

# Classification Uncertainty of Deep Neural Networks Based on Gradient Information

Philipp Oberdiek<sup>1</sup>, Matthias Rottmann<sup>2</sup>, and Hanno Gottschalk<sup>2</sup>

<sup>1</sup>Technische Universität Dortmund, 44227 Dortmund, Germany

philipp.oberdiek@udo.edu

<sup>2</sup>Bergische Universität Wuppertal, 42119 Wuppertal, Germany

{rottman, hano.gottschalk}@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 EMNIST digits data set that for several such metrics we achieve the same meta classification accuracy – i.e. the task of classifying predictions as correct or incorrect without knowing the actual label – as for entropy thresholding. We apply meta classification to unknown concepts (out-of-distribution samples) – EMNIST/Omniglot letters, CIFAR10 and noise – and demonstrate that meta classification rates for unknown concepts can be increased when using entropy together with several gradient based metrics as input quantities for a meta classifier. Meta classifiers only trained on the uncertainty metrics of known concepts, i.e. EMNIST digits, usually do not perform equally well for all unknown concepts. If we however allow the meta classifier to be trained on uncertainty metrics for some out-of-distribution samples, meta classification for concepts remote from EMNIST digits (then termed known unknowns) can be improved considerably.

## 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 classification uncertainty of neural networks, see e.g. Hendrycks et al. 2016. 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. Szegedy et al. 2013; I. J. Goodfellow et al. 2014; Kurakin et al. 2016; Yuan et al. 2017). 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 Hendrycks et al. 2016. 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 Hendrycks et al. 2016, an approach named ODIN, which is based on input uncertainty, was published in Liang et al. 2017. 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 Gal et al. 2016, 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 Bendale et al. 2015.

In this work we propose a different approach to measure uncertainty of a neural network based on gradient information. Technically, we compute the gradient of the negative log-likelihood of a single

sample during inference where the class argument in the log-likelihood is the predicted class. We then extract compressed representations of the gradients, e.g., the norm of a gradient for a chosen layer. E.g., a large norm of the 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 and study the performance of different gradient metrics used in two meta classification tasks: separating correct and incorrect predictions and detecting in- and out-of-distribution samples.

The closest approaches to ours are probably [Hendrycks et al. 2016](#) and [Bendale et al. 2015](#) as they also establish a self evaluation procedure for neural networks. However they only incorporate (non-gradient) metrics for particular layers close to the networks output while we consider gradient metrics extracted from all the layers. Just as [Hendrycks et al. 2016](#) and [Bendale et al. 2015](#) 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 [Liang et al. 2017](#) and [Gal et al. 2016](#), respectively.

The remainder of this work is structured as follows: First, in [section 2](#) we introduce (gradient) metrics, the concept of meta classification and threshold independent performance measures for meta classification, AUROC and AUPR, that are used in the experiments. In [section 3](#) we introduce the network architecture and the experiment setup containing the choice of data sets. We use EMNIST ([Cohen et al. 2017](#)) 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. This is followed by a performance study for the detection of in- and out-of-distribution samples (detection of unlearned concepts) in [section 4](#). 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 5](#), we insert unlearned concepts (which therefore become known

unknowns) 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 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  $f(y|x, w) \in [0, 1]$ . The entropy of the estimated class distribution conditioned on the input (also called Shannon information, [Shannon 1948](#))

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

is a well known dispersion measure and widely used for quantifying classification uncertainty of neural networks. In the following we will use the term entropy in the sense explained above. Note that this should not be confused with the entropy underlying the (not estimated and joint) statistical distribution of inputs and labels. The softmax baseline proposed by [Hendrycks et al. 2016](#) is calculated as

$$S(x, w) = \max_{y \in \mathcal{C}} f(y|x, w). \quad (2)$$

Using the maximum a posteriori principle (MAP), the predicted class is defined by

$$\hat{y}(x, w) := \arg \max_{y \in \mathcal{C}} f(y|x, w) \quad (3)$$

according to the Bayes decision rule [Berger 1980](#), or as one hot encoded label  $\hat{g}(x, w) \in \{0, 1\}^q$  with

