.choice(attributions)
for x iterations do
  foreach cluster do
    foreach attribution in cluster do
      tempMedoid = attribution;
      cost = sum(distance(attribution,
      tempMedoid));
      reassign medoid to attribution
      minimizing cost;
    end
    update cluster membership by assigning to
    closest medoid
  end
end

```

---

### Generating Global Attributions

Each of GAM’s global explanations yields the most centrally located vector of feature importances for a subpopulation. The global attribution landscape described by GAM is then the collection of global explanations for each subpopulation.

Finally, we form GAM’s global explanations across subpopulations by grouping similar normalized attributions.

We can gauge the explanatory power of each global attribution by the size of its subpopulation. Since GAM allows us to trace a global explanation to individual samples, we can also examine summary statistics for each subpopulation. In addition, by adjusting the cluster size  $K$ , we can tune the granularity of global explanations to capture more or fewer subpopulations. As a starting point,  $K$  can be chosen based the Silhouette score, which compares similarity within and across clusters (Rousseeuw 1987).

In contrast to global explanations via surrogate models,

which produce a single global explanation, GAM identifies differences in explanations among subpopulations. Furthermore, the ability to trace global explanations to individual samples, supplements global explanations of techniques such as SP-LIME with information about subpopulations.

## Validating Our Proposed Global Attribution Method

Although there has been significant research in the validation of clustering algorithms, interpreting clusters involves qualitative analysis and domain expertise. Moreover, research in interpretability inherently faces the challenge of effective and reliable validation (Doshi-Velez and Kim 2017). Attributions have no baseline truth. The task of identifying an appropriate validation methodology for a proposed approach is an open research question. In this paper, we validate the methodology in three ways.

The first validates GAM against two synthetically generated datasets with known feature importances. The second compares the global attributions generated by GAM against the feature weights of more interpretable statistical models. The third assesses the practical value of GAM in real-world settings with user studies. We also demonstrate the use of Kendall’s Tau and Spearman’s Rho rank distances as well as several local attribution techniques.

### Global Attributions Applied to Synthetic Dataset

We design a simple synthetic dataset to yield known feature importances against which to validate GAM’s global attributions. The dataset consists of a binary label, class 1 or 2, and two features, A and B. We aim to generate one group for which only feature A is predictive and another group for which only feature B is predictive.

For the group where only feature A is predictive, we sample balanced classes from uniform distributions where

$$\text{feature A} = \left\{ \begin{array}{ll} U[0, 1] & \text{if class 1} \\ U[1, 2] & \text{if class 2} \end{array} \right\}. \quad (3)$$

To ensure feature B has no predictive power for this group, we draw from a uniform distribution over  $[0, 2]$ .

Since values for feature A in  $[0, 1]$  mark class 1 and  $[1, 2]$  mark class 2 samples, feature A has full predictive power. Feature B on the other hand is sampled uniformly over the same range for both classes, yielding no predictive power.

Conversely, to produce a group where only feature B is predictive, we sample balanced classes where

$$\text{feature B} = \left\{ \begin{array}{ll} U[3, 4] & \text{if class 1} \\ U[4, 5] & \text{if class 2} \end{array} \right\}. \quad (4)$$

To ensure feature A has no predictive power for this group, we draw from a uniform distribution over  $[3, 5]$ .

We then conduct two experiments: 1. using a balanced set where feature A is predictive in 50% of cases and 2. using an unbalanced set where feature A is predictive in 75% of cases. We train a single layer feed forward neural network on a balanced dataset of 10k samples. We use a ReLU activation with a four node hidden layer and a sigmoid output. We hold

out 2000 test samples and achieve a validation accuracy of 99%.

In order to generate global attributions for the 2000 balanced test samples, we first generate local attributions using LIME. We then generate global attributions using GAM across several values of  $K$ . By computing the Silhouette Coefficient score, we find  $K = 2$  optimizes the similarity of attributions within the subpopulations. We obtain the corresponding size of the subpopulations and two explanation medoids summarizing feature importances (see Table 1).

Table 1: Global Attributions for Synthetic Data

Balanced	
Explanations	Subpopulation Size
feature A: 0.93 feature B: 0.07	953
feature A: 0.05 feature B: 0.95	1047

Table 2: Global Attributions for Unbalanced Synthetic Data

Unbalanced	
Explanations	Subpopulation Size
feature A: 0.94 feature B: 0.06	1436
feature A: 0.13 feature B: 0.87	564

For the balanced experiment, the global explanations and subpopulation proportions match expectations. The first subpopulation contains nearly half the samples and assigns a high importance weight to feature A. The second subpopulation contains the remaining half of samples and assigns a high importance weight to feature B.

