



# SLOVA: Uncertainty estimation using single label one-vs-all classifier

Bartosz Wójcik<sup>a,\*</sup>, Jacek Grela<sup>b</sup>, Marek Śmieja<sup>a</sup>, Krzysztof Misztal<sup>a</sup>, Jacek Tabor<sup>a</sup>

<sup>a</sup> Faculty of Mathematics and Computer Science, Jagiellonian University Łojasiewicza 6, 30-348 Kraków, Poland

<sup>b</sup> Faculty of Physics, Astronomy and Applied Computer Science, Jagiellonian University Łojasiewicza 11, 30-348 Kraków, Poland



## ARTICLE INFO

### Article history:

Received 16 October 2021  
Received in revised form 11 May 2022  
Accepted 19 June 2022  
Available online 30 June 2022

### Keywords:

Uncertainty  
One-vs-all  
Calibration  
Distribution shift  
Out-of-distribution

## ABSTRACT

Deep neural networks present impressive performance, yet they cannot reliably estimate their predictive confidence, limiting their applicability in high-risk domains. We show that applying a multi-label one-vs-all loss reveals classification ambiguity and reduces model overconfidence. The introduced SLOVA (Single Label One-Vs-All) model redefines typical one-vs-all predictive probabilities to a single label situation, where only one class is the correct answer. The proposed classifier is confident only if a single class has a high probability and other probabilities are negligible. Unlike the typical softmax function, SLOVA naturally detects out-of-distribution samples if the probabilities of all other classes are small. The model is additionally fine-tuned with exponential calibration, which allows us to precisely align the confidence score with model accuracy. We verify our approach on three tasks. First, we demonstrate that SLOVA is competitive with the state-of-the-art on in-distribution calibration. Second, the performance of SLOVA is robust under dataset shifts. Finally, our approach performs extremely well in the detection of out-of-distribution samples. Consequently, SLOVA is a tool that can be used in various applications where uncertainty modeling is required.

© 2022 Elsevier B.V. All rights reserved.

## 1. Introduction

Deep learning models frequently outperform human capabilities in typical computer vision or natural language processing tasks. Despite their impressive performance, neural networks tend to make overconfident decisions, which limits their applicability in high-risk fields such as medical diagnosis [1], autonomous vehicle control [2], or the financial and legal sector [3]. In other words, deep learning models often cannot correctly quantify predictive uncertainty.

We focus on three aspects that are critical in predictive uncertainty estimation. The first one is the *in-distribution calibration*, which says that the probability associated with the predicted class label should reflect its confidence. Second, machine learning models should be *robust under dataset shifts* (small data distortion) in terms of accuracy and confidence. Finally, the models should make low confidence predictions on *out-of-distribution (OOD) data*. While all of these issues are closely connected, there are not many models that deal with all of them simultaneously.

In this paper, we propose a simple, yet remarkably effective approach, a remedy for all of the described problems. Our idea relies on extending and calibrating the multi-label One-Vs-All


(OVA) classifier [4] to the single-label case. In contrast to a typical softmax classifier designed to rank classes rather than estimate confidence, OVA models use an individual scoring function for each class. If the model assigns high scores to more than one class or returns low predictions for all classes, it informs us about its uncertainty. Although OVA has already been used for uncertainty estimation [5,6], confidence was always defined as the maximal class probability, which results in high confidence even if two or more classes are very likely.

Following this reasoning, we construct a post hoc model, SLOVA (**S**ingle-**L**abel-**O**ne-**V**s-**A**ll), which aggregates the predictive probabilities of a pretrained OVA model. It returns high confidence if and only if a single class has a high probability and other probabilities are negligible. See Fig. 1 for a comparison with the typical OVA confidence in the two-class classification problem. We prove that SLOVA is guaranteed to return low confidence for out-of-distribution samples; see Theorem 3.1 and Fig. 2. Finally, to precisely align the confidence score with the model accuracy, we introduce exponential calibration. This transformation is capable of approximating any monotonic function [7] and, in consequence, fits perfectly for uncertainty estimation of in-distribution and shifted data.


Since the performance of OVA on a single label problem is comparable to the softmax model, SLOVA does not lead to a drop in accuracy. OVA is trained using standard multi-label binary cross-entropy and thus does not require any modification of the neural network architecture. The calculation of confidence scores

\* Corresponding author.

E-mail addresses: [Bartosz.Wojcik@ii.uj.edu.pl](mailto:Bartosz.Wojcik@ii.uj.edu.pl) (B. Wójcik), [grela@th.if.uj.edu.pl](mailto:grela@th.if.uj.edu.pl) (J. Grela), [marek.smieja@uj.edu.pl](mailto:marek.smieja@uj.edu.pl) (M. Śmieja), [krzysztof.misztal@uj.edu.pl](mailto:krzysztof.misztal@uj.edu.pl) (K. Misztal), [jacek.tabor@uj.edu.pl](mailto:jacek.tabor@uj.edu.pl) (J. Tabor).



**Fig. 1.** Confidence score map of OVA and SLOVA for a two-class classification problem. The axes  $p_1, p_2$  represent the likelihood of the data-point belonging to each class while the colors encode the score. Since OVA connects confidence with the maximal class probability, it may return high confidence if two classes are equally likely. In contrast, SLOVA gives high confidence if exactly one class has a high probability.



**Fig. 2.** Confidence score across the plane spanned by 3 randomly selected test CIFAR-10 samples, parametrized by two variables  $\alpha, \beta$ . These parameters are picked so that the plane passes through a single sample point at the center of each plot, resulting in a high confidence score at that point. The color indicates the confidence score of OVA (left) and SLOVA (right) averaged over 100 sample triplets. We see that SLOVA returns low confidence on OOD samples (far from the image center), while OVA predictions are close to 1.

and the calibration of SLOVA is performed as a post-processing step, so the training procedure remains precisely the same; see Fig. 3 for the model diagram. The implementation and usage of the proposed technique can be done with minimal effort.

