

Bergische Universität Wuppertal

Fachbereich Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational
Mathematics (IMACM)

Preprint BUW-IMACM 18/07

P. Oberdiek, M. Rottmann, and H. Gottschalk

Classification Uncertainty of Deep Neural Networks Based on Gradient Information

June 12, 2018

<http://www.math.uni-wuppertal.de>

Classification Uncertainty of Deep Neural Networks Based on Gradient Information

Philipp Oberdiek¹, Matthias Rottmann², and Hanno Gottschalk²

¹ Technische Universität Dortmund, 44227 Dortmund, Germany
philipp.oberdiek@udo.edu

² Bergische Universität Wuppertal, 42119 Wuppertal, Germany
{hanno.gottschalk, rottmann}@uni-wuppertal.de

Abstract. We study the quantification of uncertainty of Convolutional Neural Networks (CNNs) based on gradient metrics. Unlike the classical softmax entropy, such metrics gather information from all layers of the CNN. We show for the (E)MNIST data set that for several such metrics we achieve the same meta classification accuracy – i.e. the task of classifying correctly predicted labels as correct and incorrectly predicted ones as incorrect without knowing the actual label – as for entropy thresholding. Meta classification rates for out of sample images can be increased when using entropy together with several gradient based metrics as input quantities for a meta-classifier. This proves that our gradient based metrics do not contain the same information as the entropy. We also apply meta classification to concepts not used during training: EMNIST/Omniglot letters, CIFAR10 and noise. Meta classifiers only trained on the uncertainty metrics of classes available during training usually do not perform equally well for all the unknown concepts letters, CIFAR10 and uniform noise. If we however allow the meta classifier to be trained on uncertainty metrics including some samples of some or all of the categories, meta classification for concepts remote from MNIST digits can be improved considerably.

Keywords: Deep Learning · Uncertainty Quantification · Meta Classification.

1 Introduction

In recent years deep learning has outperformed other classes of predictive models in many applications. In some of these, e.g. autonomous driving or diagnostics in medicine, the reliability of a prediction is of highest interest. In classification tasks, the thresholding on the highest softmax probability or thresholding on the entropy of the classification distributions (softmax output) are commonly used metrics to quantify a network’s classification uncertainty, see e.g. [10]. However, misclassification is oftentimes not detected by these metrics and it is also well known that these metrics can be fooled easily. Many works demonstrated how an input can be designed to fool a neural network such that it

incorrectly classifies the input with high confidence (termed adversarial examples, see e.g. [17,9,12,18]). This underlines the need for measures of uncertainty.

A basic statistical study of the performance of softmax probability thresholding on several datasets was developed in [10]. This work also assigns proper out-of-distribution candidate datasets to many common datasets. For instance a network trained on MNIST is applied to images of handwritten letters, scaled gray scale images from CIFAR10, and different types of noise. This represents a baseline for comparisons.

Using classical approaches from uncertainty quantification for modeling input uncertainty and/or model uncertainty, the detection rate of misclassifications can be improved. Using the baseline in [10], an approach named ODIN, which is based on input uncertainty, was published in [14]. This approach shows improved results compared to pure softmax probability thresholding. Uncertainty in the weights of a neural network can be modeled using Bayesian neural networks. A practically feasible approximation to Bayesian neural networks was introduced in [8], known as Monte-Carlo dropout, which also improves over classical softmax probability thresholding.

Since the softmax removes one dimension from its input by normalization, some works also perform outlier detection on the softmax input (the penultimate layer) and outperform softmax probability thresholding as well, see [2].

In this work we propose a different approach to measure uncertainty of a neural network. The idea is to use the gradient of the log-likelihood evaluated at the CNN’s prediction in place of the actual class label. A high or distorted gradient is interpreted as a sign, that, if the prediction would be true, major re-learning would be necessary for the CNN. We interpret this ‘re-learning-stress’ as uncertainty. Technically, we back-propagate the difference of the predicted class distribution and the one hot encoded predicted class (not updating any weights). During back-propagation we extract compressed representations of the occurring gradients, e.g., the norm of a gradient for a chosen layer. We study the performance of different metrics used in two meta classification tasks: separating correct and incorrect predictions and detecting in- and out-of-distribution samples.