$$\hat{g}_k(x, w) = \begin{cases} 1, & \hat{y}(x, w) = k \\ 0, & \text{else} \end{cases} \quad (4)$$

for  $k = 1, \dots, q$ . Given an input sample  $x^i$  with one hot label  $y^i$ , predicted class label  $\hat{g}^i$  (from [eq. \(4\)](#)) 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  $\nabla_w L = \nabla_w L(f(y|x^i, w), \hat{g}^i)$ . In our experiments we use the gradient of the negative log-likelihood at the predicted class label,

which means

$$\begin{aligned}
 L &= L(f(y|x^i, w), \hat{g}^i) \\
 &= - \sum_{y \in \mathcal{C}} \hat{g}_y^i \log(f(y|x^i, w)) \\
 &= - \log(f(\hat{y}|x^i, w)) .
 \end{aligned} \tag{5}$$

We apply the following metrics to this gradient:

- 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 either be applied layerwise by restricting the gradient to those weights belonging to a single layer in the neural network or to the whole gradient on all layers.

The metrics can be sampled over the input  $X$  and conditioned to the event of either correct or incorrect classification. Let  $T(w)$  and  $F(w)$  denote the subset of correctly and incorrectly classified samples for the network  $f(y|x, w)$ , respectively. Given a metric  $M$  (e.g. the entropy  $E$  or any gradient based one), the two conditioned distributions  $M(X, w)|_{T(w)}$  and  $M(X, w)|_{F(w)}$  are further investigated. For a threshold  $t$ , we measure  $P(M(X, w) < t | T(w))$  and  $P(M(X, w) \geq t | F(w))$  by sampling  $X$ . If both probabilities are high,  $t$  gives a good separation between correctly and incorrectly classified samples. This concept can be transferred to the detection of out-of-distribution samples by defining these as incorrectly classified. We term this procedure (classifying  $M(X, w) < t$  vs.  $M(X, w) \geq t$ ) *meta classification*.

Since there are many possible ways to compute thresholds  $t$ , we compute our results threshold independent by using **Area Under the Receiver Operating Curve** (AUROC) and **Area Under the Precision Recall curve** (AUPR). For any chosen threshold  $t$  we define

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

and 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 performance measures see [Davis et al. 2006](#).

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 multiplied by  $-1$  to switch between AUPR-In and AUPR-Out as in [Hendrycks et al. 2016](#).

### 3 Meta Classification – a Benchmark between Softmax Probability and Gradient Metrics

We perform all our statistical experiments on the EMNIST data set [Cohen et al. 2017](#), which contains  $28 \times 28$  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 [Krizhevsky 2009](#), 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  $3 \times 3$  each, with a stride of 1, as well as a dense layer with a 10-way softmax output. Each of the first two conv layers are equipped with leaky ReLU activations

$$\text{LeakyReLU}(x) = \begin{cases} x, & x > 0 \\ 0.1x, & x < 0 \end{cases} \quad (6)$$

and followed by  $2 \times 2$  max pooling. We employ  $L^2$  regularization with a regularization parameter of  $10^{-3}$ . Additionally, dropout [Srivastava et al. 2014](#) is applied after the first and third conv layer. The dropout rate is 33%.

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 [Chollet et al. 2015](#) with Tensorflow backend [Abadi et al. 2015](#). Note, that the parameters were chosen from experience and not tuned to any extent. The goal is not to achieve a high accuracy, but to detect the uncertainty of a neural network reliably.

In this section, we study the performance of gradient metrics, the softmax baseline and the entropy in terms of AUROC and AUPR for EMNIST test data, thus considering the *error and success prediction* problem, formulated in [Hendrycks et al. 2016](#).

First of all we demonstrate that gradient metrics are indeed able to provide good separations. Results for the entropy, euclidean norm and minimum are shown in [fig. 2](#) (green and red). 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.