We conduct extensive experiments on computer vision tasks following recent benchmarks [8,9]. We demonstrate that SLOVA is competitive with the state-of-the-art on in-distribution calibration, robustness under dataset shifts, and detection of out-of-distribution data. An additional ablation study shows the impact of the aggregation scheme and model calibration on overall performance.

Our contribution can be summarized as follows:

- We introduce SLOVA – a post hoc method for uncertainty modeling, which can be applied at test time to a OVA classifier. Our theoretical result shows that the confidence score of SLOVA is close to 0 for samples that are sufficiently far from in-distribution data.
- We construct a novel and powerful exponential calibration to align SLOVA confidence with model accuracy. The introduced technique allows us to obtain more reliable confidence estimates on in-distribution data.
- Experimental results demonstrate that SLOVA obtains performance comparable to state-of-the-art methods on three fundamental uncertainty tasks: in-distribution calibration, robustness under dataset shift, and uncertainty prediction on OOD data.

The paper is organized as follows. First, in Section 2 we describe the relevant literature. Then, in Section 3, we discuss the

main features of the proposed model. We introduce the SLOVA confidence score, explain its application to out-of-distribution tasks, and describe the exponential calibration. In Section 4 we discuss SLOVA performance in three experiments – the in-distribution classification, robustness under dataset shift, and the identification of out-of-distribution data. The experiments contain a discussion of the results, a statistical analysis of the performance of the SLOVA model, and an ablation study that focuses on identifying the most significant elements of the model. The work is concluded in Section 5 with a discussion and future work recommendations.

## 2. Related work

The topic of uncertainty estimation is extensive, with a vast amount of published literature. A complete survey on uncertainty modeling in neural networks can be found in [10]. Below, we recall the papers that are the most closely related to our approach. In particular, we focus on the Bayesian approach as the most theoretically sound but not scalable, the role of calibration methods, alternative approaches to out-of-distribution detection, and the relevance of OVA models.

### 2.1. Bayesian methods

One group of methods uses the Bayesian framework, which naturally expresses the uncertainty of the prediction [11]. The main idea relies on estimating the probability distribution on the weights of the neural network, making the model robust to perturbations. Since exact Bayesian inference on the weights of a neural network is intractable, several approximations have been introduced [12]. Although theoretically attractive, these methods do not scale easily to larger models [13]. Gal and Ghahramani [14] propose using dropout as an approximation of Bayesian inference to obtain uncertainty estimates. In contrast to Bayesian neural networks, this approach achieves a similar effect at no additional cost. Maddox et al. [15] improve this even further by approximating the posterior with averaged SGD (stochastic gradient descent) iterates. More precisely, by making use of the information contained in the SGD trajectory, this model approximates the posterior distribution over the weights of the neural network in the form of a Gaussian distribution. Although the above Bayesian approaches obtain good results, they require significant modifications to the training procedure and are computationally expensive compared to standard (non-Bayesian) neural networks.

Lakshminarayanan et al. [16] proposed an alternative to Bayesian methods, which is based on creating an ensemble (committee) of neural networks. If the base models are sufficiently diverse, their aggregated predictions significantly improve the uncertainty estimates. The method can be additionally improved using adversarial training. An approximate Bayesian version of ensembling was introduced by Pearce et al. [17]. Although the ensemble methods still achieve state-of-the-art results in most benchmarks [8,18], it should be noted that training and evaluating many models is also computationally demanding.

### 2.2. Calibration methods

It is possible to use a standard model and then add a post-processing step to calibrate the outputs (which aligns predictive probabilities with model accuracy). Gao et al. [19] show that modern neural networks are highly miscalibrated and explore multiple calibration methods. One class of methods considered there learns a logit transformation applied before they are fed to the softmax function. Temperature scaling, where a single scalar

multiplicative factor is trained, stands out among multi-class calibration methods evaluated for its results, surpassing both vector scaling and matrix scaling. Kull et al. [20] introduce a method to directly calibrate the softmax outputs instead. Zhang et al. [21] propose to create an ensemble of calibration transformations of the same type and additionally suggest composing it together with a non-parametric calibration method. It is also possible to use a function approximated by a neural network as a parametric calibration method, provided that it preserves the internal order of the predictions [22]. Finally, Gaussian processes were also examined as calibration methods [23,24].

Another recent approach to uncertainty estimation relies on replacing the conventional cross-entropy function by focal loss [25]. The idea behind focal loss is to direct the network's attention to samples for which it is currently predicting a low probability for the correct class. Seo et al. [26] enhance the typical cross-entropy loss between predictions and ground truth with the cross entropy between the predictions and the uniform distribution, forcing the network to construct as uniform distribution as possible. Kumar et al. [27] design a kernel-based function, which can be optimized during training and serves as a surrogate for the calibration error.

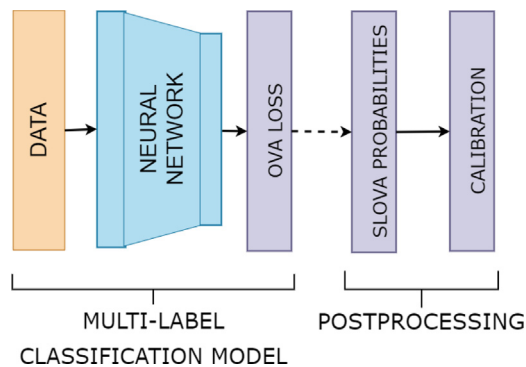
### 2.3. Out-of-distribution detection

Although calibration allows for aligning predictive probabilities with model accuracy on in-distribution data, this technique may not be adequate for detecting OOD data [28]. It was shown that out-of-distribution examples can have higher maximum softmax probabilities than in-distribution samples [29]. Multiple methods attempt to correct that by using postprocessing to adjust softmax outputs [30,31]. Alternative approaches integrate generative models, such as Generative Adversarial Networks (GANs) or Variational Autoencoders (VAEs), into the training procedure to discriminate in-distribution from OOD data [32,33]. It was later argued that generative models might return high confidence for inputs outside of the class they are supposed to model [34,35]. To avoid uncertain decisions, a rejection option was introduced into the classifier [36,37]. In the case of images, the authors of [9] enforce low confidence on OOD inputs with a procedure similar to adversarial training. These authors additionally prove that ReLU networks with softmax output produce high-confidence predictions far away from the training data. A detailed survey of OOD detection methods can be found in [38].