The remainder of this work is structured as follows: First, in [section 2](#), we give a comparison with related work. Then we introduce (gradient) metrics in [section 3](#) which are studied in meta classification tasks later on. In [section 4](#) we introduce the network architecture, the experiment setup and the required theory. The setup contains the choice of data sets. We use EMNIST ([6]) digits as a known concept on which the CNN is trained and EMNIST letters, CIFAR10 images as well as different types of noise as unknown/unlearned concepts. Then we statistically investigate the separation performance of our metrics for correct vs. incorrect classifications provided by CNNs. The separation performance is studied threshold independent using AUROC and AUPR measures. This is followed by a performance study for the detection of in- and out-of-distribution samples (detection of unlearned concepts) in [section 5](#). Therefore we also combine available metrics for training and comparing different meta classifiers. In

this section meta classifiers are trained only using known concepts, i.e., EMNIST digits. Afterwards, in [section 6](#), we insert unlearned concepts (which therefore become known unknown) into the training of the meta classifiers. While the softmax baseline achieves an AUROC value of 95.83% our approach gains 0.81% in terms of AUROC and even more in terms of AUPR values.

2 Related work

The closest approaches to ours are probably [\[10\]](#) and [\[2\]](#) as they establish a self evaluation procedure for neural networks. However they only incorporate metrics for particular layers close to the networks output. None of the presented metrics are based on gradients from back-propagation. In this work we give a comprehensive study for gradient metrics applied to different layers of a neural network. Just as [\[10\]](#) and [\[2\]](#) our approach does not make use of input or model uncertainty. However these approaches, as well as our approach, are somewhat orthogonal to classical uncertainty quantification and should be potentially combinable with input uncertainty and model uncertainty, as used in [\[14\]](#) and [\[8\]](#), respectively.

3 Entropy, Softmax Baseline and Gradient Metrics

Given an input $x \in \mathbb{R}^n$, weights $w \in \mathbb{R}^p$ and class labels $y \in \mathcal{C} = \{1, \dots, q\}$, we denote the output of a neural network by $f(y|x, w) \in [0, 1]$. Using the maximum a posteriori principle (MAP), the predicted class is defined by

$$\hat{y}(x, w) := \arg \max_{y \in \mathcal{C}} f(y|x, w) \quad (1)$$

according to the Bayes decision rule [\[3\]](#), or as one hot encoded label $\hat{g}(x, w) \in \{0, 1\}^q$ with

$$\hat{g}_k(x, w) = \begin{cases} 1, & \hat{y}(x, w) = k \\ 0, & \text{else} \end{cases} \quad (2)$$

for $k = 1, \dots, q$. The entropy (also called Shannon information, [\[15\]](#))

$$E(x, w) = -\frac{1}{\log(q)} \sum_{y \in \mathcal{C}} f(y|x, w) \log(f(y|x, w)), \quad (3)$$

is a well known dispersion measure and widely used for quantifying classification uncertainty of neural networks. The softmax baseline proposed by [\[10\]](#) is calculated as

$$S(x, w) = \max_{y \in \mathcal{C}} f(y|x, w). \quad (4)$$

Given an input sample x^i with predicted class label \hat{g}^i (from [eq. \(2\)](#)) and a loss function $L = L(f(y|x^i, w), y^i)$, we can calculate the gradient of the loss function with respect to the weights $\nabla_w L = \nabla_w L(f(y|x^i, w), \hat{g}^i)$.

