

ON NETWORK SCIENCE AND MUTUAL INFORMATION FOR EXPLAINING DEEP NEURAL NETWORKS

Brian Davis^{1*}, Umang Bhatt^{1,2*}, Kartikeya Bhardwaj^{1,3*}, Radu Marculescu^{1,4}, and José M.F. Moura¹

¹Carnegie Mellon University, ²University of Cambridge, ³Arm Inc., ⁴The University of Texas at Austin

ABSTRACT

In this paper, we present a new approach to interpret deep learning models. By coupling mutual information with network science, we explore how information flows through feedforward networks. We show that efficiently approximating mutual information allows us to create an information measure that quantifies how much information flows between any two neurons of a deep learning model. To that end, we propose NIF, Neural Information Flow, a technique for codifying information flow that exposes deep learning model internals and provides feature attributions.

Index Terms— deep learning, information theory, network science, interpretability

1. INTRODUCTION

As deep learning gains popularity, there has been an influx in methods that attempt to explain how deep learning begets its predictive power. Most approaches to interpret deep neural networks are model agnostic and make a local approximation in the feature space around the datapoint to be explained [1]. However, such techniques fail to capture global *model-specific* behavior that is crucial to understand if the function learned by a deep learning model aligns well with human intuition. Moreover, current noisy approximations neglect the topological structure of the model used for prediction [2].

It is easy to forget the network structure of deep learning models, particularly feed forward models, that resemble directed acyclic graphs. However, understanding the topological structure of different models can not only help decide the architecture best suited for the task at hand but also help expose the internal interactions between neurons at inference time. While the existing interpretability techniques [3] shed light on which input features are responsible for a given prediction, prior art still fails to quantify how information flows through a deep network at the neuron-level. This prevents answering one of the most fundamental questions in deep learning: *How much information flows through a deep network from input features to each subsequent neuron?*

To address this question, we consider two types of interpretability notions: (i) model interpretability via attribution

to input features, and (ii) network architecture interpretability with respect to how information flows from neuron to neuron for a given pretrained model. We believe, addressing notion (ii) from a fundamental information theory standpoint will automatically reveal insights about the precise decision-making process followed by the model (i.e., notion (i)).

Using an information theoretic measure, we model the flow of information via *Neural Information Flow* (NIF) between neurons in consecutive layers to expose how simple deep learning models can learn complex functions of input features. We further analyze information flow between neurons from a network science [4] perspective, where each neuron in the deep network becomes a node in the network. NIF recovers an information-theoretic feature attribution, namely a rank of feature importance to a given class. Combining an information measure with the ability to propagate information through the network can help us visualize information flow and feature attributions simultaneously.

2. BACKGROUND

2.1. Network Science

Network science has gained a lot of interest for many biological and social science applications. However, to the best of our knowledge, network concepts have *not* been used to understand the inner workings of deep neural networks. To that end, several ideas from network science can be used for better understanding deep network architectures.

Betweenness Centrality: Given a network $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, betweenness centrality $B(v)$ of a node $v \in \mathcal{V}$ is a measure of how central a node is in the network. Specifically, $B(v)$ computes how many shortest paths between different pairs of nodes in the network pass through v . Mathematically, $B(v)$ is expressed as:

$$B(v) = \sum_{s \neq t \neq v} \frac{\sigma_{st}(v)}{\sigma_{st}} \quad (1)$$

where, σ_{st} is the number of shortest paths between nodes $s, t \in \mathcal{V}$, and $\sigma_{st}(v)$ are the shortest paths through v [5].