### 2.4. One-vs-all models

One-vs-all classification models are promising in estimating uncertainty prediction because they naturally encode the "none of the above" class. In most cases, the confidence score of an OVA classifier is calculated as the maximum class probability [5,39,40]. Our paper shows that inspecting the probability of a single class is not sufficient for quantifying uncertainty because a high probability of at least two classes leads to overall high confidence of OVA. To reduce this negative effect, the authors of [6] build an ensemble of OVA and softmax models. As in a typical ensemble approach, the confidence is high only if both models agree on the predictions. OVA classifier can also be combined with distance-based logits [5] which take the density of data into account and, in consequence, work better for detecting OOD samples [41].

In contrast to these works, we combine the predictive probabilities of all classes to define the confidence score. The proposed confidence score represents the probability that the model returns a given class and does not return other classes. Consequently, the confidence is high if and only if one class has a high probability and the others are close to zero. To deal with in-distribution samples, we additionally introduce exponential calibration, which is capable of approximating any monotonic function.



**Fig. 3.** SLOVA is a post hoc method consisting of a OVA classifier of choice (for example a pretrained neural network) and a post-processing step where the OVA predictive probabilities are transformed to SLOVA calibrated probabilities. The SLOVA probabilities reflect the confidence of in-distribution and out-of-distribution data.

## 3. Model description

In this section, we provide a detailed description of the SLOVA method. First, we formulate the classification task and explain the issues that arise when the distribution of the dataset changes. Then, we briefly describe the building blocks of the SLOVA approach. The next subsection demonstrates how the proposed method outputs desirable confidence in out-of-distribution datasets. The last subsection describes exponential calibration as the final postprocessing step that aims to correct both the overconfidence of the classifier and the underestimation arising from a multiplicative form of the SLOVA confidence score.

### 3.1. Problem statement

We consider a multi-class classification problem, in which every example  $x \in \mathbb{R}^D$  is associated with a class  $y \in \{1, \dots, K\}$ . Given the input  $x$ , we use a neural network to model the predictive distribution  $p_\theta(y|x)$  over the classes, where  $\theta$  are the parameters of the neural network. Throughout the paper, we drop the index  $\theta$  to simplify the notation.

In practice, the neural network is trained on a subset  $X_{in} \subset \mathbb{R}^D$ , which represents the in-distribution (domain) data. In consequence, during training the model learns the in-domain distribution  $p(y|x, x \in X_{in})$ , which might not generalize well to the out-of-distribution (OOD) images  $X_{out}$ . Even if the model is evaluated on the in-distribution data, we frequently observe overconfidence, which means that the model is not well calibrated. In other words, the confidence of the model related to the predictive distribution significantly exceeds its accuracy.

We are interested in constructing a model whose predictive distribution is well calibrated to the model accuracy. This distribution has to be robust to dataset (covariate) shifts. Consequently, the predictive distribution should translate into a confidence score that indicates whether the model can make the correct decisions. Such a property allows one to deal with OOD data.

### 3.2. Model overview

To achieve these goals, we propose the SLOVA model shown schematically in Fig. 3. It is a two-stage post hoc method consisting of: (a) a One-Vs-All (OVA) classifier and (b) a postprocessing step that transforms OVA into SLOVA probabilities.

Firstly, we employ the OVA classification model, in which the predictive distribution of a given class is independent of other classes. OVA obtains comparable accuracy to softmax models but

is much more flexible and can be used in multi-label classification problems, where the example can be labeled with more than one class. Since we are working in a single-label situation (only one class is correct), we redefine the predictive probabilities of OVA. Secondly, SLOVA transforms the probabilities by compensating for values of other classes, reducing certainty when the model returns multiple classes or does not return any class (as depicted in Fig. 1). In Theorem 3.1, we prove that SLOVA improves the confidence score on OOD data.

While the introduced upgrade allows us to return low confidence for OOD data, we additionally need to calibrate the returned probabilities to match the model confidence with its accuracy. For this purpose, we introduce the exponential calibration function, whose form suits well the multiplicative form of SLOVA predictive probabilities.

The following parts describe the details of the proposed model.

### 3.3. Multi-label One-Vs-All classifier

Let us recall that the predictive probabilities in OVA are described by  $K$  sigmoid functions:

$$p(k|x) = \frac{1}{1 + \exp(-f_k(x))},$$

where  $f(x) = [f_1(x), \dots, f_K(x)]$  represents the network embedding of data point  $x \in \mathbb{R}^D$  into  $K$ -dimensional logit space. The  $k$ th sigmoid defines an individual Bernoulli distribution representing the probability that  $x$  belongs to the  $k$ th class. Since OVA allows for encoding more than one positive label for a given example, it is commonly used in multi-label problems. In the case of a typical single-label classification, where exactly one label is correct, OVA is trained by minimizing the standard cross-entropy function:

$$\begin{aligned} L(x, y) &= -\log \left( p(y|x) \prod_{k \neq y} (1 - p(k|x)) \right) \\ &= -\log p(y|x) - \sum_{k \neq y} \log(1 - p(k|x)), \end{aligned} \quad (1)$$

for input  $(x, y)$ .