We apply the following metrics to this gradient:



Fig. 1. Different concepts used for our statistical experiments

- Absolute norm ($\|\nabla_w L\|_1$)
- Euclidean norm ($\|\nabla_w L\|_2$)
- Minimum ($\min(\nabla_w L)$)
- Maximum ($\max(\nabla_w L)$)
- Mean ($\text{mean}(\nabla_w L)$)
- Skewness ($\text{skew}(\nabla_w L)$)
- Kurtosis ($\text{kurt}(\nabla_w L)$)

These metrics can also be applied layerwise by restricting the gradient to those weights belonging to a single layer in the neural network.

4 Meta Classification – a Benchmark between Maximum Softmax Probability and Gradient Metrics

We performed all our statistical experiments on the EMNIST data set [6], which contains 28×28 gray scale images of 280 000 handwritten digits (0 – 9) and 411 302 handwritten letters (a – z, A – Z). We train the CNNs only on the digits, in order to test their behavior on untrained concepts. We split the EMNIST data set (after a random permutation) as follows:

- 60,000 digits (0 – 9) for training
- 20,000 digits (0 – 9) for validation
- 200,000 digits (0 – 9) for testing
- 20,000 letters (a – z, A – Z) as untrained concepts

Additionally we included the CIFAR10 library [11], shrunk and converted to gray scale, as well as 20,000 images generated from random uniform noise. All concepts can be seen in [fig. 1](#).

The architecture of the CNNs consists of three convolutional (conv) layers with 16 filters of size 3×3 each, with a stride of 1, as well as a dense layer with a 10-way softmax output. The first two conv layers are followed by 2×2 max pooling. Each layer is equipped with the leaky ReLU activation function

$$\text{LeakyReLU}(x) = \begin{cases} x, & x > 0 \\ 0.1x, & x < 0 \end{cases}, \quad (5)$$

and L^2 regularization with a regularization parameter of 10^{-3} . Additionally dropout [16] is applied after the first conv layer as well as after the dense layer. The dropout rate is 33%.

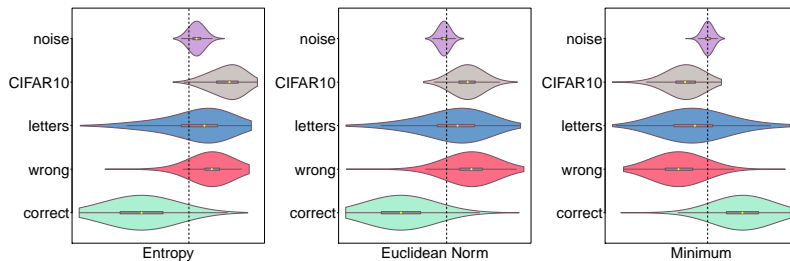


Fig. 2. Empirical distribution for entropy, euclidean norm and minimum applied to correctly predicted and incorrectly predicted digits from the test data (green and red) of one CNN. Further distributions are generated from EMNIST samples with unlearned letters (blue), CIFAR10 images (gray) and uniform noise images (purple).

The models are trained using stochastic gradient descent with a batch size of 256, momentum of 0.9 and categorical cross entropy as cost function. The initial learning rate is 0.1 and is reduced by a factor of 10 every time the average validation accuracy stagnates, until a lower limit for the learning rate of 0.001 is reached. All models were trained and evaluated using Keras [5] with Tensorflow backend [1].

Consider the EMNIST test dataset as a random sample of size $n = 200\,000$ drawn from independently identically distributed random variable X_i on some probability space $X_i : (\Omega, \mathcal{A}, P) \rightarrow \mathbb{R}^{28 \times 28}$ taking values in the space of images with resolution 28×28 . Furthermore, consider another copy X of the sampling process. We are then interested in the conditional distribution of the entropy $E(X, w^{(k)})$ as well as in the ones for the gradient metrics $\text{abs}(X, w^{(k)})$, $\text{euclid}(X, w^{(k)})$, \dots , $\text{kurt}(X, w^{(k)})$, conditioned to the event of correct classification

$$T(w) = \left\{ Y = \arg \max_{y \in \mathcal{C}} f(y|X, w) \right\} \subseteq \Omega, \quad (6)$$

where X denotes an input image and Y a class label (both random). We intend to compare these conditional distributions with the distributions conditioned to the complementary set $F(w) = \Omega \setminus T(w)$ that contains all misclassified samples.