Community Structure: Communities in a network refer to groups of tightly connected nodes. Intuitively, a community

*Equal Contribution

can be defined as a group of nodes where the number of connections within this group is significantly higher than what we would expect in a randomly connected network. Mathematically, communities can be computed by maximizing a modularity function [5]:

$$\max_{\mathbf{g}=\{g_1, g_2, \dots, g_k\}} \frac{1}{2m} \sum_{ij} \left[A_{ij} - \frac{1}{\gamma} \cdot \frac{k_i k_j}{2m} \right] \delta(g_i, g_j) \quad (2)$$

where, m is the number of edges, k_i is the degree (number of connections) of node i , A_{ij} is the weight of the link between nodes i and j , and δ is Kronecker delta. The idea is to find groups of tightly connected nodes, $\mathbf{g} = \{g_1, g_2, \dots, g_k\}$, that map the nodes \mathcal{V} to k communities. The $k_i k_j / 2m$ factor represents the number of links one would expect in a randomly connected network. Finally, γ controls the resolution of communities: lower gamma will detect a higher number of smaller communities.

2.2. Interpretability

Current interpretability techniques fall into two classes. The first class consists of gradient-based methods that compute the gradient of the output with respect to the input, treating the gradient flow as a saliency map [2]. The other type leverages perturbation-based techniques to approximate a complex model using a locally additive model, thus explaining the difference between test output-input pair and some reference output-input pair. For instance, [6] proposed SHAP, a class of methods that randomly draws points from a kernel centered at the test point and fits a sparse linear model to locally approximate the decision boundary. While gradient-based techniques like [2] consider infinitesimal regions on the decision surface and take the first-order term in the Taylor expansion as the additive model, perturbation-based additive models use the finite difference between the input and a reference vector.

2.3. Information Theory

Mutual information has proven to be a valuable tool for feature selection at training time leveraging dimensionality reduction [7]; it can be represented as:

$$I(X; Y) = \mathbb{E}_{(x,y)} \left[-\log \left(\frac{\mathbb{P}(x,y)}{\mathbb{P}(x)\mathbb{P}(y)} \right) \right]$$

where $\mathbb{E}_{(x,y)}$ is the expectation over x and y , and $\mathbb{P}(x,y)$, $\mathbb{P}(x)$, $\mathbb{P}(y)$ represent their joint and marginal distributions.

More recent work has represented deep neural networks as Markovian chains to create an information bottleneck theory for deep learning [8]. However, these works do not tackle the interpretability problem directly. Previous works look to find $I(X; Y)$, the mutual information between an input vector and the output vector. In order to explain the conditional distribution of the output vector given the input vector, some

have developed an efficient variational approximation to mutual information [3]. However, this model fails to recover the per-feature mutual information, a requisite of our model to explain how information flows through all possible paths. As such, we leverage *point-wise mutual information* (PMI) estimates that would assign a real-valued quantity to each edge for an individual sample; essentially, we remove the expectation from a normal mutual information formulation and are left with the negative log of the joint over the marginals of the observed values. Using PMI also helps us recover realistic feature attributions, as shown qualitatively hereafter. To estimate the mutual information between two empirical continuous distributions (i.e., the activations at nodes), we use the mixed continuous-discrete KSG estimator [9], the EDGE estimator [10], and the MINE estimator [11].

3. PROPOSED APPROACH

Our proposed NIF approach transforms a traditional deep learning model into a representation that actually captures the information-theoretic relationship between nodes (Fig 1). Our approach extends the approach in [11] and decomposes




Fig. 1. Traditional Model to NIF Network. Color of nodes corresponds to communities. Size of nodes corresponds to betweenness centrality.

their approximation of $I(\mathcal{X}; \mathcal{Z})$ to give us $I(\mathcal{X}_i; \mathcal{Q}_k)$, where X_i is a dimension of X (specifically the i^{th} feature of the input vector), and \mathcal{Q}_k is any quantity of interest (perhaps the k^{th} neuron in a hidden layer or a class of the output vector). Assuming that mutual information is composable and entropy is non-decreasing, we can calculate the mutual information for any feature \mathcal{Q}_k by leveraging a tractable approximation from [7]:

$$\mathcal{NIF} = I(\mathcal{X}_i; \mathcal{Q}_k) = I(\mathcal{X}; \mathcal{Q}_k) - \beta \sum_{j=1}^{i-1} I(\mathcal{X}_i; \mathcal{X}_j) \quad (3)$$

where parameter β can be used to tune the removal of mutual information between features. The first term is the *relevance* of \mathcal{X} to \mathcal{Q}_k and the second term is *redundancy* in interactions between dimensions of the input. β is chosen via cross-validation and was set at 5×10^{-4} for our experiments. Note that the complexity of our estimator scales linearly by network depth, but quadratically by layer dimensionality.