For the unbalanced experiment, the global attributions contain two subpopulations with the expected proportions (see Table 2). The first assigns a high attribution to feature A and contains 72% of samples. The second assigns a high attribution to feature B and contains 28% of samples, reflecting the feature importances inherent to the synthetic dataset. See Appendix for discussion on DeepLIFT and Integrated Gradients results.

Overall, the global attributions produced by GAM match the known feature importance and proportions of both synthetic datasets.

### Global Attributions Applied to the Mushroom Dataset

Next, we apply GAM to a dataset with a larger number of categorical features: The Audubon Society’s mushroom dataset (Dheeru and Karra Taniskidou 2017). The task is to classify mushroom samples as poisonous or not poisonous using 23 categorical features such as cap-shape, habitat, odor and so on. Here we highlight the use of gradient-based neural network attribution techniques (DeepLIFT and Integrated Gradients) to generate local attributions and compare

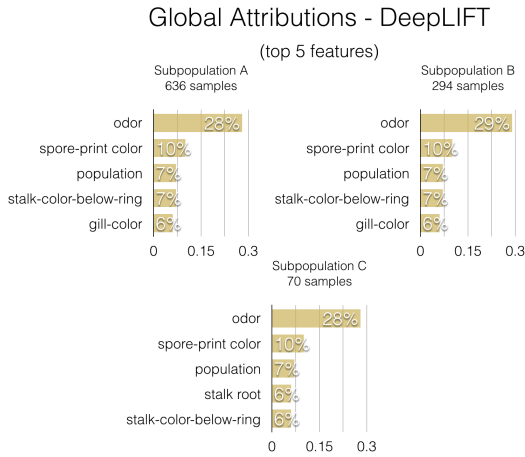


Figure 1: DeepLIFT- Five most important features for each global attribution explanation and the corresponding number of samples in each subpopulation.

the results against LIME. We also choose weighted Spearman’s Rho squared rank distance to compare attributions for computational efficiency given the larger number of features and samples.

For the classification task, we train a two hidden-layer feed-forward neural network on 4784 samples (withholding 2392 for validation) based on the architecture proposed in the Kaggle Kernel (Kim 2018). We one-hot encode the categorical features into a 128-unit input layer. The network optimizes binary cross entropy with dropout using sigmoid activation functions. We obtain a validation accuracy of approximately 99%. This allows us to produce more accurate local attributions by minimizing inaccuracies due to model error.

We apply GAM on three sets of local attributions on 1000 samples using DeepLIFT, Integrated Gradients, and LIME (Ancona et al. 2018). We identify three global explanations based on each technique.

The global attributions based on DeepLIFT produce three subpopulations with similar top features (see Figure 1). The two largest subpopulations contain 85% of samples and rank odor, spore-print-color, population, stalk-color below-ring, and gill-color as the top features with similar weights. The remaining subpopulation ranks stalk root slightly higher than gill-color and stalk-color below-ring. This implies that for 15% of predictions stalk root is a slightly more predictive feature. Examining the explanations beyond the top five features, we find other slight distinctions among subpopulations. For example, habitat appears to be a slightly more important feature for Subpopulation B than it is for Subpopulation A and C (see Appendix). Overall however, the three explanations rank and weigh the top features similarly.

The global attributions based on Integrated Gradients produce similar subpopulations. All three explanations rank odor, spore-print-color, and population as most predictive with nearly identical weights compared to those based on DeepLIFT: 27-28% versus 28-29% for odor, 10% versus

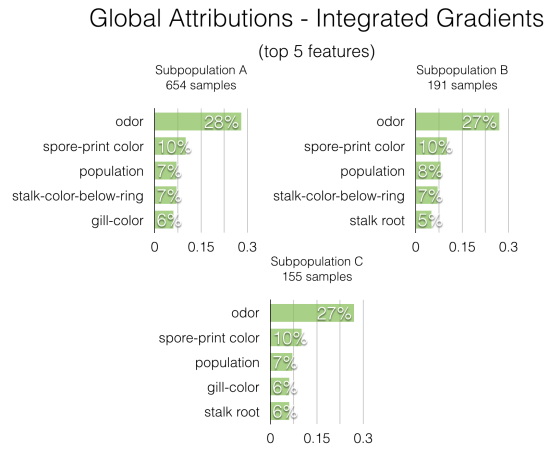


Figure 2: Integrated Gradients - Global Attributions

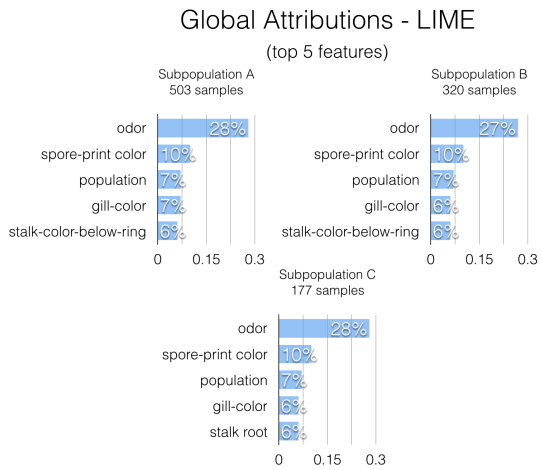


Figure 3: LIME - Global Attributions

10% for spore-print-color, and 7-8% versus 7% for population (see Figure 2). Other top features, such as gill-color and stalk root, also have matching attribution weights.