In what follows we define EMNIST<sub>c</sub> as the set containing all correctly classified samples of the EMNIST test set and EMNIST<sub>w</sub> as the set containing all incorrectly classified ones. From now on we resample the data splitting and use ensembles of CNNs. More precisely, the random splitting of the 280,000 digit images in training, validation and test data is repeated 10 times and we train one CNN for each splitting. 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](#) shows that the softmax baseline as well as some selected gradient metrics exhibit comparable performance on the test set in the error and success prediction task. Column one corresponds to the empirical distributions depicted in [fig. 2](#) for 200,000 test images.

In a next step we aggregate entropy and all gradient based metrics (evaluated on the gradient of each layer in the CNN) 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 and a regularization parameter  $\lambda_1 = 1$ , the ridge regression with a  $L^2$  regularization term and a regularization parameter of  $\lambda_2 = 1$  and finally the Elastic net with one half  $L^1$  and one half  $L^2$  regularization, which means  $\lambda_1 = \lambda_2 = 0.5$ . For details about these methods, cf. [Hastie et al. 2001](#).

To include a non linear classifier we train a feed forward NN with one hidden layer containing 15 rectified linear units (ReLUs) 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



Figure 1: Different concepts used for our statistical experiments

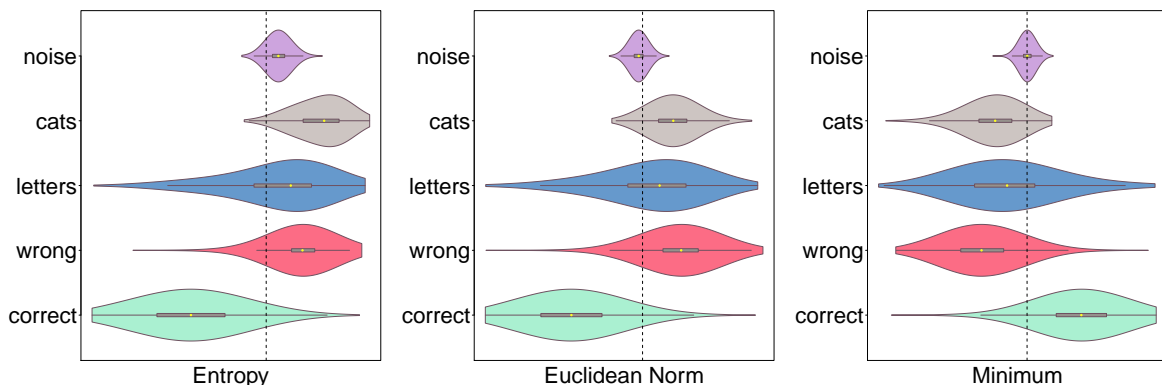


Figure 2: Empirical distribution for entropy, euclidean norm and minimum applied to correctly predicted and incorrectly predicted digits from the test data (green and red) of one CNN. Further distributions are generated from EMNIST samples with unlearned letters (blue), CIFAR10 images (gray) and uniform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

classifiers are trained on the EMNIST validation set. Results for the logistic regression techniques can be seen in [table 2](#) (column one) and those 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.

#### 4 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  $28 \times 28$  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 [Hendrycks et al. 2016](#), but are still detecting misclassifications in the prescribed label space. The



Metric	EMNIST <sub>c</sub> / EMNIST <sub>w</sub>	EMNIST <sub>c</sub> / EMNIST letters	EMNIST <sub>c</sub> / CIFAR10	EMNIST <sub>c</sub> / uniform noise
<b>AUROC</b>				
Softmax Baseline	<b>97.82</b>	87.62	99.13	92.95
Entropy	97.74	<b>88.44</b>	<b>99.42</b>	<b>93.52</b>
Absolute Norm	97.77	87.22	98.07	90.66
Euclidean Norm	97.78	87.27	98.34	91.05
Minimum	97.78	87.30	98.32	90.50
Maximum	97.70	86.92	98.34	87.05
Standard Deviation	97.78	87.26	98.34	90.98
<b>AUPR-In</b>				
Softmax Baseline	<b>99.97</b>	98.39	99.98	99.31
Entropy	<b>99.97</b>	98.38	<b>99.99</b>	<b>99.36</b>
Absolute Norm	<b>99.97</b>	<b>98.42</b>	99.95	99.07
Euclidean Norm	<b>99.97</b>	<b>98.42</b>	99.96	99.11
Minimum	<b>99.97</b>	95.20	99.96	99.05
Maximum	<b>99.97</b>	95.03	99.96	98.67
Standard Deviation	<b>99.97</b>	95.04	99.96	99.11
<b>AUPR-Out</b>				
Softmax Baseline	39.96	59.04	77.10	40.10
Entropy	<b>95.56</b>	<b>60.36</b>	<b>86.07</b>	<b>42.46</b>
Absolute Norm	95.28	58.39	41.70	33.08
Euclidean Norm	95.30	58.27	47.82	34.03
Minimum	95.36	58.76	49.32	33.00
Maximum	95.32	55.01	55.69	26.84
Standard Deviation	95.30	58.26	47.72	33.88