3.1. Feature Attribution

NIF recovers a feature attribution, as outlined in Algorithm 1. We find all the possible paths between a feature of interest x_i and any of the outputs y_1, \dots, y_c . To find out the value of a path, we take the product of all NIF calculations along the path. We then sum over all of the possible values to find $\mathcal{A}_{i,j}$, our desired feature attribution for feature x_i and class y_j . Mathematically, the element $\mathcal{A}_{i,j}$ of our *attribution matrix* $\mathcal{A} \in \mathbb{R}^{n \times c}$ (where n is the number of features and c is the number of classes) is:

$$\mathcal{A}_{i,j} = \sum_{p \in \mathbb{P}} \prod_{l \in \mathbb{L}} \mathcal{NIF}_p(l) \quad (4)$$

where \mathbb{P} is the set of all directed paths from input x_i to class y_j in the neural information flow network, and \mathbb{L} is the set of links on each path $p \in \mathbb{P}$.

Algorithm 1 NIF Attribution for predictor f

Input: NIF model for learnt predictor $\mathcal{NIF}_{\hat{f}}$, feature index i , class of interest j

Find the set of links \mathbb{L} (where each link contains two nodes) for all possible paths \mathbb{P} between feature i and class j (that is $x_i, y_j \in p \subset \mathbb{P}$)

for path $p \in \mathbb{P}$ **do**

for link $l \in \mathbb{L}_p$ **do**

 Compute the NIF between the two links $\mathcal{NIF}_p(l)$

end for

end for

Output: Find the sum of $\mathcal{NIF}_p(l)$ across all paths \mathbb{P} multiplied over all links on the path \mathbb{L}_p

3.2. Extension to Convolution Neural Networks

Moving beyond feed forward networks, we have only addressed global explanations (e.g., which neuron affects which other neuron) because we estimate mutual information that, by definition, averages across all samples. However, for feature attribution in image classification problems via CNNs, we need instance-wise explanations such as saliency maps [12]. Therefore, we formulate a point-wise mutual information (PMI)-based technique to generate saliency map-based explanation for individual samples. Recall that PMI can be obtained by removing the mean over all samples (see Section 2.3) to preserve sample-specific information. Since the activations at channels of a CNN are continuous, we need a continuous estimate of the PMI. Towards this end, KSG estimator [9] estimates mutual information as a mean (over samples) of the log of the Radon-Nikodym derivative:

$$\hat{I}(X; Q) = \frac{1}{N} \sum_{i=1}^N -\log \left(\frac{d\mathbb{P}(x_i, q_i)}{d\mathbb{P}(x_i)\mathbb{P}(q_i)} \right) \quad (5)$$

Hence, PMI can be directly obtained by the negative log of the Radon-Nikodym derivative above. Hence, from [9], we can estimate PMI between two quantities of interest Q_j and Q_k (e.g., between channels Q_j and Q_k) for sample i is:

$$\mathcal{NIF}_{Q_j, Q_k} = \psi(k) + \log(N) - \alpha \quad (6)$$

where $\alpha = \log(n_{q_j} + 1) + \log(n_{q_k} + 1)$, $\psi(\cdot)$ is a digamma function, k is a parameter to specify the number of nearest neighbors, N is the total number of samples, n_{q_j} ($n_{q_j, i}$) is the number of samples of Q_j within some distance ρ (see [9]). Next, we describe how we obtain saliency maps for CNNs.