The confidence score of the softmax model (or other single-label models) is usually defined as the maximal class probability [19]. Following this idea directly, previous approaches to using OVA for uncertainty estimation [5,40] define OVA confidence by

$$\text{conf}_{\text{OVA}}(x) = \max_{k=1, \dots, K} p(k|x). \quad (2)$$

Observe, however, that the above formula does not take into account the fact that precisely one label is correct for every example. As mentioned, OVA suits well to multi-label situations and can return a high probability for more than one class. If more than one class has a high probability in single-label classification, then the model confidence should be low because the model is unsure which class to pick. Nevertheless, the formula (2) ignores such a situation and returns high confidence if at least one class is highly probable.

### 3.4. SLOVA confidence

Motivated by the above reasoning, we calculate the correct value of the probability that  $x$  belongs to the  $k$ th class in the single-label case. This is the probability that the model assigns  $x$  to the  $k$ th class and not to any other classes, which is given by:

$$P_{\text{SLOVA}}(k|x) = p(k|x) \prod_{j \neq k} (1 - p(j|x)). \quad (3)$$

This formula uses the information that we work in a single-label case and is consistent with the loss function (1). Consequently, the SLOVA confidence score is defined by:

$$\text{conf}_{\text{SLOVA}}(x) = \max_{k=1, \dots, K} P_{\text{SLOVA}}(k|x). \quad (4)$$

This score is high if only one class has high probability while the others are negligible.

**Remark 3.1.** Our idea of confidence for an OVA classifier can also be extended to the multi-label case. Recall that the OVA classifier returns a positive label at the  $k$ th position if  $p(k|x) \geq 1/2$ . As a consequence, the confidence score in multi-label situation can be naturally defined by:

$$\prod_{k: p(k|x) \geq 1/2} p(k|x) \cdot \prod_{j: p(j|x) < 1/2} (1 - p(j|x)).$$

In the extreme case, the OVA classifier may not return any label if  $p(k|x) < 1/2$  for all  $k = 1, \dots, K$ . The probability that “none class is correct” is thus given by:

$$P_{\text{SLOVA}}(\text{none}|x) = \prod_{k=1}^K (1 - p(k|x)).$$

We leave the further discussion on quantifying the confidence in multi-label case for future work.

In the next subsection, we continue with a theoretical analysis which shows that the SLOVA method correctly outputs low confidence scores when evaluated on the out-of-distribution data.

### 3.5. Theoretical analysis on OOD data

It is well known that softmax models tend to have overconfident predictions [19], especially for examples far away from the data distribution. In particular, it was shown that the confidence of neural networks with ReLU activation and softmax output converges to 1 for OOD data [9, Theorem 3.1]. It partially follows from the fact that the softmax model always spreads the whole probability mass to all classes because there is no event like “none class is correct”.


We show that such a limiting behavior does not hold for SLOVA (4), but is likely to occur for the typical OVA confidence score (2).

**Theorem 3.1.** Let  $f$  be a neural network with ReLU activation functions on hidden layers and sigmoid output implementing  $K$ -class OVA model. We assume that the distribution of classes is uniform.

Then for almost all  $x \in \mathbb{R}^D$ , we have  $\text{conf}_{\text{SLOVA}}(\alpha x) \rightarrow 1$  with probability  $\frac{1}{2^k}$ , as  $\alpha \rightarrow \infty$ . In the same case,  $\text{conf}_{\text{OVA}}(\alpha x) \rightarrow 1$  with probability  $1 - \frac{1}{2^k}$ , as  $\alpha \rightarrow \infty$ .

**Proof.** Let us first recall that a feedforward neural network  $f : \mathbb{R}^D \ni x \rightarrow (f_1(x), \dots, f_K(x)) \in \mathbb{R}^K$  that uses piecewise affine activation functions (e.g. ReLU, leaky ReLU) and is linear in the output layer, can be rewritten as continuous piecewise affine functions. Thus there exists a finite set of polytopes  $\{Q_l\}_{l=1}^R$  such that  $\bigcup_{l=1}^R Q_l = \mathbb{R}^D$  and  $f(x) = V^l x + a^l$  is the piecewise affine representation of the output of a ReLU network on  $Q_l$ .

By [9, Lemma 3.1], there exists a region  $Q_t$  with  $t \in \{1, \dots, R\}$  and  $\beta$  such that for all  $\alpha \geq \beta$  we have  $\alpha x \in Q_t$ . Let  $f(x) = V^l x + a^l$  be the affine form of the ReLU classifier  $f$  on  $Q_t$ . By making use of [9, Theorem 3.1] and the fact that the sigmoid is a special case of the softmax function for two classes, we get that the sigmoid output  $p(k|\alpha x)$  is arbitrarily close either to 1 or 0, for all  $k = 1, \dots, K$  and sufficiently large  $\alpha$ .



**Fig. 4.** Confidence scores of OVA and SLOVA on OOD data illustrating [Theorem 3.1](#). We show both the mean and the standard deviation of confidence scores for both models evaluated at an input  $x + \alpha x_{\text{rand}}$  composed of a CIFAR-10 test point  $x$  and a random vector  $x_{\text{rand}}$ . The  $\alpha$  parameter plotted on the horizontal axis controls the strength of the distributional shift. In contrast to OVA, whose confidence saturates at high values of  $\alpha$ , the confidence of SLOVA converges to 0 on OOD data as expected.