In addition, we find both gradient-based global attributions match global attributions based on LIME. Global attributions for LIME also rank odor, spore-print-color, population, and gill-color/stalk-color-below-ring as most important. The three explanations produce a very similar distribution of weights and roughly similar subpopulation proportions (see Figure 3) compared to those based on DeepLIFT and Integrated Gradients. Despite slight deviations, all three local techniques produce consistent top global feature importances.

Finally, we validate global attributions across all three local techniques against the known feature importances of a gradient boosting tree classifier. We use the XGBoost classifier implementation and assign feature importances using each feature’s F1 score (Chen and Guestrin 2016). Five out of the six most important features of XGBoost consistently appear with similar rank and weight distribu-

## XGBoost Feature Importance




Figure 4: Feature importances for XGBoost as measured by each feature’s F1 score.

tions across global attributions for all three local techniques. Odor, spore-print, population, gill-color, and stalk root consistently appear as the most important features in the explanation weights (see Figure 4). Furthermore, these top features appear with a similar distribution of weights. Odor ranks as most important with 2.8x the weight of spore-print color, which consistently ranks second. Population, gill-color, and stalk-root appear as the next most important features also with weights between 5-7%. Although global XGBoost feature importances corroborate the feature rankings and weights of our global attribution technique, we qualify these findings by noting that two equally performing models can learn a different set of feature interactions. Nevertheless, in this case the global attributions’ remarkable similarity to known global feature importances and consistency across local techniques suggest the explanations generated effectively uncover associations among features and the classification task.

We found similar results using GAM on the Iris dataset with global explanations matching the weights of a logistic classifier (see Appendix: GAM Applied to the Iris Dataset).

### User Studies

We conduct several user studies on FICO Home Equity Line Credit applications to assess the practical value of GAM in real-world settings. The dataset provides information on credit candidate applications and risk outcome characterized by 90-day payment delinquencies.

First, we simulated a scenario where a practitioner would use GAM to conduct feature selection. We selected seven predictive features and added in three additional uniformly distributed noise features. We dropped all feature names, trained a neural network to make credit decisions, produced a GAM model explanation, and presented the explanation to 55 Amazon Mechanical Turk users. We provided users with

GAM explanations in the form of a bar chart and asked users to remove the three least important features. In 94% of cases, respondents selected the features with the lowest importance corresponding to the random noise features.

Second, we polled credit modeling experts to assess whether GAM’s results are as intuitive as those of the popular SP-LIME global attribution method. We present both explanations to a group of 55 credit modeling practitioners and asked which explanation allows them to more intuitively improve the model (see Appendix). We found respondents had nearly equal preference among the three categories: GAM, SP-LIME, and no preference.

Our sample sizes are relatively small, but the results suggest that GAM has real world practical value in machine learning tasks like feature selection and produces results that are intuitive to human practitioners.

### Conclusion

In this paper, we develop GAM, a methodology for generating global attributions to supplement existing interpretability techniques for neural networks. GAM’s global explanations describe the non-linear representations learned by neural networks. GAM also provides tunable subpopulation granularity along with the ability to trace global explanations to specific samples. We demonstrate on both real and synthetic datasets that GAM illuminates global explanation patterns across learned subpopulations. We validate that global attributions produced by GAM match known feature importances and are insightful to humans through user studies. With explanations across subpopulations, neural network predictions are more transparent. A possible next step is to explore how global interpretability can help ensure fairness in AI algorithms.

### References

- Ancona, M.; Ceolini, E.; Öztireli, C.; and Gross, M. 2018. Towards better understanding of gradient-based attribution methods for deep neural networks. In *International Conference on Learning Representations*.