Table 1: AUROC, AUPR-In (EMNIST<sub>c</sub> as positive case) and AUPR-Out (EMNIST<sub>c</sub> 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 splittings of the training data. Values in brackets are the standard deviation of the mean in percentage. To get the standard deviation within the sample, multiply by  $\sqrt{10}$ .

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 regression results in [table 2](#) one can see that the GLM is inferior to the other methods. Regression techniques with a regularization term like LASSO, Ridge and Elastic net are performing best. We get similar AUROC values as for the threshold classification with single

metrics, but can improve between 5% and 14.08% over the softmax baseline in terms of AUPR-Out values for unknown concepts, showing a better generalization.

## 5 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 extending the training set of the meta classifier with the entropy

Metric / Regression technique	EMNIST <sub>c</sub> / EMNIST <sub>w</sub>	EMNIST <sub>c</sub> / EMNIST letters	EMNIST <sub>c</sub> / CIFAR10	EMNIST <sub>c</sub> / uniform noise
<b>AUROC</b>				
Softmax Baseline	<b>97.82</b> (0.03)	87.62 (0.20)	99.13 (0.11)	92.95 (1.86)
Entropy	97.74 (0.04)	88.44 (0.21)	<b>99.42</b> (0.10)	93.52 (1.83)
GLM	94.76 (0.70)	85.94 (0.46)	80.26 (5.46)	89.41 (2.90)
LASSO	97.75 (0.03)	<b>89.34</b> (0.17)	99.23 (0.03)	93.86 (1.04)
Ridge	97.59 (0.03)	88.63 (0.11)	98.93 (0.02)	<b>94.08</b> (0.67)
Elastic net	97.79 (0.06)	89.27 (0.24)	98.82 (0.06)	93.47 (0.67)
<b>AUPR-In</b>				
Softmax Baseline	<b>99.97</b> (0.00)	<b>98.39</b> (0.03)	99.98 (0.00)	99.31 (0.19)
Entropy	99.97 (0.00)	98.38 (0.04)	<b>99.99</b> (0.00)	<b>99.36</b> (0.19)
GLM	99.81 (0.05)	96.80 (0.21)	95.51 (1.15)	97.81 (0.84)
LASSO	99.97 (0.00)	98.30 (0.06)	99.95 (0.00)	99.33 (0.12)
Ridge	99.97 (0.00)	97.86 (0.04)	99.93 (0.00)	99.36 (0.08)
Elastic net	99.97 (0.00)	98.26 (0.09)	99.92 (0.00)	99.29 (0.08)
<b>AUPR-Out</b>				
Softmax Baseline	39.96 (0.57)	59.04 (0.37)	77.10 (1.90)	40.10 (4.87)
Entropy	<b>95.56</b> (0.05)	60.36 (0.42)	86.07 (1.76)	42.46 (5.38)
GLM	31.27 (0.79)	57.72 (0.74)	62.90 (6.77)	46.43 (5.24)
LASSO	36.27 (0.32)	<b>64.04</b> (0.26)	<b>91.18</b> (0.44)	<b>48.38</b> (3.12)
Ridge	38.17 (0.34)	61.92 (0.18)	82.95 (0.61)	47.30 (2.20)
Elastic net	38.71 (0.65)	63.43 (0.62)	79.56 (1.76)	45.03 (1.92)