Calculating the entropy and gradient metrics for all test samples and conditioning into the groups $T(w)$ and $F(w)$, we obtain pairs for the empirical distributions for $E(X, w)|_{T(w)}$ vs. $E(X, w)|_{F(w)}$, $\text{abs}(X, w^{(k)})|_{T(w)}$ vs. $\text{abs}(X, w^{(k)})|_{F(w)}$, \dots , $\text{kurt}(X, w^{(k)})|_{T(w)}$ vs. $\text{kurt}(X, w^{(k)})|_{F(w)}$. The results for the entropy, euclidean norm and minimum are shown in fig. 2 (green and red) (a violin plot is a density estimation of the underlying values). In what follows we define EMNISTc as the set containing all correctly classified samples of the EMNIST test set and EMNISTw as the set containing all incorrectly classified ones.

In this section, we predict correct classification and misclassification for EMNIST test data based on a simple threshold criterion, thus solving the *error and success prediction* problem, formulated in [10]. To introduce the idea of

meta classification, we are showing an exemplary threshold calculation for the entropy. For the entropy the threshold \hat{t} could be calculated as

$$\hat{t}_E = \arg \max_{t \in [0,1]} \left[\frac{1}{2} \hat{\text{cdf}}_{E|T}(t) + \frac{1}{2} \left(1 - \hat{\text{cdf}}_{E|F}(t) \right) \right], \quad (7)$$

such that the estimated classification rate is maximized for balanced test data. Here $\hat{\text{cdf}}_{E|T}(t)$ is the empirical cumulative distribution function of the entropy for the training data conditioned to $T(w)$ and evaluated at $t \in [0, 1]$. If $E(x, w) > \hat{t}_E$ for a given sample x from the test set, the classification $\hat{y}(x, w)$ is predicted to be presumably incorrect, whereas for $E(x, w) \leq \hat{t}_E$ it is classified as presumably correct. The same kind of decision rule is applied to the comparisons of the gradient metrics with their corresponding thresholds.

There are many different ways to compute suitable thresholds, but we will evaluate our results with the **Area Under the Receiver Operating Curve** (AUROC) and **Area Under the Precision Recall curve** (AUPR), which are threshold independent. By defining

$$\begin{aligned} TP &= \#\{\text{correctly predicted positive cases}\}, \\ TN &= \#\{\text{correctly predicted negative cases}\}, \\ FP &= \#\{\text{incorrectly predicted positive cases}\}, \\ FN &= \#\{\text{incorrectly predicted negative cases}\}, \end{aligned}$$

we can compute the quantities

$$\begin{aligned} R = TPR &= \frac{TP}{TP + FN} && \text{(True positive rate or Recall)}, \\ FPR &= \frac{FP}{FP + TN} && \text{(False positive rate)}, \\ P &= \frac{TP}{TP + FP} && \text{(Precision)}. \end{aligned}$$

When dealing with threshold dependent classification techniques, one calculates TPR (R), FPR and P for many different thresholds in the value range of the variable. The AUROC is the area under the receiver operating curve, which has the FPR as ordinate and the TPR as abscissa. The AUPR is the area under the precision recall curve, which has the recall as the ordinate and the precision as abscissa. For more information on these metrics see [7].

The AUPR is in general more informative for datasets with a strong imbalance in positive and negative cases and is sensitive to which class is defined as the positive case. Because of that we are computing the AUPR-In and AUPR-Out, for which the definition of a positive case is reversed. In addition the values of one variable are getting multiplied by -1 to switch between AUPR-In and AUPR-Out as in [10].

