Deep Weighted Averaging Classifiers

Dallas Card¹

dcard@cmu.edu

¹Machine Learning Department
Carnegie Mellon University
Pittsburgh, Pennsylvania

Michael Zhang²

mjqzhang@cs.washington.edu

²Paul G. Allen School of Computer
Science and Engineering
University of Washington
Seattle, Washington

Noah A. Smith^{2,3}

nasmith@cs.washington.edu

³Allen Institute for Artificial
Intelligence
Seattle, Washington

ABSTRACT

Recent advances in deep learning have achieved impressive gains in classification accuracy on a variety of types of data, including images and text. Despite these gains, however, concerns have been raised about the calibration, robustness, and interpretability of these models. In this paper we propose a simple way to modify any conventional deep architecture to automatically provide more transparent explanations for classification decisions, as well as an intuitive notion of the *credibility* of each prediction. Specifically, we draw on ideas from nonparametric kernel regression, and propose to predict labels based on a weighted sum of training instances, where the weights are determined by distance in a learned instance-embedding space. Working within the framework of *conformal methods*, we propose a new measure of nonconformity suggested by our model, and experimentally validate the accompanying theoretical expectations, demonstrating improved transparency, controlled error rates, and robustness to out-of-domain data, without compromising on accuracy or calibration.

CCS CONCEPTS

• Computing methodologies → Supervised learning.

KEYWORDS

interpretability, credibility, conformal methods

ACM Reference Format:

Dallas Card¹, Michael Zhang², and Noah A. Smith^{2,3}. 2019. Deep Weighted Averaging Classifiers. In *FAT* '19: Conference on Fairness, Accountability, and Transparency*, January 29–31, 2019, Atlanta, GA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3287560.3287595>

1 INTRODUCTION

For any domain involving complex, structured, or high-dimensional data, deep learning has rapidly become the dominant approach for training classifiers. Although deep learning is somewhat vaguely defined, we will use it here to refer to any architecture which makes a prediction based on the output of a function involving a series of linear and non-linear transformations of the input representation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAT '19, January 29–31, 2019, Atlanta, GA, USA*

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6125-5/19/01...\$15.00

<https://doi.org/10.1145/3287560.3287595>

While the details of these transformations differ by domain (for example, two-dimensional convolutions are often used for images, whereas sequential models with attention are more common for text), most models for binary or multiclass classification include a final softmax layer to produce a properly normalized probability distribution over the label space. In this paper, we explore an alternative to the softmax, yielding what we call a *deep weighted averaging classifier* (DWAC), and evaluate its potential to deliver equally accurate predictions, while offering greater transparency, interpretability, and robustness.¹

Despite the success of deep learning as a framework, concerns have been raised about the properties of deep models [48]. In addition to typically requiring large amounts of labeled data and computational resources, the parameters tend to be relatively difficult to interpret, compared to more traditional (e.g., linear) methods. In addition, some deep models tend to be poorly calibrated relative to simpler models, despite being more accurate [14]. Finally, careful evaluation and an explosion of work on adversarial examples have demonstrated that many deep models are more brittle than test-set accuracies would suggest [12, 15, 37, 41].

Particularly in light of recent controversy and legislation, such as the General Data Protection Regulation (GDPR) in Europe, there has been rapidly growing interest in developing more interpretable models, and in finding ways to provide explanations for predictions made by machine learning systems. Although there is currently an active debate in the field about how best to conceptualize and operationalize these terms [11, 13], recent research has broadly fallen into two camps. Some work has focused on models that are inherently interpretable, such that an explanation for a decision can be given in terms that are easily understood by humans. This category includes classic models that can easily be simulated by humans, such as decision lists, as well as sparse linear models, where the prediction is based on a weighted sum of features [6, 24, 29, 49, 50, 56]. Other work, meanwhile, has focused on developing methods to provide explanations that approximate the true inner workings of more complex models, in a way that provides some utility to the user or developer of a model beyond what is attainable through more direct means [2, 27, 30, 42–44, 46].

In this paper, we propose a method which, like those in the former category, offers an explanation that is transparent (in that the complete explanation is in terms of a weighted sum of training instances), but also explore ways to approximate this explanation by using only a subset of the relevant instances. While this approach retains some of the inherent complexity of typical deep models (in that it is still difficult to explain *why* the model has weighted

¹Our implementation is available at <https://github.com/dallascard/DWAC>

the training instances as it has for a particular test instance), the mechanism behind the prediction is far more transparent than softmax-based models, and the individual instance weights provide a way for a user to examine the basis of the prediction, and evaluate whether or not the model is doing something reasonable. Similarly, while looking at the nearest neighbors of a test point is a commonly-used heuristic to attempt to understand what a model is doing, that approach is only an approximation for models which map each instance directly to a vector of probabilities.

There is, of course, a long history in machine learning of making predictions directly in terms of training instances, including nearest neighbor methods [9], kernel methods, including support vector machines [5, 8], and transductive learners more broadly [52]. The main novelty here is to adapt any existing deep model to make predictions explicitly in terms of the training data using only a minor modification to the model architecture, and arguing for and demonstrating the advantages offered by this approach.

As we will describe in more detail below, we propose to learn a function which maps from the input representation to a low-dimensional vector representation of each input. Predictions on new instances are then made in terms of a weighting of the label vectors of the training instances, where the weights are a function of the distance from the instance to be evaluated to all training instances, in the low-dimensional space. This is closely related to a long line of past work on *metric learning* [3, 10, 23, 58, 60], but rather than trying to optimize a particular notion of distance (such as Mahalanobis distance), we make use of a fixed distance function, and allow the architecture of standard deep models to do the equivalent work for us. This idea is also related to models which use neural networks to learn a similarity function for specific applications, such as face recognition [7] or text similarity [33]; here we show how this is a more generally applicable way to train models, and we emphasize the connection to interpretability.

Such an approach comes with distinct advantages:

- (1) A precise explanation of why the model makes a specific prediction (label or probability) can be given in terms of a subset of the training examples, rank ordered by distance. Moreover, the weight on each training instance implicitly captures the degree to which the model views the two instances as similar. The explanation is thus given in terms of exemplars or *prototypes*, which have been shown to be an effective approach to interpretability [18].
- (2) In §5.3, we show that, in many cases, a very small subset of training instances can be used to provide an approximate explanation with high fidelity to the complete explanation.
- (3) In addition, it is possible to choose the size of the learned output representation so as to trade off between performance and interpretability. For example, we can use a lower dimensional output representation if we wish to make it easy to directly visualize the embedded training data.
- (4) Even in cases where revealing the training data is not feasible, it is possible to provide an explanation purely in terms of weights and labels. Although this does not reveal the way in which a new instance is viewed (by the model) as similar to past examples, it still provides a quantifiable notion of how unusual the new example is. The form of our model suggests

a natural metric of nonconformity, and in §3.4, we formalize this using the notion of *conformal methods*, describing how the relevant distances can be used to either provide bounds on the error rate (for data drawn from the same distribution), or robustness against outliers.