Observe that  $\text{conf}_{\text{SLOVA}}(\alpha x) \rightarrow 1$  iff. there exists exactly one  $k$  such that  $p(k|\alpha x) \rightarrow 1$  and  $p(i|\alpha x) \rightarrow 0$ , for all  $i \neq k$ , as  $\alpha \rightarrow \infty$ . Assuming a uniform distribution of classes, this situation occurs with probability  $\frac{1}{2^K}$ . On the other hand,  $\text{conf}_{\text{OVA}}(\alpha x) \rightarrow 1$ , as  $\alpha \rightarrow \infty$ , with probability  $1 - \frac{1}{2^K}$ , because the probability that all  $K$  sigmoids converge to 0 equals  $\frac{1}{2^K}$ .  $\square$

To illustrate the implications of the above result, let us consider a typical task with 10 classes such as the CIFAR-10 dataset. For OOD data  $x \in X_{\text{out}}$ , which are sufficiently distant from in-distribution samples, we have  $\text{conf}_{\text{SLOVA}}(x) \approx 1$  with probability  $\frac{1}{2^{10}} < 0.001$ . To construct such OOD samples it is sufficient to select almost any direction  $z \in \mathbb{R}^D$  and put  $x = \alpha z$ , where  $\alpha \in \mathbb{R}$  is a sufficiently large scalar. For analogical OOD samples,  $\text{conf}_{\text{OVA}}(x) \rightarrow 1$  with probability greater than 0.999. Consequently, a typical OVA classifier behaves analogously to a softmax classifier, while SLOVA does not inherit this drawback.

[Fig. 4](#) shows the mean confidence of OVA and SLOVA for inputs localized on a fixed 1-dimensional affine subspace  $x + \alpha x_{\text{rand}}$ . The middle point  $\alpha = 0$  is a test sample  $x$ , which represents the in-distribution data. The plot is averaged over many random subspaces  $x_{\text{rand}}$ . It is evident that OVA, on average, returns high confidence for large  $\alpha$ , while SLOVA converges to a small value. Observe the rapid drop of confidence around the test point for OVA and SLOVA, which is caused by the decrease in the predictive probability of the dominant class (the probabilities of other classes are still small). Next, the predictive probabilities of some other classes are increasing, resulting in a slight increase of SLOVA confidence. Finally, we observe a decrease of SLOVA confidence because predictive probabilities of most classes converge to 0 while others converge to 1. OVA confidence converges to 1 since it relies only on the dominant class. Analogical comparison for the two-dimensional subspace is presented in [Fig. 2](#).

### 3.6. Exponential calibration

We now describe calibration as the final postprocessing step of the SLOVA method. Using  $P_{\text{SLOVA}}(k|x)$  as predictive probabilities for  $k = 1, \dots, K$  reduces the overconfidence of a typical OVA classifier when more than one class has a high probability. However,  $\sum_{k=1}^K P_{\text{SLOVA}}(k|x) \leq 1$  because various multi-label events are additionally encoded; see [Remark 3.1](#). As a consequence, there exists a discrepancy between the predicted probabilities (confidence) and the true probabilities (accuracy). The middle point in [Fig. 4](#) illustrates the underestimation of SLOVA predictive probabilities on in-distribution data. Rather than normalizing the obtained predictive probabilities, which is only a partial solution, we instead use calibration techniques that align the expected confidence with accuracy.

To formally define the calibration, let  $\hat{p}(k|x)$  be the probability that the model returns the class  $k$  for the input example  $x$  and let  $\mathbb{P}$  be the *true* probability of classes. We say that the model

is calibrated if the true probability  $\mathbb{P}$  of class  $k$  given the output probability  $\hat{p}$  coincides with the output probability for that class:

$$\mathbb{P}(k|\hat{p}(x)) = \hat{p}(k|x), \text{ for } k \in \{1, \dots, K\}. \quad (5)$$

A typical way to calibrate the classification model relies on transforming the output probabilities  $\hat{p}$  into calibrated ones  $p$ . Based on the multiplicative form of  $P_{\text{SLOVA}}$ , we use the exponential calibration, which, we argue in [Appendix A](#), is capable of “inverting” the multiplications. Formally, we use the following parametric transformation:

$$p_k = c(\hat{p}_k) = \sum_{i=1}^M \beta_i \hat{p}_k^{\alpha_i}, \quad (6)$$

where  $\alpha_i > 0, \beta_i \in (0, 1)$  are trainable parameters such that  $\sum_{i=1}^M \beta_i = 1$ . Due to the choice of  $\alpha_i, \beta_i$  in the above ranges, the introduced transformation is monotonic and is a map on the unit interval. In the [Appendix A](#) we argue that the exponential map is a natural choice to invert the multiplicative character of SLOVA probabilities. Moreover, it was shown that the above exponential transformation is capable of approximating any completely monotonic function [7,42].

To calibrate SLOVA, we form a calibration training set  $X_{\text{cal}}$  from the validation dataset  $X_{\text{val}}$ . For each example  $(x, y) \in X_{\text{val}}$ , we gather all  $K$  pairs of uncalibrated output probabilities and their corresponding one-hot encoded binary labels  $(P_{\text{SLOVA}}(k|x), \delta_{k,y}) \in X_{\text{cal}}$ , which results in  $|X_{\text{cal}}| = K|X_{\text{val}}|$  pairs in total.

Next, we sort the pairs by the output probabilities  $P_{\text{SLOVA}}$ :

$$\begin{cases} P := \text{sort}(P_{\text{SLOVA}}(k|x)), & k \in (1 \dots K), \quad (x, y) \in X_{\text{val}}; \\ D := \text{sort}_{P_{\text{SLOVA}}(k|x)}(\delta_{k,y}), \end{cases}$$

and apply sliding windows of size  $n$  to compute moving averages of the sorted data:

$$\bar{\pi}_i = \frac{1}{n} \sum_{j=i}^{i+n-1} P_j, \quad \bar{b}_i = \frac{1}{n} \sum_{j=i}^{i+n-1} D_j.$$

This results in  $N = |X_{\text{cal}}| - n + 1$  averaged points, and we fix the window size  $n = |X_{\text{cal}}|/100$  so that there are enough points within each window. Intuitively,  $\bar{b}_i$  is the fraction of positive labels, while  $\bar{\pi}_i$  represents the average confidence of the model in that window. Lastly, we select an equally spaced subset of the averaged pairs  $(\bar{\pi}, \bar{b})$  so that the total number of pairs in the dataset is  $n_b$ . The parameters  $\alpha, \beta$  from (6) are directly optimized by minimizing the mean square error:

$$\sum_{i=1}^{n_b} (c(\bar{\pi}_i) - \bar{b}_i)^2.$$

Note that calibration methods are normally trained and evaluated on in-distribution data. This means that calibration usually

does not improve and can actually increase confidence on out-of-distribution data. To prevent this, we additionally augment  $X_{val}$  with random noise samples that are labeled “none”, that is,  $\delta_k = 0$  for  $k \in (1..K)$  (only zeros in one-hot representation). We use  $[0, 1]$  uniform distribution for the image data as we find that it works adequately and add  $\lfloor 0.1|X_{val}| \rfloor$  of such samples.