Saliency Maps for CNNs: The saliency maps show which part of the input image are important for the final prediction. Hence, the NIF links for CNNs must map the input image to the channels. Therefore, because the first convolutional layer maps images to multiple channels, we compute NIF from individual pixels to output channels for the first layer. This stage incurs most of the computation. Note that, in CNNs, the interpretation is w.r.t. output *channels* (and not individual neurons – that would be merely a pixel of an output channel¹). Therefore, to get Q_j, Q_k in the rest of the network, we use averaged outputs of channels in each convolutional layer. Consequently, creating NIF links for a CNN becomes analogous to that in a feedforward neural network (since each channel now represents a single averaged output), except for the very first layer, that maps input pixels to the first channels. Finally, Equation (4) is used to traverse the NIF links throughout the CNN and to obtain a saliency map. Specifically, the links map the output to the input pixels and the resulting attribution matrix is used as a saliency map.

4. RESULTS

In order to test the fidelity of NIF, we run a few experiments that validate our proposed technique. We run all experiments on UCI datasets, namely Iris and Banknote authentication, both of which provide us with a small enough feature space to interpret and visualize the network [14].

4.1. Network Visualization

We start by visualizing NIF for a one layer perceptron trained on the Banknote dataset with ReLU activations and optimized via ADAM [15], and a two layered network with ReLU activation on the Iris dataset.

In Figure 2, we show the NIF network created using Equation (6). For both NIF models, we normalize the information flow per layer to ease visualization of the edges. The thickness of an edge denotes how much information is flowing between any two nodes: the thicker the connection, the more information travelling from one node to the next. The size of the node denotes its centrality: the bigger the node, the more

¹In the context of image classification problems, we always talk about what a channel represents in the final convolutional layer of a CNN, e.g., for a CNN that classifies between cats and dogs, some channels activate for dog-faces, others activate for cat-stripes, etc. [13]




Fig. 2. Visualizing NIF models trained with ReLU activation and an ADAM optimizer [15] (a) NIF network for an Iris perceptron with the communities highlighted (b) NIF network for a multi-layer perceptron trained on the Iris dataset

central it is for information to propagate through the network. The color of the node denotes which community the node is a member of: for standard resolution of $\gamma = 1$, we use Equation (2) to find the communities outlined in Figure 2(a). In Figure 2(b), of the five hidden neurons in the final hidden layer, only three are central to the model’s final prediction. This makes sense as the ReLU activation at those nodes is zero: thus, ReLU effectively stifles information from flowing through the network.

4.2. Network Pruning

It is worthwhile to note that both of the models described above received upwards of 96% accuracy on a held-out test set. Though the state of the art network pruning work has been established [16, 17], we show how pruning arises from a NIF model and compares to the norm in the pruning literature. We ran experiments wherein we zero out the weights and biases of the original model for zero activation neurons in the original model.




Fig. 3. Accuracy vs Number of Zeroed Weights plot for each of the NIF models in Figure 2 respectively

In Figure 3, we report accuracy as we make the weight matrix, W , of the original network more sparse. We find that as we zero out the low NIF important links, we do not lose much accuracy; however, we find a decrease in accuracy that levels out to around the 33% we would expect from a random model for Iris’ three classes.

4.3. Feature Attribution

To obtain a NIF feature attribution for a CNN trained on MNIST using Equation (6), we compare NIF against a few

ATTRIBUTION	K-S STATISTIC	P-VALUE
NIF	1.0	0.011
SHAP	0.75	0.107
INTEGRATED GRADIENTS	0.25	0.996

Table 1. Feature attribution comparison




Fig. 4. NIF attribution: Figure 4(a) contains the original input and Figure 4(b) contains a feature attribution found using Algorithm 1 (red means negative and blue means positive)

popular feature attribution techniques: SHAP [6] and Integrated Gradients [2] in Table 1. Using the two sample Kolmogorov-Smirnov test for goodness of fit between two empirical distributions (here, the raw mutual information attribution between the input and output classes and the attribution in question), we find that NIF outperforms these existing techniques, which means NIF is likely drawn from the same distribution as the raw mutual information. In Figure 4, we validate NIF’s performance by providing a qualitative example of a NIF attribution, obtained via Equation (6). Thus, an information theoretic feature attribution is viable.