- (5) Finally, although this model does entail a slight cost in terms of increased computational complexity, the difference in terms of speed and memory requirements at test time can be minimized by pre-computing and storing only the low-dimensional representation of the training data (from the final layer of the model). The cost during training will in most cases be dominated by the other parts of the network, and it is still possible to train such models on large datasets without difficulty. Moreover, in our experiments, this choice seemingly involves no loss in accuracy or calibration.

Our **deep weighted averaging classifiers** (DWACs) are ideally suited to domains where it is possible to directly inspect the training data, such as controlled settings like social science research, medical image analysis, or managing customer feedback. In these domains, DWACs can create a more transparent and interpretable version of any successfully developed deep learning architecture. Although the advantages are diminished in domains where privacy is a concern, presenting information solely in terms of weights and labels still provides a useful way to quantify the credibility of a prediction, even without allowing direct inspection of the original training data.

2 BACKGROUND

2.1 Scope and Notation

In this paper we will be concerned with the problem of classification. In general, we will assume a set of m training instances, $\mathbf{x}_i$ for $i = 1, \dots, m$, with corresponding labels in some categorical label space, $y_i \in \mathcal{Y}$, for $i = 1, \dots, m$, where $c = |\mathcal{Y}|$ is the number of classes. We also assume we will eventually be given a set of n test instances, $\mathbf{x}_i, y_i$, for $i = 1 + m, \dots, n + m$. We will use $\mathbf{h}_i$ to refer to the output representation of our model for instance i . Square brackets $[\cdot]_k$ will denote the k th element of a vector.

2.2 Nonparametric Kernel Regression

We propose to build on a classic method from nonparametric regression, known as Nadaraya-Watson (NW). The original use case of NW was in regression, with $y_i \in \mathbb{R}$, where it is assumed that $y_i = m(\mathbf{x}_i) + \varepsilon$, where $\mathbb{E}[\varepsilon] = 0$ and $\mathbb{V}[\varepsilon] = \sigma^2$. The goal is to estimate the mean function, $m(\mathbf{x})$, which can be expressed in terms of the joint density as

$$m(\mathbf{x}) = \mathbb{E}[Y | X = \mathbf{x}] = \int y P(y | \mathbf{x}) dy = \frac{\int y P(\mathbf{x}, y) dy}{\int P(\mathbf{x}, y) dy}. \quad (1)$$

By approximating the joint density using kernel density estimation, it is possible to re-express this as a weighted sum of training instances, i.e.,

$$\hat{m}(\mathbf{x}) = \frac{\sum_{i=1}^m y_i K(\mathbf{x}, \mathbf{x}_i)}{\sum_{j=1}^m K(\mathbf{x}, \mathbf{x}_j)}, \quad (2)$$

where $K(\mathbf{x}, \mathbf{x}_i)$ is a kernel, such as a Gaussian [35, 57]. It is easy to see that this corresponds to a *linear smoother*, in that it will predict

outputs as a weighted sum of training instances, i.e.,

$$\hat{m}(\mathbf{x}) = \sum_{i=1}^m y_i \alpha_i(\mathbf{x}), \quad (3)$$

where $\alpha_i(\mathbf{x}) = \frac{K(\mathbf{x}, \mathbf{x}_i)}{\sum_{j=1}^m K(\mathbf{x}, \mathbf{x}_j)}$.

This method can easily be adapted to classification by predicting the *probability* of an output label as a weighted sum of training labels, i.e.,

$$P_{NW}(y = k \mid \mathbf{x}) = \sum_{i=1}^m \mathbb{I}(y_i = k) \alpha_i(\mathbf{x}), \quad (4)$$

where $\mathbb{I}(\cdot)$ equals 1 if the condition holds or 0 if not.

The primary limitation on NW is due to the curse of dimensionality: as the dimensionality of $\mathbf{x}$ grows, sparse data becomes a problem, and the notion of proximity becomes problematic [1]. In this work, we show how we can avoid this problem by using neural networks to embed inputs into a space with reasonable dimensionality, and proceed to compute weights over training instances in that space.

2.3 Conformal Methods

Conformal methods refer to a broad set of ideas which aim to provide theoretical guarantees on error rates in classification or regression [26, 45, 47, 54]. Conformal methods can be used with any base classifier or predictor, and work by introducing a generic notion of *nonconformity*. As will be explained in detail below, each possible prediction (i.e., label or value) that can be made for a given test instance can be evaluated in terms of its nonconformity. By comparing these with the nonconformity scores of either all data in a leave-one-out manner, or with a held-out *calibration* set, the equivalent of a *p*-value can be associated with each possible prediction, allowing for thresholding in a way that provides a guarantee on the error rate (for i.i.d. or exchangeable data).

Because of the high computational cost of the leave-one-out approach to conformal predictors, we will focus here on the approach based on a held-out calibration set.² In particular, we will begin by shuffling our training data, and partitioning it into a proper training set ($i = 1, \dots, t$), and a calibration set ($i = t + 1, \dots, m$).³

The fundamental choices in conformal methods are a base classifier and a measure of nonconformity, A . The latter concept, which intuitively corresponds to how *atypical* an instance is, maps a bag of examples (i.e., the proper training set, with labels), and one additional instance ($\mathbf{x}$, the observed features) with one possible label, $k \in \mathcal{Y}$, to a scalar $\eta \in \mathbb{R}$, i.e.,

$$\eta(\mathbf{x}, k) = A\left(\{\mathcal{Z}(\mathbf{x}_i, y_i)\}_{i=1}^t, (\mathbf{x}, k)\right), \quad (5)$$

where $\{\cdot\}$ denotes a *bag*, i.e., a multi-set (potentially containing duplicate instances).