The choice of hyperparameter  $M$  affects the capacity and computational cost of the transformation. Note, however, that the cost is negligible compared to the cost of the forward pass of the base model. In Appendix B we provide stability studies of exponential calibration as we vary the number of adjustable parameters  $M$  and the size of the calibration dataset  $n_b$ .

## 4. Experiments

To evaluate both the confidence calibration and the out-of-distribution aspects of uncertainty estimation, we focus on three tasks presented in the following sections. First, we measure the model calibration on in-distribution data. Then we repeat the same experiment under distribution shift. We then proceed to examine the behavior of our model on out-of-distribution data. Lastly, we analyze the significance of the calibration step in the final section.

### 4.1. In-distribution and shifted data

*Setup.* To examine the model calibration on in-distribution and distorted data, we follow a recent benchmark [8] and compare SLOVA with the following baselines:

- *Softmax*: Maximum softmax probability [43],
- *Temp Scaling*: Softmax with calibration defined by the temperature scaling [19],
- *Dropout*: Monte Carlo Dropout [14],
- *Ensembles*: Ensemble of  $M$  neural networks trained independently on the entire dataset [16],
- *SVI*: Stochastic Variational Bayesian Inference for deep learning [11,44],
- *LL SVI*: Mean field stochastic variational inference on the last layer only [45],
- *LL Dropout*: Dropout applied only to activations before the last layer [45],
- *OVA DM*: OVA classifier with distance-based logits [5].

In this setting, we train all models on the CIFAR-10 train set and subsequently calibrate them on the validation set. In-distribution calibration is evaluated on the original test set. To investigate the robustness under dataset shift, we follow Hendrycks & Dietterich [43] and use CIFAR-10-C produced by applying 16 types of distortions to CIFAR-10 images, each taken with 5 degrees of intensity.

We use the ResNet20 architecture [46] that achieves state-of-the-art results on typical computer vision tasks, but is otherwise an arbitrary choice to compare methods. We train it with the Adam optimizer [47] for 250 epochs with batch size 512. The train dataset is augmented by random crops and rotations, the learning rate is piecewise constant with two 0.1 drops in epochs 125 and 185. Before training, we perform a random hyperparameter search of 20 trials with learning rates sampled uniformly from an interval  $[10^{-3}, 1]$ , momentum  $\beta_1 \in [0.85, 0.99]$ , and stability parameter  $\epsilon \in [10^{-8}, 10^{-5}]$ . Calibration was performed for  $M = 20$  parameters trained for 20 epochs with the size of the calibration dataset  $n_b = 4000$ .

*Metrics.* We investigate the classification accuracy as well as three calibration-related metrics. The Expected Calibration Error (ECE) [48] measures the discrepancy between predicted probabilities and empirical accuracy by grouping model predictions into  $M$  interval bins  $B_1, \dots, B_M$  and calculating the expected difference between the accuracy  $\text{acc}(B_i)$  and confidence  $\text{conf}(B_i)$  over the bins:

$$ECE = \sum_{i=1}^M \frac{|B_i|}{M} |\text{acc}(B_i) - \text{conf}(B_i)|.$$

We also consider the Brier score [49] – the squared error of the predicted probability vector and the one-hot encoded true response; and the negative log-likelihood (NLL), which is also a proper scoring rule [50], but a less objective one, because it can overemphasize tail probabilities [51].

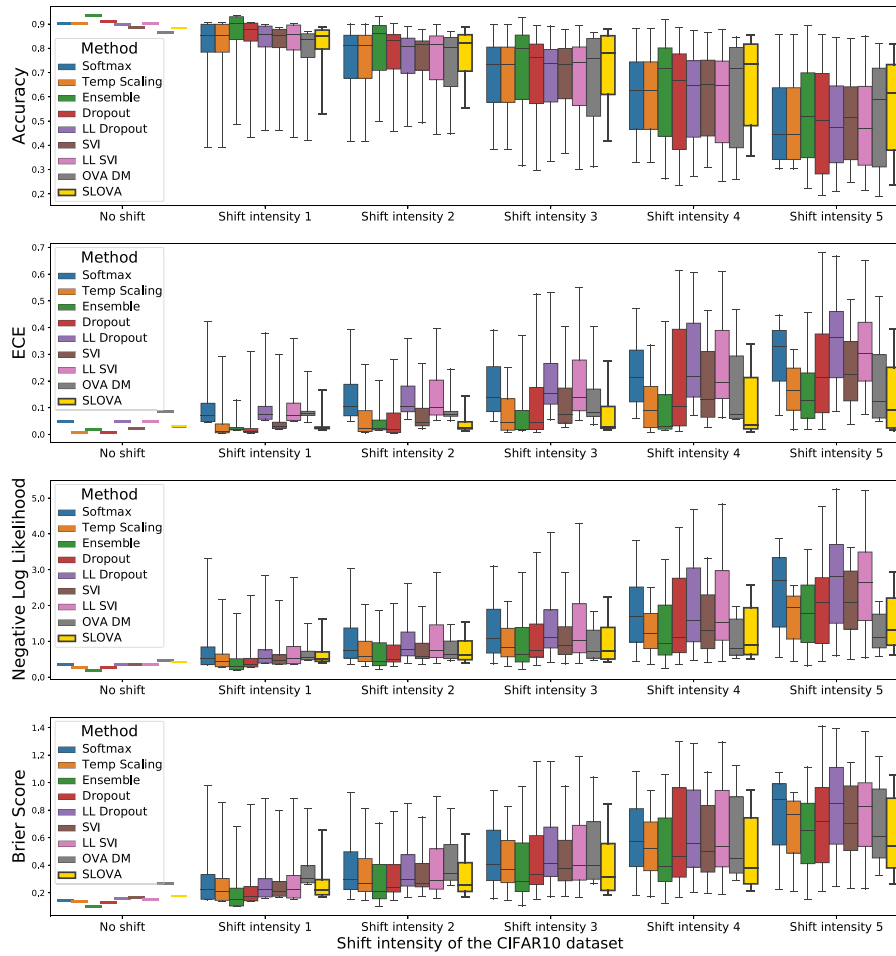
*Results.* The results presented in Fig. 5 demonstrate that SLOVA gives comparable accuracy and calibration scores to vanilla softmax on in-distribution data, but is more robust to dataset shifts. At the same time, it presents a notable improvement over softmax and its calibrated variant (Temperature Scaling) on highly distorted data in terms of all metrics. In the case of high intensity dataset shifts (level 4 and 5), the calibration scores (ECE, NLL, Brier) of SLOVA are frequently better than the ones obtained by ensembles, which is currently considered the state-of-the-art method in uncertainty modeling. Note that complex methods, such as Ensemble, are significantly more computationally demanding than SLOVA. While SLOVA relies on using a single OVA model, Ensemble requires training and evaluation of multiple classifiers, which are also trained using adversarial examples.