Table 1. AUROC, AUPR-In (EMNISTc as positive case) and AUPR-Out (EMNISTc as negative case) values for the threshold classification on the softmax baseline, entropy as well as selected gradient metrics. All values are in percentage and averaged over 10 differently initialized CNNs with distinct permutations of the training data.

Metric	EMNISTc / EMNISTw	EMNISTc / EMNIST letters	EMNISTc / CIFAR10	EMNISTc / uniform noise
AUROC				
Softmax Baseline	97.82	87.62	99.13	92.95
Entropy	97.74	88.44	99.24	93.52
Absolute Norm	97.77	87.22	98.19	90.66
Euclidean Norm	97.78	87.27	98.38	91.05
Minimum	97.78	87.30	98.40	90.50
Maximum	97.70	86.92	98.31	87.05
Standard Deviation	97.78	87.26	98.38	90.98
AUPR-In				
Softmax Baseline	99.97	98.39	99.98	99.31
Entropy	99.97	98.38	99.95	99.36
Absolute Norm	99.97	98.42	99.89	99.07
Euclidean Norm	99.97	98.42	99.90	99.11
Minimum	99.97	95.20	99.90	99.05
Maximum	99.97	95.03	99.89	98.67
Standard Deviation	99.97	95.04	99.90	99.11
AUPR-Out				
Softmax Baseline	39.96	59.04	77.10	40.10
Entropy	95.56	60.36	91.27	42.46
Absolute Norm	95.28	58.39	66.62	33.08
Euclidean Norm	95.30	58.27	70.81	34.03
Minimum	95.36	58.76	72.72	33.00
Maximum	95.32	55.01	74.59	26.84
Standard Deviation	95.30	58.26	70.75	33.88

The random splitting of the 280,000 digit images in training, validation and test data is repeated 10 times. In this way we train 10 CNNs that differ with respect to initial weights, training, validation and test data. We then repeat the above meta classification for each of the CNNs. With this non parametric bootstrap, we try to get as close as possible to a true sampling of the statistical law underlying the EMNIST ensemble of data and obtain results with statistic validity. [Table 1](#) (column one) gives the respective figures for the classification of the 200,000 test images. Note that we have left out the mean, skewness and kurtosis metric, as their violin plots showed, that they are not suitable for a threshold meta classifier.

[Table 1](#) shows that gradient metrics individually are performing equally well or better than the softmax probability. When comparing the threshold classification performance for metrics on different layers, most of them appear to perform better on layers that are closer to the softmax. This could be due to the fact, that early layers are detecting lower order features like edges, which may be also present in images of unknown concepts, whereas layers that are close to the softmax output are specialized for detecting concepts seen in the training set. In conclusion the softmax baseline as well as some selected gradient metrics exhibit comparable performance on the test set in the error and success prediction task.

In a next step we aggregate entropy and gradient based metrics in a more sophisticated classification technique. Therefore we choose a variety of regularized and unregularized logistic regression techniques, namely a **Generalized Linear Model** (GLM) equipped with the logit link function, the **Least Absolute Shrinkage and Selection Operator** (LASSO) with a L^1 regularization term, the ridge regression with a L^2 regularization term and finally the Elastic net with one half L^1 and one half L^2 regularization.

To include a non linear classifier we train a feed forward NN with one hidden layer containing 15 rectified linear units (ReLU) with L^2 weight decay of 10^{-3} and 2-way softmax output. The neural network is trained in the same fashion as the CNNs with stochastic gradient descent. Both groups of classifiers are trained on the EMNIST validation set. Results for the logistic regression techniques can be seen in [table 2](#) (column one) and the ones for the neural network in [table 3](#) (first row of each evaluation metric). For comparison we also include the entropy and softmax baseline in each table. The regression techniques perform equally well or better compared to the softmax baseline. This is however not true for the NN. For the logistic regression types including more features from early layers did not improve the performance, the neural network however showed improved results. This means the additional information in those layers can only be utilized by a non linear classifier.