The idea of a measure of nonconformity is quite general, but in practice, the most common approach is to convert a bag of examples into a *model*, and then compare the *prediction* of that model on the training instance ($\mathbf{x}$) with the hypothesized label, k . In particular,

for any probabilistic model, the simplest measure of model-based nonconformity is any inverse monotonic transform of the predicted probability of the hypothesized label, such as the inverse or negation. For example, the following is a valid measure of nonconformity:

$$\eta(\mathbf{x}, k) = -P_{\mathcal{M}(\{\mathcal{Z}(\mathbf{x}_i, y_i)\}_{i=1}^t)}(y = k \mid \mathbf{x}), \quad (6)$$

where $\mathcal{M}(\{\mathcal{Z}(\mathbf{x}_i, y_i)\}_{i=1}^t)$ represents a model trained on the proper training set, $i = 1, \dots, t$.

Conformal methods work by comparing the nonconformity score of each possible prediction for each test instance to the nonconformity scores of all instances in the calibration set. Specifically, we will compute a *p*-value for each hypothetical test instance label equal to the proportion of the calibration instances that have a higher nonconformity score (i.e., a value between 0 and 1 indicating how *conforming* a possible label for a new instance is, relative to the calibration data). More precisely, for a test instance $\mathbf{x}$ and hypothesized label, k ,

$$p(\mathbf{x}, k) = \frac{\sum_{j=t+1}^m \mathbb{I}(\eta(\mathbf{x}_j, y_j) \geq \eta(\mathbf{x}, k))}{m - t}, \quad (7)$$

where η is computed relative to the proper training set for all instances (both calibration and test).

Unlike traditional classifiers, conformal methods may predict anywhere from zero to c labels for a given instance. We will revisit this choice below, but for the moment, the decision rule will be that for each test instance, we make a positive prediction for all labels k for which $p(\mathbf{x}, k) > \varepsilon$, where $\varepsilon \in [0, 1]$ is chosen by the user. By the properties of conformal predictors, these predictions will be asymptotically valid; that is, both theoretically and empirically, for i.i.d. data, this will produce a set of predicted labels for each test instance such that at least $1 - \varepsilon$ of the predicted label sets include the true label, with high probability. Moreover, this property holds for any measure of nonconformity.⁴ Of course, this property is trivially easy to satisfy by predicting all labels for all test instances. In practice, however, better choices of measures of nonconformity will be more or less *efficient*, in that they will tend to produce smaller predicted label sets without compromising on error rate.

In addition, for any given test instance, we can compute two properties of the label that is predicted with the highest probability, which Saunders et al. called *confidence* and *credibility* [45, 47]. Confidence is equal to the largest $1 - \varepsilon$ such that the predicted label set includes only a single label (i.e., one minus the second-largest *p*-value among the possible labels). This corresponds to the probability, according to the model, that the label which is predicted to be most likely is correct. For example, for i.i.d. data, we would expect 92% of predictions with 92% confidence to be correct, and 8% to be incorrect.

Credibility, by contrast, is equal to the largest ε for which the predicted label set is empty (i.e., the largest *p*-value among the possible labels). This corresponds to one minus the model's confidence that *none* of the possible labels are correct. In other words, predictions

²The leave-one-out approach may be superior in terms of statistical efficiency, but it is computationally infeasible for all but very small datasets.

³Our calibration set will also serve as our validation set for early stopping during training.

⁴More precisely, for any measure of nonconformity, and i.i.d. data drawn from any distribution, the unconditional probability that the true label of a random test instance is not in the predicted label set does not exceed ε for any ε . The broader statement follows from the law of large numbers. It is also possible to bound the conditional probability of error (conditional on the training data) with slightly weaker guarantees. For more details, please refer to [53, 54].

with low credibility indicate that even the most likely prediction is relatively nonconforming compared to the calibration instances (with their true labels), and that we should therefore be skeptical of this prediction.

The idea of not predicting any label is somewhat foreign to conventional approaches, as we can only obtain better accuracy by hazarding at least some guess, and in some applications it would be reasonable to still predict the most probable label according to the model, with an associated confidence and credibility. In other circumstances, however, there may be valuable information in an empty label set. For example, if we set ϵ to 0.2, and find that 20% of the test instances end up with empty predicted label sets, this is a strong indication that the predictions for most of the remaining instances are correct (because the overall error rate is expected to be 20%). In practice, this would be an unusual outcome, but we will explore the usefulness of this sort of signal in the context of outlier detection.

3 DEEP WEIGHTED AVERAGING CLASSIFIERS

We now turn to our new method, deep weighted averaging classifiers.

3.1 Model Details

Most neural network models for classification take the form

$$P(y = k | \mathbf{x}) = \frac{\exp([\mathbf{h}]_k)}{\sum_{j=1}^c \exp([\mathbf{h}]_j)}, \quad (8)$$

where $\mathbf{h} = \mathbf{W} \cdot f(\mathbf{x}) + \mathbf{b}$. In this architecture, $f(\mathbf{x})$ embeds the input in a model-specific way (e.g., a convolutional or recurrent network) into a lower-dimensional vector representation. Although $\mathbf{W}$ and $\mathbf{b}$ could be folded into $f(\mathbf{x})$, we make them explicit to emphasize that $\mathbf{W}$ is required to project the output of $f(\mathbf{x})$ (which is a vector of arbitrary length) down to a vector of length c (the number of classes). The softmax (Eq. 8) then projects this vector onto the simplex.

Our proposed alternative is to leave $f(\mathbf{x})$ unchanged, eliminate the softmax, and to redefine the predicted probability as

$$P(y = k | \mathbf{x}) = \frac{\sum_{i=1}^t \mathbb{I}(y_i = k) w(\mathbf{h}, \mathbf{h}_i)}{\sum_{j=1}^t w(\mathbf{h}, \mathbf{h}_j)}, \quad (9)$$

i.e., a weighted sum over the labels of the instances in the proper training set, where $\mathbf{h}$ is defined as above and $w(\mathbf{h}, \mathbf{h}_i)$ is a function of the similarity between the embeddings of $\mathbf{x}$ and $\mathbf{x}_i$ in the low-dimensional space, according to some metric. In this architecture, the dimensionality of $\mathbf{h}$ is arbitrary, and the size of $\mathbf{W}$ and $\mathbf{b}$ can be modified as necessary.⁵

An obvious choice of weight function is a Gaussian kernel operating on Euclidean distance, i.e.,

$$w(\mathbf{h}, \mathbf{h}_i) = \exp\left(\frac{-\|\mathbf{h} - \mathbf{h}_i\|_2^2}{2\sigma}\right). \quad (10)$$

⁵We could of course dispense with $\mathbf{W}$ and $\mathbf{b}$ here, and compute $w(\cdot, \cdot)$ directly on the output of $f(\mathbf{x})$, but we wish to remain as close as possible to the softmax model for the purpose of comparison, while allowing for the possibility of varying the size of $\mathbf{h}$.