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 but excluding the softmax baseline. The values are averaged over 10 CNNs and displayed in percentage. The values in brackets are the standard deviations of the mean in percentage. To get the standard deviation within the sample, multiply by  $\sqrt{10}$ .

and gradient metric values of a few unlearned concepts and labeling them as false, i.e., incorrectly predicted. As in the previous sections we then train meta classifiers on the metrics. For this we use the same data sets as [Hendrycks et al. 2016](#), namely the omniglot handwritten characters set [Lake et al. 2015](#), the notMNIST dataset [Bulatov 2011](#) consisting of letters from different fonts, the CIFAR10 dataset [Krizhevsky 2009](#) 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 3](#) 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 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

Training set for the neural network meta classifier						
Wrong Datasets	Entropy	Softmax Baseline Hendrycks et al. 2016	EMNIST validation	EMNIST validation+ uniform noise	EMNIST validation+ uniform noise+ CIFAR10	EMNIST validation+ uniform noise+ CIFAR10+ omniglot
<b>AUROC</b>						
EMNISTw	97.74	<b>97.84</b>	94.59	96.51	96.69	96.68
Omniglot	<b>98.05</b>	97.84	94.38	97.29	97.44	97.84
notMNIST	<b>95.41</b>	95.24	85.90	93.22	94.49	94.86
CIFAR10	<b>99.24</b>	99.03	81.19	96.27	99.12	99.09
Normal noise	94.36	94.49	56.09	<b>98.37</b>	98.34	98.17
Uniform noise	94.31	93.87	86.77	94.22	93.87	<b>94.42</b>
All	96.04	95.83	80.55	95.49	96.36	<b>96.64</b>
<b>AUPR-In</b>						
EMNISTw	99.97	<b>99.97</b>	99.89	99.95	99.96	99.96
Omniglot	<b>99.84</b>	99.82	99.04	99.73	99.75	99.80
notMNIST	<b>99.45</b>	99.43	95.86	98.83	99.19	99.29
CIFAR10	<b>99.95</b>	99.94	95.47	99.41	99.94	99.93
Normal noise	99.59	99.60	92.72	<b>99.89</b>	99.89	99.88
Uniform noise	99.65	99.62	98.05	99.56	99.53	<b>99.57</b>
All	98.66	98.59	84.98	97.72	98.53	<b>98.71</b>
<b>AUPR-Out</b>						
EMNISTw	35.83	<b>39.94</b>	32.95	36.02	35.98	35.33
Omniglot	<b>83.48</b>	80.45	74.17	80.36	81.46	83.40
notMNIST	74.86	14.59	64.57	71.53	74.91	<b>75.13</b>
CIFAR10	<b>91.27</b>	87.38	54.45	73.93	90.84	89.82
Normal noise	54.98	57.32	18.57	<b>68.73</b>	67.89	65.12
Uniform noise	37.97	36.63	56.66	58.56	56.59	<b>59.53</b>
All	89.17	88.07	75.64	89.38	91.23	<b>91.55</b>

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 supplied features are all gradient based metrics as well as the entropy. 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 and the values in brackets are the standard deviations of the mean in percentage. To get the mean within the sample multiply by  $\sqrt{10}$ .

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 a 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%).



## 6 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. But still, our new metrics allow contributions of different layers and weights to the total uncertainty. 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. We also thank

the referees for their comments and criticism helping us to improve the paper.