- Binder, A.; Montavon, G.; Bach, S.; Müller, K.; and Samek, W. 2016. Layer-wise relevance propagation for neural networks with local renormalization layers. *CoRR* abs/1604.00825.
- Chen, T., and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’16, 785–794. New York, NY, USA: ACM.
- Datta, A.; Sen, S.; and Zick, Y. 2016. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *Security and Privacy (SP), 2016 IEEE Symposium on*, 598–617. IEEE.
- Dheeru, D., and Karra Taniskidou, E. 2017. UCI machine learning repository: Mushroom data set.
- Doshi-Velez, F., and Kim, B. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

Fredricks, G. A., and Nelsen, R. B. 2007. On the relationship between spearman’s rho and kendall’s tau for pairs of continuous random variables. *Journal of Statistical Planning and Inference* 137(7):2143 – 2150.

Frosst, N., and Hinton, G. 2017. Distilling a neural network into a soft decision tree. *arXiv preprint arXiv:1711.09784*.

Kim, Y. 2018. Kaggle: Modeling on mushroom dataset.

Kindermans, P.-J.; Hooker, S.; Adebayo, J.; Schutt, K. T.; Alber, M.; Dahne, S.; Erhan, D.; and Kim, B. 2018a. The (un)reliability of saliency methods.

Kindermans, P.-J.; Schutt, K. T.; Alber, M.; Muller, K.-R.; Erhan, D.; Kim, B.; and Dahne, S. 2018b. Learning how to explain neural networks: Patternnet and patternattribution.

Knight, W. R. 1966. A computer method for calculating kendall’s tau with ungrouped data. *Journal of the American Statistical Association* 61(314):436–439.

Lakkaraju, H.; Bach, S. H.; and Leskovec, J. 2016. Interpretable decision sets: A joint framework for description and prediction. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1675–1684. ACM.

LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *nature* 521(7553):436.

Lee, P. H., and Yu, P. L. H. 2010. Distance-based tree models for ranking data. *Comput. Stat. Data Anal.* 54(6):1672–1682.

Lee, P., and Yu, P. 2012. Mixtures of weighted distance-based models for ranking data with applications in political studies. 56:2486–2500.

Lundberg, S., and Lee, S. 2017. A unified approach to interpreting model predictions. *CoRR* abs/1705.07874.

Montavon, G.; Bach, S.; Binder, A.; Samek, W.; and Müller, K. 2015. Explaining nonlinear classification decisions with deep taylor decomposition. *CoRR* abs/1512.02479.

Park, H.-S., and Jun, C.-H. 2009. A simple and fast algorithm for k-medoids clustering. *Expert Syst. Appl.* 36(2):3336–3341.

Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. ”why should I trust you?”: Explaining the predictions of any classifier. *CoRR* abs/1602.04938.

Rousseeuw, P. J. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics* 20:53–65.

Shieh, G. S.; Bai, Z.; and Tsai, W.-Y. 2000. Rank tests for independence — with a weighted contamination alternative. *Statistica Sinica* 10(2):577–593.

Shrikumar, A.; Greenside, P.; and Kundaje, A. 2017. Learning important features through propagating activation differences. *CoRR* abs/1704.02685.

Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR* abs/1312.6034.

Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. *CoRR* abs/1703.01365.

Xu, W.; Hou, Y.; Hung, Y.; and Zou, Y. 2013. A comparative analysis of spearman’s rho and kendall’s tau in normal and contaminated normal models. *Signal Processing* 93(1):261 – 276.

Yang, C.; Rangarajan, A.; and Ranka, S. 2018. Global model interpretation via recursive partitioning. *arXiv preprint arXiv:1802.04253*.

## Appendix

### Background on Local Interpretability Techniques

#### Locally Interpretable Model-Agnostic Explanations (LIME)

LIME produces interpretable representations of complicated models by optimizing two metrics: interpretability  $\Omega$  and local fidelity  $\mathcal{L}$ . Defined as,

$$\zeta(x) = \operatorname{argmin}_{g \in G} \mathcal{L}(f, g, \pi_x) + \Omega(g). \quad (5)$$

In Equation 1, the complexity or comprehensibility  $\Omega$  of an interpretable function  $g \in G$  (e.g. non-zero coefficients of a linear model) is optimized alongside the fidelity  $\mathcal{L}$  or faithfulness of  $g$  in approximating true function  $f$  in local neighborhood  $\pi_x$ .