Typically, in using NW, or other kernel smoothers, one needs to choose the bandwidth, equivalent to σ in equation (10). However, because we will assume that we will be building this classifier on top of a high-capacity embedding network, $f(\mathbf{x})$, we will simply fix $\sigma = \frac{1}{2}$, and force $f(\mathbf{x})$ to adapt to this distance function.

3.2 Training

In order to learn $\mathbf{W}$, $\mathbf{b}$, and all parameters of $f(\mathbf{x})$, we use stochastic gradient descent to optimize the log loss on the training data. Because it is impractical to compute the exact probabilities according to the model during training (because they depend on all training instances), we instead rely on an approximation based on the other instances within each minibatch. Specifically, on each epoch of training, we shuffle all instances in the proper training set into minibatches of size B . For each minibatch, $\mathcal{B}$, we then minimize

$$\mathcal{L}(\mathcal{B}) = \frac{1}{B} \sum_{j \in \mathcal{B}} \sum_{k=1}^c -\mathbb{I}(y_j = k) \log \hat{P}(y_j = k | \mathbf{x}_j), \quad (11)$$

where

$$\hat{P}(y_j = k | \mathbf{x}_j) = \frac{\sum_{i \in \mathcal{B} \setminus \{j\}} \mathbb{I}(y_i = k) w(\mathbf{h}_j, \mathbf{h}_i)}{\sum_{l \in \mathcal{B} \setminus \{j\}} w(\mathbf{h}_j, \mathbf{h}_l)}. \quad (12)$$

As usual, the loss function in equation (11) can be augmented with a regularization term if desired.

Although computing all pairwise distances between many points is relatively expensive, this can be done efficiently for minibatches using standard matrix operations on a GPU. Specifically, a forward pass through the last layer of a DWAC model (i.e., computing probabilities for one minibatch) requires $O(B^2 h)$, where B is the size of the minibatch and $h = \text{len}(\mathbf{h})$. This will typically be larger than the last layer of the softmax classifier, which is $O(Bhc)$, where c is the number of classes, but in most cases will be dominated by the cost of computing $f(\mathbf{x})$. In practice, the only significant increase in training time is due to the need to embed all training instances in order to estimate performance on a validation set after each epoch. As such, the training runtime will tend to be no worse than twice that of training a softmax model. Similarly, at test time, the computational cost of making a prediction on one test instance is dominated by the cost of embedding the training data. However, this can be pre-computed after training, and only the low-dimensional $\mathbf{h}$ vectors need be stored.⁶

3.3 Prediction and Explanations

Once the model has been trained, predictions can easily be made using the entire training dataset, rather than using a subset, as when computing the loss during training. The complete and explicit explanation for why the model predicts a particular label or probability can then be given explicitly in terms of the training instances, along with the weight on each instance. Moreover, if we consider the closest points (which will be most heavily weighted) as being the most similar, relevant, or important, it is reasonable to provide a sorted list of examples as the explanation. Because the later examples will carry less weight, in many cases only a subset of instances needs to be provided (because for many instances,

⁶Note that we only need to compute distances in the low-dimensional space, for which we can choose an appropriately small dimensionality.

the lower-weighted training instances will be unable to affect the prediction, no matter what their labels may be).

If we wish to provide an even simpler but approximate explanation, we can also choose to provide only the closest k examples as the explanation, which is a commonly used heuristic for trying to understand model behavior. Although we would not expect that using only a small set of examples would provide a well-calibrated *probability*, it could still provide a reasonable approximate explanation for why the model predicted a particular *label*, assuming that there is strong agreement between predictions made using such a subset and the full model (which we will empirically evaluate in the experiments below).

3.4 Confidence and Credibility

As discussed in section §2.3 above, for any probabilistic classifier, we can use any monotonic transformation of the predicted probabilities which reverses their order (such as $-P(\mathbf{x})$ or $1/P(\mathbf{x})$) as a valid measure of nonconformity. However, our proposed architecture suggests another measure, namely the negated unnormalized weighted sum of training labels of the hypothesized class, i.e.,

$$\eta(\mathbf{x}, k) = A\left(\mathcal{I}(\mathbf{x}_i, y_i)_{i=1}^t, (\mathbf{x}, k)\right) = -\sum_{i=1}^t \mathbb{I}(y_i = k) w(\mathbf{h}, \mathbf{h}_i). \quad (13)$$

Because of the properties of conformal methods, this measure of nonconformity is automatically *valid*. It may, however, be more or less *efficient* than other measures, such as ones based on probabilities. We note, however, that this proposed measure has an intuitive explanation in terms of how close the training points of the predicted class are (in the embedded space) to the instance for which we wish to make a prediction. Naturally the absolute distance has no meaning in the embedded space, but we avoid this problem by scaling the measure of nonconformity relative to calibration data in order to obtain a p -value, as is always the case in conformal methods.

When using probability as the basis of nonconformity, the farther a point is from the decision boundary, the higher will be its predicted probability, and therefore its credibility. Under the measure we propose in equation (13), by contrast, predictions will only be associated with high credibility when the embedded representation of that instance is relatively close to the embedded training instances. That is, if we encounter an instance that is unlike anything seen in the training data, and if the model embeds that instance such that it is far away from all embedded training instances, then this measure will tell us that it is highly nonconforming for all classes, which will result in the model's prediction having very low credibility. As we show below, this is a useful way to quantify the degree to which we should be skeptical of the model's prediction.

4 EXPERIMENTS

To demonstrate the potential of DWAC models, we provide a range of empirical evaluations on a variety of datasets. In addition to showing that our proposed approach is capable of obtaining equivalently accurate predictions, we also compare to a baseline in terms of calibration and robustness to outliers, and illustrate the sorts of explanations offered by a DWAC model. We also empirically validate that the theoretical guarantees claimed by conformal methods

Dataset	Type	# classes	# instances	# features
Adult Income	Tabular	2	45000	103
Coverttype	Tabular	7	581000	54
Lending Club	Tabular	2	626000	157
Fashion MNIST	Image	10	60000	784
CIFAR-10	Image	10	60000	3072
Subjectivity	Text	2	10000	20000
Stack Overflow	Text	20	20000	15000
IMDB	Text	2	50000	112000
Amazon	Text	5	190000	144000

Table 1: Properties of datasets used in this paper. For text data, we report vocabulary size as the number of features.

hold for both the softmax and the DWAC model, both with and without our proposed measure of nonconformity.