5 Recognition of Unlearned Concepts

A (C)NN, being a statistical classifier, classifies inside the prescribed label space. In this section, we empirically test the hypothesis that test samples out of the label space will be all misclassified, however at a statistically different level of entropy or gradient metric, respectively. We test this hypothesis for three cases: First we feed the CNN with images from the EMNIST letter set and determine the entropy as well as the values for all gradient metrics for each of it. Secondly we follow the same procedure, however the inputs are gray scale CIFAR10 images coarsened to 28×28 pixels. Finally, we use uncorrelated noise that is uniformly distributed in the gray scales with the same resolution. Roughly speaking, we test empirical distributions for unlearned data that is close to the learned concept as in the case of EMNIST letters, data that represents a somewhat remote concept as in the case of CIFAR10 or, as in the case of noise, do not represent any concept at all.

We are classifying the output of a CNN on such input as incorrect label, this way we solve the *in- and out-of-distribution detection problem* from [\[10\]](#), but are still detecting misclassifications in the prescribed label space. The empirical distributions of unlearned concepts can be seen in [fig. 2](#). As we can observe, the distributions for incorrectly classified samples are in a statistical sense significantly different from those for correctly classified ones. The gradient metrics however are not able to separate the noise samples very well, but also resulting in an overall good separation of the other concepts, as for the entropy. The threshold classification evaluation metrics can be seen in [table 1](#). For the logistic

Table 2. Average AUROC, AUPR-In and AUPR-Out values for different regression types trained on the validation set and all metric features including the entropy. The values are averaged over 10 CNNs and displayed in percentage.

Metric / Regression technique	EMNISTc / EMNISTw	EMNISTc / EMNIST letters	EMNISTc / CIFAR10	EMNISTc / uniform noise
AUROC				
Softmax Baseline	97.82	87.62	99.13	92.95
Entropy	97.74	88.44	99.42	93.52
GLM	49.90	50.00	50.03	50.54
LASSO	97.58	89.33	98.70	93.90
Ridge	66.07	62.93	66.32	64.41
Elastic net	97.77	89.29	98.54	92.27
AUPR-In				
Softmax Baseline	99.97	98.39	99.98	99.31
Entropy	99.97	98.38	99.99	99.36
GLM	98.24	86.40	91.28	85.40
LASSO	99.97	98.34	99.92	99.33
Ridge	99.22	92.16	95.86	92.42
Elastic net	99.97	98.34	99.91	99.15
AUPR-Out				
Softmax Baseline	39.96	59.04	77.10	40.10
Entropy	95.56	60.36	86.07	42.46
GLM	2.23	16.47	11.47	23.54
LASSO	35.47	64.63	85.60	48.75
Ridge	3.02	20.58	14.75	21.90
Elastic net	46.18	64.57	75.52	42.53

regressions one can see that GLM and ridge are not able to separate EMNISTc from EMNISTw. Regression techniques which produce sparsity in the input features like LASSO and Elastic net are performing best. We get similar AUROC values as for the threshold classification with single metrics, but can improve between 5.59% and 8.65% over the softmax baseline in terms of AUPR-Out values for unknown concepts, showing a better generalization.

6 Meta Classification with Known Unknowns

In the previous section we trained the meta classifier on the training or validation data only. This means it has no knowledge of entropy or metric distributions for unlearned concepts, hence we followed a puristic approach treating out of distribution cases as unknown unknowns. The classification accuracy could be improved, by supplying the meta classifier with the entropy and gradient metric values of a few unlearned concepts. For this we use the same data sets as [10], namely the omniglot handwritten characters set [13], the notMNIST dataset [4] consisting of letters from different fonts, the CIFAR10 dataset [11] coarsened and converted to gray scale as well as normal and uniform noise. In order to investigate the influence of unknown concepts in the training set of the meta classifier, we used the LASSO regression and the NN introduced in section 4 and supplied them with different training sets, consisting of