In practice, an input point  $x$  is perturbed by random sampling in a local neighborhood and a simpler linear model is fit with the newly constructed synthetic data set. The method is model agnostic, which means it can be applied to neural networks or any other uninterpretable model. The now explainable linear model’s weights can be used to interpret a particular model prediction.

#### Integrated Gradients

Sundararajan et. al developed an approach that satisfies two proposed axioms for attributions. First, sensitivity states that for any two inputs that differ in only a single feature and have different predictions, the attribution should be non-zero. Second, implementation invariance states that two networks that have equal outputs for all possible inputs should also have the same attributions. In this paper, the authors develop an approach that takes the path integral of the gradient for a particular point  $x_i$  and the model’s inference  $F$  on the path ( $\alpha$ ) between a zero information baseline  $x'$  and the input  $x$ .

$$IG_i(x) ::= (x_i - x'_i) \int_{\alpha=0}^1 \frac{\partial f(x' + \alpha \times (x - x'))}{\partial x_i} d\alpha$$

The path integral between the baseline and the true input can be approximated with Riemman sums. An attribution vector for a particular observation (local) is produced as a result. The authors found that 20 to 300 steps can sufficiently approximate the integral within 5%. Selecting a baseline vector remains an open research question. A domain-specific heuristic (e.g. a black image in image classification) has been recommended for baseline selection.

## Deep Learning Important Features (DeepLIFT)

Much like Integrated Gradients, DeepLIFT seeks to explain the difference in output  $\Delta t$  between a 'reference' or 'baseline' input  $t^0$  and the original input  $t$  for a neuron output:  $\Delta t = t - t^0$ . For each input  $x_i$ , an attribution score is calculated  $C_{\Delta x_i \Delta t}$  that should sum up to the total change in the output  $\Delta t$ . DeepLIFT is best defined by this described summation-to-delta property:

$$\sum_{i=1}^n C_{\Delta x_i \Delta t} = \Delta t$$

Shrikumar et. al then construct a function analogous to a partial derivative and use the chain rule to trace the attribution score from the output layer to the original input. This method also requires a reference vector and the authors recommend leveraging domain knowledge to select an optimal baseline input.

## Integrated Gradients and DeepLIFT Analysis for Synthetic Dataset

While our local attributions using LIME did vary slightly depending our choice of kernel-width, we find local attributions to be extremely sensitive to our choice of a baseline input. Particularly, the number of points in each cluster varies greatly depending on our choice of baseline for Integrated Gradients. Both feature importances and subpopulations varied greatly for DeepLIFT. Unfortunately in this context there is no natural choice of baselines. To select a baseline for Integrated Gradients, we ran a grid search across our range of inputs within  $\pm 0.005$  to produce a neutral prediction. We obtain two subpopulations using this baseline for Integrated Gradients. The first contains 610 samples and assigns an attribution of 89% to feature A. The second contains the remaining 1390 and assigns a 99% attribution to feature B. We demonstrate that our global attribution technique is local-technique agnostic by generating new attributions using Integrated Gradients and DeepLIFT.

Both feature importances and subpopulations varied greatly for DeepLIFT. To select a baseline for Integrated Gradients, we ran a grid search across our range of inputs within  $\pm 0.005$  to produce a neutral prediction. We obtain two subpopulation using this baseline for Integrated Gradients. The first contains 610 samples and assigns an attribution of 89% to feature A. The second contains the remaining 1390 and assigns a 99% attribution to feature B.

## GAM Applied to the Iris Dataset

We also evaluate the proposed global attribution method on the well known Iris dataset (Anderson, 1936; Fisher, 1936). The task is to classify Iris flowers into one of three species (setosa, virginica, and versicolor) based on the length and width of each sample's petals and sepals. We train a 2-hidden layer feed forward network on 75% of the Iris data and validate on the remaining. We use a softmax activation function as the final output layer and ReLU activation functions for all other layers. The trained network achieves a validation accuracy of 90%.

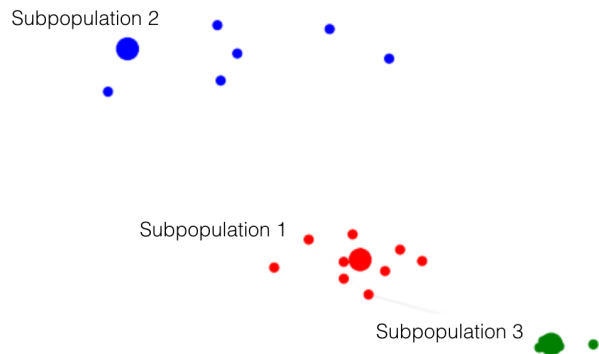


Figure 5: **Attribution Subpopulations in Kendall Tau's Rank Distance Space.** The attributions are represented as nodes in a force-directed graph where the distance between any two nodes is proportional to their Kendall Tau's rank distance. This allows us to visualize the relative rank distance among local attributions and their cluster medoids (represented as nodes with larger diameter). Subpopulation 1 and Subpopulation 2 contain attributions associated with predictions for the setosa and versicolor flower species. Subpopulation 3 contains attributions associated with virginica predictions and a single versicolor prediction.