In all cases we report accuracy and calibration (i.e., the accuracy of the predicted probabilities), measuring the latter in terms of mean absolute error (MAE), using the adaptive binning approach of Nguyen and O'Connor [38]. For the initial comparison between models, we only consider conventional prediction (i.e., only using the top-scoring label predicted by each model), and separately evaluate conformal prediction in §5.4. Note that our purpose here is not to demonstrate state-of-the-art performance on any particular task, but rather to show that our proposed modification works for a wide variety of architectures.

In order to evaluate robustness to outliers, we consider two approaches. The first is to drop one class from a multiclass dataset, and treat the held-out class as out-of-domain data. The other approach is to find a dataset that has a similar input representation, but is fundamentally different in terms of content, and again, treat these instances as out-of-domain.

4.1 Datasets

For our experiments we make use of datasets of three different types (tabular, image, and text), including both binary and multiclass problems. For tabular data, we use the familiar Adult Income [20] and Coverttype [4] datasets available from the UCI machine learning repository, as well as the Lending Club loan dataset available through Kaggle. For images we use CIFAR-10 [21] and Fashion MNIST [59]. For text we use paragraph-length product reviews (Amazon; 1–5 stars) [32] and movie reviews (IMDB; positive or negative) [31], sentences extracted from movie reviews and labeled in terms of *subjectivity* [40], and a dataset of Stack Overflow question titles sampled from 20 different categories [61]. Table 1 summarizes the most important properties of these datasets; more details and URLs can be found in supplementary material.

4.2 Models and Training

In all cases we choose a base model appropriate to the data. For the tabular data, we use a simple three-layer multi-layer perceptron. For images, we use multi-layer convolutional neural networks. For text datasets, we use a shallow convolutional model with attention [34]. In all cases we compare DWAC and softmax models of equivalent size, but also explore varying the dimensionality of $\mathbf{h}$ in the DWAC model. We use the standard train/test split where available, and

Dataset	Accuracy $\uparrow$	
	Softmax	DWAC
Adult Income	0.851 (± 0.002)	0.850 (± 0.002)
Covertypes	0.774 (± 0.003)	0.760 (± 0.001)
Lending Club	0.956 (± 0.001)	0.955 (± 0.001)
Fashion MNIST	0.928 (± 0.002)	0.927 (± 0.002)
CIFAR-10	0.898 (± 0.004)	0.897 (± 0.009)
Subjectivity	0.948 (± 0.002)	0.946 (± 0.004)
Stack Overflow	0.869 (± 0.005)	0.866 (± 0.008)
IMDB	0.905 (± 0.002)	0.904 (± 0.001)
Amazon	0.740 (± 0.002)	0.738 (± 0.002)

Dataset	Calibration (MAE) $\downarrow$	
	Softmax	DWAC
Adult Income	0.012 (± 0.006)	0.018 (± 0.002)
Covertypes	0.005 (± 0.001)	0.010 (± 0.001)
Lending Club	0.007 (± 0.001)	0.014 (± 0.001)
FashionMNIST	0.006 (± 0.001)	0.003 (± 0.001)
CIFAR-10	0.011 (± 0.001)	0.009 (± 0.002)
Subjectivity	0.020 (± 0.006)	0.023 (± 0.006)
Stack Overflow	0.009 (± 0.001)	0.010 (± 0.001)
IMDB	0.029 (± 0.006)	0.024 (± 0.010)
Amazon	0.008 (± 0.002)	0.004 (± 0.001)

Table 2: Accuracy (higher is better) and calibration (lower is better) on various datasets using the single best-scoring predicted label from softmax and DWAC models of equivalent size, with standard deviations in parentheses.

otherwise sample a random 10% of the data for a test set, and always use a random 10% of the training data as a validation/calibration set. For measuring accuracy and calibration on test data, we average over 5 trials with different splits of the training data into a proper training set and a validation/calibration set, with the same split being given to both models. For both models we use Adam [19] with an initial learning rate of 0.001 and early stopping. For more details, please refer to supplementary material.

5 RESULTS

5.1 Classification Performance

As shown in Table 2, except for one dataset (Covertypes) the performance of our model is indistinguishable from a softmax model of the same size in terms of accuracy. As noted above, these results are based on the single best-scoring label from each model, without yet incorporating the idea of conformal prediction. For calibration, the DWAC model is sometimes slightly better and sometimes slightly worse, although we note in passing that, at least for these models and datasets, the predictions from both models are quite well calibrated, such that the predicted probabilities are relatively reliable, at least for in-domain test data. As expected, runtime during training was approximately 50% longer per epoch than for the equivalent softmax model, with a similar number of epochs required.

One advantage of DWAC is the freedom to choose the dimensionality of the final output layer, and Figure 1 illustrates the impact of this choice on the performance of our model. While using the same dimensionality as the softmax model gives equivalent performance,

Figure 1: Accuracy of DWAC in comparison to a softmax model on the 10-class Fashion MNIST dataset for varying dimensionality of the output layer. Performance is indistinguishable for a DWAC model of the same size, but accuracy drops if we decrease the size of the output layer too much.

Figure 2: Learned embeddings of the Fashion MNIST training data when using a DWAC model with a two-dimensional output layer.

the same accuracy can often be obtained using a lower-dimensional representation (i.e., few total parameters). In some cases, however, reducing the dimensionality too much (e.g., two-dimensional output for Fashion MNIST) results in a slight degradation of performance.

On the other hand, using a two-dimensional output layer means that we are able to more easily visualize the learned embeddings, without requiring an additional dimensionality reduction step, e.g., using PCA or t-SNE [51]. In this way, we could look directly at where a test instance is being embedded, relative to training instances, with no loss of fidelity. Figure 2 shows the embeddings learned by our model for the Fashion MNIST training data using a two-dimensional output layer. Pleasingly, there is a natural semantics to this space, with all of the footwear occurring close together, and the “shirt” class being centrally located relative to related classes (t-shirt, pullover, etc.).

Figure 3: Two examples of predictions made by the DWAC model on the Fashion MNIST dataset. An approximate explanation for the model’s prediction on the test image (single image) is the four most-highly weighted training images (row of four), along with their weights.

5.2 Interpretability and Explanations

Recall that explanations for predictions made by our model are given in terms of a weighted sum of training instances. Figure 3 (top) shows an example of a partial explanation provided by our model for an image dataset. A test image is shown in the top row, along with its true label, and the four closest images from the proper training set (as measured in the embedded space) are shown below it, along with their weights. In this case, all share the same label, and all contribute approximately equally to the prediction.