- EMNIST validation set

- EMNIST validation set and 200 uniform noise images
- EMNIST validation set, 200 uniform noise images and 200 CIFAR10 images
- EMNIST validation set, 200 uniform noise images, 200 CIFAR10 images and 200 omniglot images

We are omitting the results for the LASSO here, since they are inferior to those of the NN. Including known unknowns into the training set, the NN has far

Table 3. AUROC, AUPR-In (EMNISTc is positive case) and AUPR-Out (EMNISTc is negative case) values for a NN meta classifier. "All" contains omniglot, notMNIST, CIFAR10, normal noise and uniform noise. We used 200 samples of each concept that was additionally included into the training set. The displayed values are averages over 5 differently initialized NN meta classifiers for each of the 10 CNNs trained on the EMNIST dataset. All values are in percentage.

Training set for the neural network meta classifier						
Wrong Datasets	Entropy	Softmax Baseline [10]	EMNIST validation	EMNIST validation+ uniform noise	EMNIST validation+ uniform noise+ CIFAR10	EMNIST validation+ uniform noise+ CIFAR10+ omniglot
AUROC						
EMNISTw	97.74	97.84	94.59	96.51	96.69	96.68
Omniglot	98.05	97.84	94.38	97.29	97.44	97.84
notMNIST	95.41	95.24	85.90	93.22	94.49	94.86
CIFAR10	99.24	99.03	81.19	96.27	99.12	99.09
Normal noise	94.36	94.49	56.09	98.37	98.34	98.17
Uniform noise	94.31	93.87	86.77	94.22	93.87	94.42
All	96.04	95.83	80.55	95.49	96.36	96.64
AUPR-In						
EMNISTw	99.97	99.97	99.89	99.95	99.96	99.96
Omniglot	99.84	99.82	99.04	99.73	99.75	99.80
notMNIST	99.45	99.43	95.86	98.83	99.19	99.29
CIFAR10	99.95	99.94	95.47	99.41	99.94	99.93
Normal noise	99.59	99.60	92.72	99.89	99.89	99.88
Uniform noise	99.65	99.62	98.05	99.56	99.53	99.57
All	98.66	98.59	84.98	97.72	98.53	98.71
AUPR-Out						
EMNISTw	35.83	39.94	32.95	36.02	35.98	35.33
Omniglot	83.48	80.45	74.17	80.36	81.46	83.40
notMNIST	74.86	14.59	64.57	71.53	74.91	75.13
CIFAR10	91.27	87.38	54.45	73.93	90.84	89.82
Normal noise	54.98	57.32	18.57	68.73	67.89	65.12
Uniform noise	37.97	36.63	56.66	58.56	56.59	59.53
All	89.17	88.07	75.64	89.38	91.23	91.55

better performance on the unknown concepts, even though the amount of additional training data is small. Noteworthy the validation set together with only 200 uniform noise images increases the results on the AUPR-Out values for all unknown concepts already significantly by 13.74%, even comparable to using all concepts. Together with the fact, that noise is virtually available at no cost, it is a very promising candidate for improving the generalization of the meta classifier without the need of generating labels for more datasets. The in-distribution

detection rate of correct and wrong predictions is also increased when using additional training concepts, making it only beneficial to include noise into the training set of the meta classifier. Our experiments show however that normal noise does not have such an high influence on the performance as uniform noise and is even decreasing the in-distribution meta classification performance. All in all we reach a 3.48% higher performance on the out of distribution examples compared to the softmax baseline in AUPR-Out and 0.81% in AUROC, whereas the increase in AUPR-In is marginal (0.12%).