The advantage of SLOVA on highly distorted data can be explained by the fact that SLOVA is designed not only to improve predictive uncertainty but also to detect OOD data. Strong shifts pushed the data far away from the true data distribution, which is similar to creating OOD samples. This experiment confirms our thesis that most methods are created to solve only one particular problem – in-distribution calibration, robustness under dataset shifts, or OOD detection. In contrast, SLOVA is the only method that gives satisfactory results on all of these tasks.


It is also worth noting that SLOVA compares favorably with OVA DM, which is an upgraded version of the OVA classifier. OVA DM performs poorly for the in-distribution data (drops in accuracy observed both in our experiments and in the original article [5]), a weakness that our method does not share. Nevertheless, in the case of high intensity dataset shifts, the performance of OVA DM is still lower than SLOVA, which works well in all cases.

*Statistical analysis.* To support our analysis, we perform statistical tests to verify the statistical difference between algorithms. Specifically, we combine the nonparametric statistical Friedman test with the Bonferroni–Dunn post hoc analysis [52]. To this end, we aggregated the results of all the methods from Fig. 5, giving 81 in total. Next, we ranked the methods according to their performance. Given this ranking, the analysis consists of two steps:

- The null hypothesis is made that all methods perform equally and the observed differences are merely random. The hypothesis is tested by the Friedman test, which follows a  $\chi^2$  distribution.
- Having rejected the null hypothesis, the differences in ranks are analyzed by the Bonferroni–Dunn test. The Bonferroni–Dunn tests are applied only between SLOVA and other methods.



**Fig. 5.** Accuracy and three uncertainty-aware metrics: ECE, NLL, and Brier score for the CIFAR-10 and its distorted version. On the horizontal axis we group model outputs with 5 degrees of shift intensity constructed following Hendrycks & Dietterich [43]. Despite its simplicity, SLOVA retains high accuracy and low Brier score across all shift intensities while ECE and NLL are comparable to more complex state-of-the-art models. Based on these results, in Fig. 6 we conduct a statistical performance test which ranks SLOVA as the second-best giving up the place only to the computationally demanding ensembles. The proposed model is especially robust for higher values of shift severity where other models, in general, are prone to larger drops of performance.



**Fig. 6.** Visualization of statistical comparison performed on in-distribution calibration and dataset shifts. Methods which lie inside the critical distance around SLOVA are statistically indistinguishable.

Fig. 6 visualizes the results for a significance level of  $\alpha = 0.05$ . The x-axis shows the mean rank for each method. The thick horizontal line illustrates the critical distance. Methods for which the difference in mean rank is lower than the critical distance are statistically indistinguishable.

It is evident that the mean rank of SLOVA is comparable to Dropout and Temp Scaling. Since the distance between their mean ranks is lower than the critical distance, the difference between these methods is statistically insignificant. At the same time, SLOVA is statistically better than other models, except Ensemble, which performs best across all cases.

#### 4.2. Out-of-distribution data

*Setup.* We train a classification model on one dataset for the OOD experiment and then evaluate it on other out-of-distribution datasets. We reuse the setup from [9] and compare SLOVA with the following methods:

- *Softmax:* Maximum softmax probability,
- *CEDA:* Softmax network trained with an additional loss term that penalizes non-uniform answers on additional OOD train data [9,35],

**Table 1**

Mean maximum confidence for the out-of-distribution (OOD) experiment. Low confidence on OOD data is desired and SLOVA produces top results on 9 out of 12 cases.