As a contrasting example, Figure 3 (bottom) shows an example which has poor support among the nearest training points. Although the prediction is correct, and the closest training images appear visually similar, this sort of wide-legged trouser is quite rare in the dataset (most trousers included have narrow legs). As such, the model has clearly not learned as well how to embed such images into the low-dimensional space. The low weights indicate that we should be skeptical of this prediction.

Images are relatively easy to compare at a glance; text, by contrast, may be more difficult. Nevertheless, the explanations given by our model for the predictions on text data are in some cases very meaningful. For the Stack Overflow dataset, for example, many instances are almost trivially easy, in that the label is part of the question. Not surprisingly, these examples tend to have many highly-weighted neighbors which provide a convincing explanation. Such an example is shown in Table 3 (top). In other cases, the text is

Weight	Sentence	Label
Test	Drupal 6 dynamic menu item	Drupal
0.994	drupal how to limit the display results of the calendar view	Drupal
0.994	Drupal : Different output for first item in a block	Drupal
0.994	changing user role in drupal	Drupal
Weight	Sentence	Label
Test	save data from editable division	Ajax
0.214	Pass data from workspace to a function	Matlab
0.133	upload data from excel to access using java	Apache
0.130	Finding incorrectly - formatted email addresses in a CSV file	Excel

Table 3: Two examples from the Stack Overflow dataset with approximate explanations from a DWAC model: an easy example with many close neighbours (top) and a more difficult example with no close neighbours (bottom). The first line is the test instance in both cases.

more ambiguous. Table 3 (bottom) shows an example with very little support, for which a user might rightly be skeptical of the model, based on both the weights and the explanation.

Finally, because we can use any deep model to compute $f(\mathbf{x})$, we are free to choose one with characteristics we prefer, including interpretability. For our text classification experiments, we chose a base model originally proposed for interpretable classification of medical texts [34]. To further unpack the explanation given for a prediction, we could, for example, inspect the attention weights for a particular pair of sentences to understand the importance of each word in context.

For a comparable pair of examples from a tabular dataset, please refer to supplementary material.

5.3 Approximate Explanations

Because the weights on training instances decrease exponentially with distance, the closest training instances will contribute the most to the prediction. In some cases, only relatively few training instances will be required to fully determine the model’s predicted label (because beyond this the remaining instances will lack sufficient weight to alter which class will be most highly-weighted). In practice, most test instances tend to require a substantial proportion of the nearest training instances in order to cross this threshold. However, even considering a much smaller number of the closest training instances may still result in high agreement with the prediction based on all of the data. Table 4 shows the agreement with the full model if we only consider the top k neighbors to each test instance. For most datasets, this agreement is very high, even for a very small number of neighbors.

Dataset	k=1	k=5	k=10	k=100
Adult Income	0.85	0.91	0.93	0.96
Coverttype	0.76	0.77	0.77	0.83
Lending Club	0.96	0.98	0.99	0.99
Fashion MNIST	0.95	0.97	0.98	0.99
CIFAR-10	0.96	0.98	0.98	0.99
Subjectivity	0.98	0.99	0.99	1.00
Stack Overflow	0.81	0.84	0.85	0.89
IMDB	0.95	0.98	0.98	0.99
Amazon	0.78	0.87	0.90	0.94

Table 4: Impact of considering a subset of the training instances as an approximate explanation: the columns show agreement with the full model on the single most-probable label when basing the prediction on only the k closest training instances.

5.4 Confidence and Credibility

The above results only considered the single top-scoring label predicted for each instance; here we extend both models to the conformal setting. To verify the theoretical expectations of conformal methods, we show that our proposed measure of nonconformity correctly works to maintain a desired error rate in making predictions. Figure 4 shows the results for the Fashion MNIST dataset, where we vary ϵ (the desired maximum error rate) from 0 to 0.2. The top subfigure (a) shows the proportion of predictions on the test set which are *correct* (that is, which contain the true label), for the softmax model using negative probability (probs) as a measure of nonconformity, our model using the same measure, and our model using our proposed measure of nonconformity (weights) given in equation (13). As can be seen, all three demonstrate correct coverage, with all lines close to but not exceeding the expected proportion $(1 - \epsilon)$ across the full range (shown as a dashed line on the top subfigure). Note that this is not the same as accuracy, as some predictions may contain multiple labels.

The second and third subfigures show the same lines for the proportion of predicted label sets that are empty (b) or contain multiple labels (c). The bottom figure shows the mean number of labels in all non-empty predicted label sets. In all cases, the dashed line represents an optimal outcome (that is, a proportion of predicted label sets equal to ϵ are empty, all other predictions are correct, and no predictions contain multiple labels).

As can be seen, the softmax and DWAC models give indistinguishable results when using the same measure of nonconformity. Our proposed measure of nonconformity, by contrast, appears to be slightly less efficient, producing slightly more predicted label sets with multiple labels, but also slightly more empty label sets, which represent identifiable errors contributing to the proportion of incorrect predictions.

The advantage of our proposed measure, however, comes in robustness to out-of-domain data. If we train a model on the Fashion MNIST dataset, and then ask it to make predictions on the original MNIST digits dataset (which has the same size and data format, but consists of hand-written digits rather than items of clothing), we would hope that a good model would predict relatively low probability for all such instances. In fact, as has been previously observed

Figure 4: Coverage of various models on the Fashion MNIST test data, as we vary the desired maximum error rate (ϵ). From top to bottom, the subfigures show (a) the proportion of predicted label sets that are correct (contain the true label); (b) that are empty (make no prediction); (c) that contain multiple labels; and (d) the mean number of labels in non-empty prediction sets. The softmax and DWAC models give nearly identical results when using negative probability as a measure of nonconformity (probs). Our proposed measure (weights) has an indistinguishable error rate, but is slightly less efficient. The dashed line in each figure represents an optimal response.

[37], deep models tend to predict relatively high probabilities, even for out-of-domain data, and this is also true of DWAC models.

Fortunately, the *credibility* score from a conformal predictor provides a meaningful estimate of how much we should trust the corresponding prediction.⁷ Both the softmax model (using negative

⁷Specifically, as mentioned above, it is equal to one minus the model probability that none of the labels should be given to this instance.

Figure 5: Empirical distribution of credibility scores from the softmax (top) and DWAC (bottom) models when trained on Fashion MNIST and tested on MNIST digits (which have the same input format but different content), with the latter using our proposed measure of nonconformity.

probability as a measure of nonconformity), and the DWAC model (using our proposed measure of nonconformity) give low credibility to the vast majority of out-of-domain examples, as shown in Figure 5. The credibility scores from DWAC however, are noticeably shifted closer to zero, indicating that the sum of the weights of the corresponding class is a better measure when we are concerned about the possibility of out-of-domain data. (For in-domain data, the credibility values will be approximately uniformly distributed).