As with the synthetic dataset, we generate local attributions using LIME. We did not obtain attributions using Integrated Gradients or DeepLIFT because that method relies on the use of a baseline vector to determine the attributions, which is not as clearly defined for the Iris dataset and defining one is beyond the scope of this paper. We then conduct a silhouette analysis to determine the optimal number of clusters for our global attribution method. We find the peak silhouette score of 0.95 is achieved with three clusters, in contrast to 0.65 for two clusters and 0.78 for four clusters. We then apply our global attribution method using three subpopulations and extract three explanations to summarize the global attribution pattern.

To validate our results, we first plot the global attribution clusters in a force-directed graph to visualize Kendall Tau's rank distance among attributions. We find the global attributions form well-isolated clusters in rank distance space (see Figure 5). Each cluster groups a set of similar attributions in rank space, which is succinctly summarized using each cluster's medoid.

We find each attribution cluster corresponds to a target class. While this structure need not arise in the context of every problem, the attributions here can be explained nicely in terms of the classes associated with each prediction (a single exception being one versicolor prediction in Subpopulation 3). We did find slight variation in the assignment of nearby attribution vectors, namely those in the region between Subpopulation 1 and Subpopulation 3, but the cluster structure appeared stable across many iterations with distinct clusters.

Each subpopulation's explanation succinctly summarizes the feature importances for each species' predictions (see



Table 3: Attribution Subpopulation Explanations

Feature	Subpop. A	Subpop. B	Subpop. C
sepal length	[0.01]	[0.15]	[0.11]
sepal width	0.61	0.31	0.01
petal length	0.03	0.37	0.40
petal width	[0.35]	[0.17]	[0.48]

Table 3). Sepal and petal width are most important for setosa predictions (Subpopulation A), whereas sepal width and petal length are most important for versicolor predictions (Subpopulation B). Petal features turn out to be most important for virginica predictions (Subpopulation C). We can then compare these features against the weights of a baseline classifier to assess whether the medoids reasonably capture the correct set of feature importances.

Table 4: Absolute Weights of Logistic Classifiers

Feature	setosa	versicolor	virginica
sepal length	[0.87]	[0.04]	[0.33]
sepal width	1.31	1.23	0.17
petal length	1.66	0.91	1.39
petal width	1.45	[0.83]	2.32

We additionally validate our approach by comparing the set of feature importances from our global attributions against the weights learned from a baseline one-vs-rest logistic regression classifier. We find the coefficients from the logistic regression classifiers correspond directly to the weights we obtain from our global attribution method. We find sepal width, petal length, and petal weights to have the largest coefficients across the three classes, similar to what we found across all three attribution medoids (see Table 4). Furthermore, we find the largest absolute coefficients across all classes (except for petal length for setosa) correspond directly to the feature importances uncovered by our global attribution method. For example, the largest weights for virginica predictions are the same petal features highlighted by the global attribution medoid for Subpopulation C.

It is worth noting that two classifiers may learn a different set of feature interactions. In this case, the feature interactions surfaced by our global attribution method correspond to the feature weights learned by a one-vs-rest logistic regression classifier. This provides evidence not only that the number of clusters and cluster patterns are meaningful, but also that the specific feature importances uncovered in each cluster effectively describe global model attribution.

### Mushroom Global Attributions

Global Attributions with explanations for all features across the three techniques.