7 Conclusion and Outlook

We introduced a new set of metrics that measures the uncertainty of deep CNNs. These metrics have a comparable performance with the widely used entropy and maximum softmax probability to meta-classify whether a certain classification proposed by the underlying CNN is presumably correct or incorrect. Here the performance is measured by AUROC, AUPR-In and AUPR-Out. Entropy and softmax probability perform equally well or slightly better than any single member of the new gradient based metrics for the detection of unknown concepts like EMNIST letters, gray scale converted CIFAR10 images and uniform noise where simple thresholding criteria are applied. Combining the gradient metrics together with entropy in a more complex meta classifier increases the ability to identify out-of-distribution examples, so that in some cases these meta classifiers outperform the baseline. Additional calibration by including a few samples of unknown concepts increases the performance significantly. Uniform noise proved to raise the overall performance, without the need of more labels. Overall the results for the classification of correct or incorrect predictions increased when the meta classifier was supplied with more distinct concepts in the training set. It seems that the higher number of uncertainty metrics helps to better hedge the correctly classified samples from the variety of out of sample classes, which would be difficult, if only one metric is available. Note that this increase in meta classification is particularly valuable, if one does not want to deteriorate the classification performance of the underlying classifier by additional classes for the known unknowns.

As future work we want to evaluate the performance and robustness of such gradient metrics on different tasks in pattern recognition. Further features could be generated by applying the metrics to activations rather than gradients. One could also investigate the possibility of generating artificial samples, labeled as incorrect, for the training set of the meta classifier in order to further improve the results.

Acknowledgement. We thank Fabian Hüger and Peter Schlicht from Volkswagen Group Research for discussion and remarks on this work.

References

1. Abadi, M., Agarwal, A., Barham, P., et al.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), <http://tensorflow.org/>, software available from tensorflow.org
2. Bendale, A., Boulton, T.E.: Towards open set deep networks. CoRR **abs/1511.06233** (2015), <http://arxiv.org/abs/1511.06233>
3. Berger, J.O.: Statistical Decision Theory and Bayesian Analysis. Springer - Verlag N.Y. (1980)
4. Bulatov, Y.: notmnist dataset (2011), yaroslavvb.blogspot.de/2011/09/notmnist-dataset.html
5. Chollet, F., et al.: Keras. <https://github.com/fchollet/keras> (2015)
6. Cohen, G., Afshar, S., Tapson, J., van Schaik, A.: Emnist: an extension of mnist to handwritten letters (2017), <http://arxiv.org/abs/1702.05373>
7. Davis, J., Goadrich, M.: The relationship between precision-recall and ROC curves. In: Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006. pp. 233–240 (2006). <https://doi.org/10.1145/1143844.1143874>, <http://doi.acm.org/10.1145/1143844.1143874>
8. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48. pp. 1050–1059. ICML'16, JMLR.org (2016), <http://dl.acm.org/citation.cfm?id=3045390.3045502>
9. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint (2014), <http://arxiv.org/abs/1412.6572>
10. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. CoRR **abs/1610.02136** (2016), <http://arxiv.org/abs/1610.02136>
11. Krizhevsky, A.: CIFAR-10 (2009), <https://www.cs.toronto.edu/~kriz/cifar.html>
12. Kurakin, A., Goodfellow, I.J., Bengio, S.: Adversarial examples in the physical world. CoRR **abs/1607.02533** (2016), <http://arxiv.org/abs/1607.02533>
13. Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction (2015)
14. Liang, S., Li, Y., Srikant, R.: Principled detection of out-of-distribution examples in neural networks. CoRR **abs/1706.02690** (2017), <http://arxiv.org/abs/1706.02690>
15. Shannon, C.E.: A Mathematical Theory of Communication, vol. 27, pp. 379–423, 623–656 (1948), <http://math.harvard.edu/~ctm/home/text/others/shannon/entropy/entropy.pdf>
16. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research **15**, 1929–1958 (2014), <http://jmlr.org/papers/v15/srivastava14a.html>
17. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I.J., Fergus, R.: Intriguing properties of neural networks. CoRR **abs/1312.6199** (2013), <http://arxiv.org/abs/1312.6199>
18. Yuan, X., He, P., Zhu, Q., Bhat, R.R., Li, X.: Adversarial examples: Attacks and defenses for deep learning. CoRR **abs/1712.07107** (2017)