We see a similar result when training on CIFAR-10 images, and predicting on the Tiny Images dataset, and an even more extreme difference in the case of the Covertypes dataset, where we treat one out of seven classes as out-of-domain data, and train a model on only the six remaining classes. For that setup, the mean credibility score from the softmax model on the out-of-domain data is 0.9, whereas for the DWAC model with our measure of nonconformity it is 0.2 (figures in supplementary material). This indicates that to the softmax model, the held-out images “look like” even more extreme versions of one of the other classes, whereas the DWAC model correctly recognizes that the held-out images are relatively unlike the training instances (relative to the calibration data).

6 DISCUSSION AND FUTURE WORK

The idea of interpretability has always been important in statistics and machine learning, but has taken on a renewed urgency with the increased expressive power of deep learning models and the expanded deployment of machine learning systems in society. No single approach to interpretability is likely to solve all problems; here we have focused on adapting existing models so as to make their predictions more transparent, by decomposing them into a sum over training instances, the support for each of which can be inspected. As emphasized by papers that have made use of user studies [16, 22, 36], careful empirical work is required to evaluate

the effectiveness of explanations, and we leave such an evaluation of this approach for future work.

A few recent papers have also sought to formalize the question of when to trust a classifier. Jiang et al. [17] present a method for determining how much any existing classifier should be trusted on any given test point, based on the relative distance to the most-likely and second-most-likely clusters, with clusters based on a pruned training set. This appears to be a theoretically well-motivated and empirically effective technique, but is more focused on trust than interpretability. In an unpublished paper, Papernot and McDaniel [39] propose to train a conventional deep model, but make predictions using a k -nearest neighbours approach, with distance computed using all internal nodes of the network. We, by contrast, propose to train a model using the same form as will be used for prediction, to make predictions based on a weighted sum over the entire training set, and rely on similarities computed in the low-dimensional space of the final layer. Wallace et al. [55] apply the method from Papernot and McDaniel to the problem of text classification and explore the implications for interpretability.

There have also been several papers focused on the problem of predicting whether data is in-domain or out-of-domain [25, 28]. Many of these build on Hendrycks and Gimpel [15], who observed that the predicted probabilities contain some useful signal as to whether data came from in-domain or out-of-domain, and proposed to use this to differentiate between the two by thresholding these probabilities. The authors did not, however, make the connection to conformal methods, which offer a more theoretically sound basis on which to make these decisions, as well as greater flexibility of metrics, beyond just predicted probability.

There are several natural extensions to this work which could be pursued, such as applying a similar architecture to regression or multi-label problems, as well as extending the idea of nonconformity to provide class-conditional guarantees.

7 CONCLUSIONS

In this paper we have demonstrated that even for sophisticated deep learning models, it is possible to create a nearly identical model, with all of the same desirable properties, that nevertheless provides an explanation for any prediction in terms of a weighted sum of training instances. In domains where the training data can be freely inspected, this provides greater transparency by revealing the many components that are explicitly contributing to a model’s prediction, each of which can in principle be inspected and interrogated. Moreover, this method can build on top of other approaches to interpretability, by choosing an appropriate base model. When an approximate explanation will suffice, then using only a small subset of the training instances provides a natural high-fidelity approximation. More importantly, representing the prediction in this manner suggests a natural alternative measure of nonconformity, which, as we have shown, provides a more effective measure for detecting out-of-domain examples. Even in cases where training data cannot be shared (for privacy reasons, for example), this use of conformal methods still allows us to assert a quantitative estimate of the credibility of an individual prediction, one that is far more meaningful than the model’s predicted probability.

ACKNOWLEDGMENTS

We acknowledge helpful comments from the anonymous reviewers. This research was funded in part by a University of Washington Innovation award and an REU supplement to NSF grant IIS-1562364.