## Global Attributions with DeepLIFT

### Global Attributions for DeepLIFT

```

[{"odor", 0.2819175505146053}, {"spore-print-color", 0.09988245611040582}, {"population", 0.0734925776767589}, {"stalk-color-below-ring", 0.06761223754619289}, {"gill-color", 0.0612280312866792}, {"stalk-root", 0.05811876114919356}, {"stalk-color-above-ring", 0.0508643463993431}, {"gill-size", 0.027929794396024}, {"habitat", 0.03614860758494721}, {"stalk-shape", 0.0364257485177965}, {"ring-type", 0.03392389850735}, {"cap-color", 0.0328608184353151}, {"stalk-surface-above-ring", 0.03010193666159343}, {"gill-spacing", 0.0275864098264906}, {"stalk-surface-below-ring", 0.02046430958955}, {"cap-shape", 0.015988254002975616}, {"bristles", 0.010228697145034588}, {"ring-number", 0.0078950959092361499}, {"cap-surface", 0.00581878031419271}, {"gill-attachment", 0.0058267854921983}, {"veil-color", 0.004587809115534257}, {"veil-type", 0.0}

[{"odor", 0.28613172969611}, {"spore-print-color", 0.0998295262133686}, {"population", 0.0712711203169187}, {"stalk-color-below-ring", 0.066736120737785}, {"gill-color", 0.0567516677321895}, {"stalk-root", 0.056442988347589194}, {"stalk-color-above-ring", 0.0420074803242949}, {"ring-type", 0.0409731046632541}, {"gill-size", 0.038933440983005}, {"habitat", 0.0382242455321979}, {"stalk-shape", 0.03821004930305221}, {"stalk-surface-above-ring", 0.03642569892102365}, {"cap-color", 0.03103603979153031}, {"stalk-surface-below-ring", 0.0277668060817938}, {"stalk-surface-above-ring", 0.02192276226490263}, {"cap-shape", 0.0188553171181378}, {"bristles", 0.013505713202368}, {"cap-surface", 0.007646534987348}, {"ring-number", 0.007467347854207311}, {"gill-attachment", 0.00283991010293986}, {"veil-color", 0.001564174739328}, {"veil-type", 0.0}

[{"odor", 0.2786058992002732}, {"spore-print-color", 0.1023203203029687}, {"population", 0.07301881514219}, {"stalk-root", 0.063424608292548}, {"stalk-color-below-ring", 0.0428222434459276}, {"stalk-color-above-ring", 0.0293487316900484}, {"ring-type", 0.027449862460376}, {"stalk-shape", 0.0303337808080846}, {"stalk-surface-above-ring", 0.0293454508305315}, {"habitat", 0.031097126952167}, {"cap-color", 0.03245814924282071}, {"gill-spacing", 0.0291079954282824}, {"stalk-surface-below-ring", 0.0213662053214087}, {"cap-shape", 0.01925537348619075}, {"bristles", 0.0126916677647947}, {"ring-number", 0.00932328528662043}, {"cap-surface", 0.004455596228907}, {"veil-color", 0.0040645104515133}, {"gill-attachment", 0.0022523412094962}, {"veil-type", 0.0}

```

## Global Attributions with Integrated Gradients

### Global Attributions for Integrated Gradients

```

[{"odor", 0.27294447026378257}, {"spore-print-color", 0.0959787420733831}, {"population", 0.075137889672198}, {"stalk-color-below-ring", 0.06916192155564631}, {"stalk-root", 0.0549135810751718}, {"gill-color", 0.0538662992304901}, {"stalk-color-above-ring", 0.04222412414039}, {"ring-type", 0.04222412414039}, {"gill-size", 0.0386927883380541}, {"stalk-shape", 0.0364503912550713}, {"habitat", 0.0366176730877991}, {"stalk-surface-above-ring", 0.0340869367854133}, {"gill-spacing", 0.0310953952928099}, {"cap-color", 0.0294890504427458}, {"stalk-surface-below-ring", 0.0227239841963747}, {"cap-shape", 0.0204776922124369}, {"bristles", 0.0126721554554588}, {"veil-color", 0.00421732844009812}, {"ring-number", 0.00617264420751216}, {"cap-surface", 0.003318631899169325}, {"gill-attachment", 0.00147484007221977}, {"veil-type", 0.0}

[{"odor", 0.2819175505146053}, {"spore-print-color", 0.09982841140582}, {"population", 0.0734925776767589}, {"stalk-color-below-ring", 0.06761223754619289}, {"gill-color", 0.0612280312866792}, {"stalk-root", 0.05811876114919356}, {"stalk-color-above-ring", 0.0508643463993431}, {"gill-size", 0.027929794396024}, {"habitat", 0.03614860758494721}, {"stalk-shape", 0.0364257485177965}, {"ring-type", 0.03392389850735}, {"cap-color", 0.0328608184353151}, {"stalk-surface-above-ring", 0.03010193666159343}, {"gill-spacing", 0.0275864098264906}, {"stalk-surface-below-ring", 0.02046430958955}, {"cap-shape", 0.015988254002975616}, {"bristles", 0.010228697145034588}, {"ring-number", 0.0078950959092361499}, {"cap-surface", 0.00581878031419271}, {"gill-attachment", 0.0058267854921983}, {"veil-color", 0.004587809115534257}, {"veil-type", 0.0}