5. CONCLUSION

We proposed NIF, Neural Information Flow, a new technique for measuring information flow through deep learning models. Merging mutual information estimation with network science, we find that NIF not only provides insight into which pathways are crucial within a network, but also allows us to leverage fewer parameters at inference time, since we can remove parameters deemed useless by the NIF without loss of accuracy. Finally, NIF recovers an information theoretic feature attribution that aligns with existing techniques.

For future work, NIF provides a unique approach to generic function learning. Suppose we want to learn $f^* : \mathbb{R}^d \rightarrow \{0, 1\}$ from data X and labels Y . Unfortunately, we cannot tell if a trained predictor, \hat{f} , learned the desired, f^* . NIF can clear that doubt by providing insight into the properties of the function learned; particularly, NIF can prune the model class of f . Providing theoretical analysis of how certain network properties (lack of centrality or low clustering coefficient) can give us a more restricted model class for f by decreasing the VC-dimension of the hypothesis space over which we trained the predictor \hat{f} . As such, NIF can help assess which models provide sounder approximations of f^* .

6. REFERENCES

- [1] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, “Why should I trust you?: Explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.
- [2] Mukund Sundararajan, Ankur Taly, and Qiqi Yan, “Axiomatic attribution for deep networks,” in *Proceedings of the 34th International Conference on Machine Learning*, International Convention Centre, Sydney, Australia, 2017, vol. 70, pp. 3319–3328, PMLR.
- [3] Jianbo Chen, Le Song, Martin J. Wainwright, and Michael I. Jordan, “Learning to explain: An information-theoretic perspective on model interpretation,” *ICML 2018*, 2018.
- [4] Albert-László Barabási and Eric Bonabeau, “Scale-free networks,” *Scientific American*, vol. 288, no. 60-69, 2003.
- [5] Mark EJ Newman, “Modularity and community structure in networks,” *Proceedings of the national academy of sciences*, vol. 103, no. 23, pp. 8577–8582, 2006.
- [6] Scott M Lundberg and Su-In Lee, “A unified approach to interpreting model predictions,” in *Advances in Neural Information Processing Systems 30*, pp. 4765–4774. 2017.
- [7] K. D. Bollacker and J. Ghosh, “Linear feature extractors based on mutual information,” in *Proceedings of 13th International Conference on Pattern Recognition*, Aug 1996, vol. 2, pp. 720–724 vol.2.
- [8] Ravid Shwartz-Ziv and Naftali Tishby, “Opening the black box of deep neural networks via information,” *CoRR*, 2017.
- [9] Weihao Gao, Sreeram Kannan, Sewoong Oh, and Pramod Viswanath, “Estimating mutual information for discrete-continuous mixtures,” in *Advances in neural information processing systems*, 2017, pp. 5986–5997.
- [10] Morteza Noshad, Yu Zeng, and Alfred O Hero, “Scalable mutual information estimation using dependence graphs,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 2962–2966.
- [11] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm, “Mutual information neural estimation,” in *Proceedings of the 35th International Conference on Machine Learning*, Jennifer Dy and Andreas Krause, Eds., Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018, vol. 80 of *Proceedings of Machine Learning Research*, pp. 531–540, PMLR.
- [12] David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller, “How to explain individual classification decisions,” *Journal of Machine Learning Research*, vol. 11, no. Jun, pp. 1803–1831, 2010.
- [13] Kartikeya Bhardwaj, Naveen Suda, and Radu Marculescu, “Dream distillation: A data-independent model compression framework,” *arXiv preprint arXiv:1905.07072*, 2019.
- [14] Dua Dheeru and Efi Karra Taniskidou, “UCI machine learning repository,” 2017.
- [15] Diederik P. Kingma and Jimmy Ba, “ADAM: A method for stochastic optimization,” 2014, arxiv:1412.6980: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- [16] Christos Louizos, Max Welling, and Diederik P. Kingma, “Learning sparse neural networks through L_0 regularization,” in *International Conference on Learning Representations*, 2018.
- [17] L. Theis, I. Korshunova, A. Tejani, and F. Huszr, “Faster gaze prediction with dense networks and fisher pruning,” arXiv:1801.05787, 2018.