REFERENCES

- [1] Charu C. Aggarwal, Alexander Hinneburg, and Daniel A. Keim. 2001. On the Surprising Behavior of Distance Metrics in High Dimensional Spaces. In *Proceedings of International Conference on Database Theory*.
- [2] Osbert Bastani, Carolyn Kim, and Hamsa Bastani. 2017. Interpreting Blackbox Models via Model Extraction. *CoRR* abs/1705.08504 (2017).
- [3] Aurélien Bellet, Amaury Habrard, and Marc Sebban. 2013. A Survey on Metric Learning for Feature Vectors and Structured Data. *CoRR* abs/1306.6709 (2013).
- [4] Jock A. Blackard and Denis J. Dean. 1999. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and Electronics in Agriculture* 24 (1999).
- [5] Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. 1992. A Training Algorithm for Optimal Margin Classifiers. In *Proceedings of COLT*.
- [6] Leo Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. 1984. *Classification and Regression Trees*. Wadsworth.
- [7] Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. Learning a similarity metric discriminatively, with application to face verification. *Proceedings of CVPR*.
- [8] Corinna Cortes and Vladimir Vapnik. 1995. Support-Vector Networks. *Machine Learning* 20, 3 (1995), 273–297.
- [9] T. Cover and P. Hart. 1967. Nearest Neighbor Pattern Classification. *IEEE Transactions on Information Theory* 13, 1 (1967), 21–27.
- [10] Jason V. Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S. Dhillon. 2007. Information-theoretic Metric Learning. In *Proceedings of ICML*.
- [11] Finale Doshi-Velez and Been Kim. 2017. Towards A Rigorous Science of Interpretable Machine Learning. *CoRR* (2017). <https://arxiv.org/abs/1702.08608>
- [12] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *Proceedings of ICLR*.
- [13] Riccardo Guidotti, Anna Monreale, Franco Turini, Dino Pedreschi, and Fosca Giannotti. 2018. A Survey Of Methods For Explaining Black Box Models. *CoRR* abs/1802.01933 (2018).
- [14] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On Calibration of Modern Neural Networks. In *Proceedings of ICML*.
- [15] Dan Hendrycks and Kevin Gimpel. 2017. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In *Proceedings of ICLR*.
- [16] Johan Huysmans, Karel Dejaeger, Christophe Mues, Jan Vanthienen, and Bart Baesens. 2011. An Empirical Evaluation of the Comprehensibility of Decision Table, Tree and Rule Based Predictive Models. *Decis. Support Syst.* 51, 1 (2011).
- [17] Heinrich Jiang, Been Kim, and Maya R. Gupta. 2018. To Trust Or Not To Trust A Classifier. *CoRR* abs/1805.11783 (2018).
- [18] Been Kim, Cynthia Rudin, and Julie A Shah. 2014. The Bayesian Case Model: A Generative Approach for Case-Based Reasoning and Prototype Classification. In *Proceedings of NIPS*.
- [19] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* (2014). <https://arxiv.org/abs/1412.6980>
- [20] Ron Kohavi. 1996. Scaling Up the Accuracy of Naive-Bayes Classifiers: a Decision-Tree Hybrid. In *Proceedings of KDD*.
- [21] Alex Krizhevsky and Geoffrey Hinton. 2009. *Learning Multiple Layers of Features from Tiny Images*. Technical Report. University of Toronto.
- [22] Todd Kulesza, Simone Stumpf, Margaret M. Burnett, Sherry Yang, Irwin Kwan, and Weng-Keen Wong. 2013. Too much, too little, or just right? Ways explanations impact end users' mental models. *2013 IEEE Symposium on Visual Languages and Human Centric Computing* (2013).
- [23] Brian Kulis. 2013. Metric learning: A survey. *Foundations and Trends® in Machine Learning* 5, 4 (2013), 287–364.
- [24] Himabindu Lakkaraju, Stephen H. Bach, and Jure Leskovec. 2016. Interpretable Decision Sets: A Joint Framework for Description and Prediction. In *Proceedings of KDD*.
- [25] Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. 2018. Training Confidence-calibrated Classifiers for Detecting Out-of-Distribution Samples. In *Proceedings of ICLR*.
- [26] Jing Lei, James M. Robins, and Larry A. Wasserman. 2014. Distribution Free Prediction Sets. *J Am Stat Assoc* 108 (2014), 278–287. Issue 510.
- [27] Tao Lei, Regina Barzilay, and Tommi S. Jaakkola. 2016. Rationalizing Neural Predictions. In *Proceedings of EMNLP*.
- [28] Shiyu Liang, Yixuan Li, and R. Srikant. 2018. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. In *Proceedings of ICLR*.
- [29] Yin Lou, Rich Caruana, and Johannes Gehrke. 2012. Intelligible Models for Classification and Regression. In *Proceedings KDD*.
- [30] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of NIPS*.
- [31] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning Word Vectors for Sentiment Analysis. In *Proceedings of ACL*.
- [32] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of SIGIR*.
- [33] Jonas Mueller and Aditya Thyagarajan. 2016. Siamese Recurrent Architectures for Learning Sentence Similarity. In *Proceedings of AAAI*.
- [34] James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. Explainable Prediction of Medical Codes from Clinical Text. In *Proceedings of NAACL*.
- [35] Ėlizar A. Nadaraya. 1964. On Estimating Regression. *Theory of Probability and its Applications* 9 (1964), 141–142. Issue 1.
- [36] Menaka Narayanan, Emily Chen, Jeffrey He, Been Kim, Sam Gershman, and Finale Doshi-Velez. 2018. How do Humans Understand Explanations from Machine Learning Systems? An Evaluation of the Human-Interpretability of Explanation. *CoRR* abs/1802.00682 (2018).
- [37] Anh Mai Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015), 427–436.
- [38] Khanh Nguyen and Brendan O'Connor. 2015. Posterior calibration and exploratory analysis for natural language processing models.. In *Proceedings of EMNLP*.
- [39] Patrick McDaniel Nicolas Papernot. 2018. Deep k-Nearest Neighbors: Towards Confidential, Interpretable and Robust Deep Learning. *CoRR* (2018). <https://arxiv.org/abs/1803.04765>
- [40] Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of ACL*.
- [41] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. 2018. Do CIFAR-10 Classifiers Generalize to CIFAR-10? *CoRR* abs/1806.00451 (2018).
- [42] Marco Tùlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proceedings of NAACL*.
- [43] Marco Tùlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Anchors: High-Precision Model-Agnostic Explanations. In *Proceedings of AAAI*.
- [44] Marco Tùlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Semantically Equivalent Adversarial Rules for Debugging NLP models. In *Proceedings ACL*.
- [45] Craig Saunders, Alexander Gammerman, and Vladimir Vovk. 1999. Transduction with Confidence and Credibility. In *IJCAI*.
- [46] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. 2016. Grad-CAM: Why did you say that? *CoRR* abs/1611.07450 (2016).
- [47] Glenn Shafer and Vladimir Vovk. 2008. A Tutorial on Conformal Prediction. *Journal of Machine Learning Research* 9 (2008), 371–421.
- [48] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *CoRR* (2013). <http://arxiv.org/abs/1312.6199>
- [49] Rob Tibshirani. 1996. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society (Series B)* 58 (1996), 267–288.
- [50] Berk Ustun and Cynthia Rudin. 2015. Sparsely linear integer models for optimized medical scoring systems. *Machine Learning* 102 (2015), 349–391.
- [51] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *JMLR* 9 (November 2008), 2579–2605.
- [52] Vladimir Vapnik. 1998. *Statistical learning theory*. Wiley.
- [53] Vladimir Vovk. 2012. Conditional validity of inductive conformal predictors. In *Asian Conference on Machine Learning*.
- [54] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. 2005. *Algorithmic learning in a random world*. Springer.
- [55] Eric Wallace, Shi Feng, and Jordan Boyd-Graber. 2018. Interpreting Neural Networks with Nearest Neighbors. In *Proceedings of the Workshop on Analyzing and Interpreting Neural Networks for NLP at EMNLP*.
- [56] Fulton Wang and Cynthia Rudin. 2015. Falling Rule Lists. In *Proceedings of AISTATS*.
- [57] Geoffrey S. Watson. 1964. Smooth Regression Analysis. *Sankhya: The Indian Journal of Statistics, Series A* 26 (1964), 359–372. Issue 4.
- [58] Kilian Q. Weinberger, John Blitzer, and Lawrence K. Saul. 2006. Distance Metric Learning for Large Margin Nearest Neighbor Classification. In *Proceedings of NIPS*.
- [59] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. <https://arxiv.org/abs/1708.07747>
- [60] Eric P. Xing, Andrew Y. Ng, Michael I. Jordan, and Stuart Russell. 2002. Distance Metric Learning, with Application to Clustering with Side-information. In *Proceedings of NIPS*.
- [61] Jiaming Xu, Peng Wang, Guanhua Tian, Bo Xu, Jun Zhao, Fangyuan Wang, and Hongwei Hao. 2015. Short Text Clustering via Convolutional Neural Networks. In *Proceedings of NAACL*.