[{"odor", 0.2741475034424628}, {"spore-print-color", 0.0988497191645933}, {"population", 0.07480031297229}, {"stalk-color-below-ring", 0.064847526079378}, {"stalk-root", 0.0628060205408285}, {"stalk-color-above-ring", 0.0576756492272763}, {"stalk-shape", 0.05670293708867125}, {"ring-type", 0.03749424830272968}, {"gill-size", 0.0374528845499222}, {"stalk-surface-above-ring", 0.0358939497939339}, {"cap-color", 0.035210390809191}, {"habitat", 0.03471077797937}, {"stalk-surface-below-ring", 0.03186246174528336}, {"gill-spacing", 0.0276222321021979}, {"stalk-surface-above-ring", 0.0213081453374881}, {"cap-shape", 0.01409150595974317}, {"bristles", 0.0098718512113444}, {"ring-number", 0.0081874692614735}, {"cap-surface", 0.007786212449237}, {"veil-color", 0.0068777869171841}, {"gill-attachment", 0.002746159139946625}, {"veil-type", 0.0}

```

## Global Attributions with LIME

### Global Attributions for LIME

```

[{"odor", 0.28326692581650925}, {"spore-print-color", 0.1018822529519018}, {"population", 0.0739378672270212}, {"gill-color", 0.046827258079383}, {"stalk-root", 0.056389866268081}, {"stalk-color-below-ring", 0.0546182361828084}, {"stalk-color-above-ring", 0.0512250389821889}, {"habitat", 0.0379912173953958}, {"ring-type", 0.0372074829081235}, {"gill-size", 0.035860412812826}, {"stalk-surface-above-ring", 0.0352500000000000}, {"cap-color", 0.03288397161241}, {"gill-spacing", 0.0296429207949157}, {"stalk-surface-below-ring", 0.02471820460221271}, {"cap-shape", 0.0179407667498254}, {"bristles", 0.0124401376454281}, {"ring-number", 0.007446791044379876}, {"cap-surface", 0.0058975014195748}, {"veil-color", 0.00325396618550825}, {"gill-attachment", 0.000879191790134611}, {"veil-type", 0.0}


[{"odor", 0.26953359473025}, {"spore-print-color", 0.10278148567176432}, {"population", 0.07289139252629312}, {"gill-color", 0.0512827165688783}, {"stalk-root", 0.0581371666295386}, {"stalk-color-below-ring", 0.0547613056516223}, {"stalk-color-above-ring", 0.04842597144768093}, {"cap-color", 0.0420125253952786}, {"ring-type", 0.039278882788783}, {"habitat", 0.03716451736941424}, {"stalk-shape", 0.037134242026065}, {"stalk-surface-above-ring", 0.036242702787815}, {"gill-spacing", 0.033971922296419}, {"gill-size", 0.02601341245829}, {"stalk-surface-below-ring", 0.0221248890796192}, {"cap-shape", 0.022263747579783}, {"bristles", 0.01199580544270594}, {"ring-number", 0.00828186181816}, {"veil-color", 0.004703873476423834}, {"cap-surface", 0.004520474209101}, {"gill-attachment", 0.0003815282530481}, {"veil-type", 0.0}

[{"odor", 0.2742462411910523}, {"spore-print-color", 0.1020715003885014}, {"population", 0.0689144363958799}, {"stalk-color-below-ring", 0.0529484137459591}, {"stalk-root", 0.0582464086336665}, {"stalk-color-above-ring", 0.04844794176234069}, {"ring-type", 0.0388470628264576}, {"gill-size", 0.0382169745732798}, {"stalk-shape", 0.0354527080850116}, {"cap-color", 0.03448420220232467}, {"habitat", 0.0333066102866410}, {"stalk-surface-above-ring", 0.03316946878821212}, {"gill-spacing", 0.02778610896597803}, {"stalk-surface-below-ring", 0.0224455703864684}, {"cap-shape", 0.02108962537056747}, {"bristles", 0.0127480113919931}, {"veil-color", 0.00813254546025937}, {"ring-number", 0.0038618013772081}, {"gill-attachment", 0.0003489976464132936}, {"veil-type", 0.0}

```

## User Studies

We asked 55 MTurk users to identify the top five features based on GAM's global explanations with two subpopulations on the FICO dataset (see figures below).



In 97% of cases respondents selected the meaningful features for subpopulation one and 92% for subpopulation two.

We also generated identical attributions for SP-LIME with B=2 and asked 55 credit modeling practitioners which explanation they preferred. In this case, the raw feature names were also displayed without a method label (see figures). Below are aggregate anonymous responses.