## References

- Abadi, Martin, Ashish Agarwal, Paul Barham, et al. (2015). *TensorFlow: large-scale machine learning on heterogeneous systems*. Software available from tensorflow.org.
- Bendale, Abhijit and Terrance E. Boult (2015). “Towards open set deep networks”. *Corr abs/1511.06233*.
- Berger, James O. (1980). *Statistical Decision Theory and Bayesian Analysis*. Springer - Verlag N.Y.
- Bulatov, Yaroslav (2011). *Notmnist dataset*. URL: [yaroslavvb.blogspot.de/2011/09/notmnist-dataset.html](http://yaroslavvb.blogspot.de/2011/09/notmnist-dataset.html) (visited on 05/08/2018).
- Chollet, François et al. (2015). *Keras*. <https://github.com/fchollet/keras>.
- Cohen, G. et al. (2017). “Emnist: an extension of mnist to handwritten letters”.
- Cowell, F. (2011). *Measuring Inequality*. Oxford Univ. Press.
- Davis, Jesse and Mark Goadrich (2006). “The relationship between precision-recall and ROC curves”. *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, pp. 233–240.
- Efron, Bradley and Robert J. Tibshirani (1993). *An Introduction to the Bootstrap*. Chapman & Hall.
- Freund, Yoav and Robert E. Schapire (1995). “A decision-theoretic generalization of on-line learning and an application to boosting”. *Proceedings of the Second European Conference on Computational Learning Theory*. EuroCOLT '95. London, UK, UK: Springer-Verlag, pp. 23–37.
- Gal, Yarin and Zoubin Ghahramani (2016). “Dropout as a bayesian approximation: representing model uncertainty in deep learning”. *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*. ICML'16. New York, NY, USA: JMLR.org, pp. 1050–1059.
- Goodfellow, Ian J., Jonathon Shlens, and Christian Szegedy (2014). “Explaining and harnessing adversarial examples”. *Arxiv preprint*.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. MIT Press.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2001). *The Elements of Statistical Learning*. Springer Series in Statistics. New York, NY, USA: Springer New York Inc.
- Hendrycks, Dan and Kevin Gimpel (2016). “A baseline for detecting misclassified and out-of-distribution examples in neural networks”. *Corr abs/1610.02136*.
- Ioffe, Sergey and Christian Szegedy (2015). “Batch normalization: accelerating deep network training by reducing internal covariate shift”. *Corr abs/1502.03167*.
- Jain, Lalit P., Walter J. Scheirer, and Terrance E. Boult (2014). “Multi-class open set recognition using probability of inclusion”. *The European Conference on Computer Vision (ECCV)*.
- Krizhevsky, Alex (2009). *CIFAR-10*.
- Kurakin, Alexey, Ian J. Goodfellow, and Samy Bengio (2016). “Adversarial examples in the physical world”. *Corr abs/1607.02533*.
- Lake, B. M., R. Salakhutdinov, and J. B. Tenenbaum (2015). “Human-level concept learning through probabilistic program induction”.
- LeCun, Y. et al. (1998). “Gradient-based learning applied to document recognition”. *Proceedings of the IEEE* 86.11, pp. 2278–2324.
- Liang, Shiyu, Yixuan Li, and R. Srikant (2017). “Principled detection of out-of-distribution examples in neural networks”. *Corr abs/1706.02690*.
- Nguyen, Anh Mai, Jason Yosinski, and Jeff Clune (2014). “Deep neural networks are easily fooled: high confidence predictions for unrecognizable images”. *Corr abs/1412.1897*.
- Scheirer, Walter J. et al. (2011). “Meta-recognition: the theory and practice of recognition score analysis”. *Ieee transactions on pattern analysis and machine intelligence* 33, pp. 1689–1695.
- Shannon, C. E. (1948). “A mathematical theory of communication”. In: *The Bell System Technical Journal*. Vol. 27, pp. 379–423, 623–656.
- Srivastava, Nitish et al. (2014). “Dropout: a simple way to prevent neural networks from overfitting”. *Journal of machine learning research* 15, pp. 1929–1958.
- Subramanya, Akshayvarun, Suraj Srinivas, and R. Venkatesh Babu (2017). “Confidence estimation in deep neural networks via density modelling”. *Arxiv preprint*.
- Szegedy, Christian et al. (2013). “Intriguing properties of neural networks”. *Corr abs/1312.6199*.

- Xu, Bing et al. (2015). “Empirical evaluation of rectified activations in convolutional network”. *Corr abs/1505.00853*.
- Yuan, Xiaoyong et al. (2017). “Adversarial examples: attacks and defenses for deep learning”. *Corr abs/1712.07107*.
- Zheng, Stephan et al. (2016). “Improving the robustness of deep neural networks via stability training”. *Corr abs/1604.04326*.