Method	Error	Mean maximum confidence			
		CIFAR-10	CIFAR-100	LSUN	ImageNet
SVHN					
Softmax	3.53%	0.732	0.730	0.722	0.568
CEDA	3.50%	0.551	0.527	0.364	0.736
ACET	3.52%	0.435	0.414	<b>0.148</b>	0.421
ReAct	3.77%	0.671	0.676	0.673	0.489
SLOVA	3.29%	<b>0.376</b>	<b>0.387</b>	0.338	<b>0.146</b>
CIFAR-10					
SVHN					
Softmax	8.87%	0.800	0.764	0.738	0.864
CEDA	8.87%	0.327	0.761	0.735	0.660
ACET	8.44%	<b>0.263</b>	0.764	0.745	0.242
ReAct	9.19%	0.725	0.733	0.689	0.713
SLOVA	7.68%	0.644	<b>0.640</b>	<b>0.566</b>	<b>0.180</b>
CIFAR-100					
SVHN					
Softmax	31.97%	0.570	0.560	0.592	0.878
CEDA	32.74%	0.290	0.547	0.581	0.542
ACET	32.24%	<b>0.234</b>	0.530	0.554	0.601
ReAct	37.64%	0.426	0.448	0.510	0.667
SLOVA	32.01%	0.407	<b>0.402</b>	<b>0.386</b>	<b>0.098</b>

- *ACET*: Softmax network also trained with the same loss term, but with OOD examples being generated adversarially [9],
- *ReAct*: Softmax network with penultimate activations being clipped to a threshold in evaluation [53].

A ResNet model is trained on each of the three commonly used datasets: CIFAR-10, CIFAR-100, and SVHN. Random crops are used for all datasets for training, while random mirroring is additionally used for CIFAR-10 and CIFAR-100. We evaluate each model on the two remaining datasets along with the LSUN (classroom subset) and the ImageNet datasets. All inputs are  $32 \times 32$  pixels except for the ImageNet dataset, where the images are scaled to  $224 \times 224$  pixels. For ReAct, we follow the original authors and set the threshold  $c$  so that it corresponds to the value of the 90th percentile of all activations from the penultimate layer [53].

For SLOVA we slightly tune the original hyperparameters of the softmax model to avoid a large drop in the accuracy of the OVA model. This is easily achieved by altering the learning rate and the weight decay hyperparameter  $\lambda$  - we use the learning rate 0.02, 0.03, 0.02, and  $\lambda$  0.01, 0.02, 0.01 for CIFAR-10, CIFAR-100, and SVHN, respectively. The remaining hyperparameters are exactly the same as in the setup from [9].

*Metrics.* Similarly as in [9], we report the test error and the mean maximal confidence (MMC) for each model. MMC is defined as the mean of maximal model predictions taken over all data points. In the case of SLOVA, MMC is evaluated using  $\max_{k=1, \dots, K} P_{\text{SLOVA}}(k|x)$ , averaged over all  $x$ . Low MMC means low confidence, which is expected on OOD data.

*Results.* The results reported in Table 1 demonstrate that SLOVA gives significantly lower MMC scores than the typical softmax classifier on OOD detection task. It is interesting that ReAct, a recent OOD detector, performs only slightly better than softmax. At the same time, it is inferior to SLOVA in all cases and has a significant reduction in classification accuracy for the model trained on CIFAR-100.

Notably, ACET gives impressive scores in a few variants (SVHN/LSUN, CIFAR-10/SVHN, CIFAR-100/SVHN), but its results are very unstable. In particular, ACET trained on CIFAR-10 is able to almost perfectly detect SVHN and ImageNet samples, but it is unaware of examples from the CIFAR-100 and LSUN datasets.

This behavior can be explained by the fact that ACET is trained using adversarial examples, which is a complex and quite difficult procedure to apply in practice. In this case, it performs even worse than CEDA, a similar method that uses noise instead of adversarial examples.

In terms of OOD detection, SLOVA outperforms other methods in 9 out of 12 cases. Furthermore, this experiment shows that OVA classifiers provide accuracy comparable to that of softmax models in typical classification tasks. This is an important conclusion because most neural networks are based on softmax. Since SLOVA performs similarly to softmax but is better in OOD detection and predictive uncertainty, it is safe to replace the softmax model with a SLOVA classifier.

*Statistical analysis.* We repeat the statistical analysis for the OOD detection. Namely, we conduct the non-parametric statistical Friedman test with the Bonferroni–Dunn post hoc analysis for 12 combinations of datasets reported in Table 1.

As can be seen in Fig. 7, SLOVA received a rank close to 1, which confirms that it was superior in most cases. More importantly, the difference between SLOVA and other models is statistically significant.

#### 4.3. Ablation study

We perform the following ablation study to analyze the impact of the proposed confidence score and the calibration separately. We group all the cases of in-distribution data and distributional shifts together (including shift intensities). First, we examine the standard OVA classifier, then SLOVA without exponential calibration, and finally the complete SLOVA model. For comparison, we also report the vanilla softmax model.

The results reported in Fig. 8 confirm that subsequent modifications do not affect accuracy. This is expected behavior because both SLOVA predictive probabilities and exponential calibration do not change the order of probability scores. As a consequence, the accuracy of the OVA model is preserved.

It is evident that subsequent modifications of the OVA probabilities produced by SLOVA successively improve the ECE score. The impact on the Brier score is smaller, but there is a slight positive trend. In the case of NLL, there is no evident trend, but a complete SLOVA model obtains the best result with the lowest variance.

The most significant difference observed in the case of the ECE metric can result from the fact that the exponential calibration was optimized using a type of ECE objective. Although all calibration metrics are related, the selected optimization objective is crucial for the final performance. In most studies, the ECE is considered the most reliable calibration metric, which motivated our selection. It is worth emphasizing that many methods (including temperature scaling) use log-likelihood to select calibration parameters because its optimization is significantly easier.

## 5. Conclusion

### 5.1. Summary

In this work, we presented SLOVA, a simple and effective method of transforming OVA probabilities into reliable uncertainty predictions. Since OVA can be used in multi-label classification tasks, it has a natural ability to detect OOD samples by labeling them as “none of the above classes”. As a consequence, SLOVA goes beyond the properties of a typical softmax predictor and is better suited for modeling the distribution of OOD samples. We confirmed this claim theoretically. To quantify the uncertainty of the extended OVA classifier on in-distribution and shifted data, we introduced the exponential calibration, which






Fig. 7. Visualization of statistical comparison performed on OOD detection. The difference between SLOVA and other methods are statistically significant, which is indicated by the critical distance.

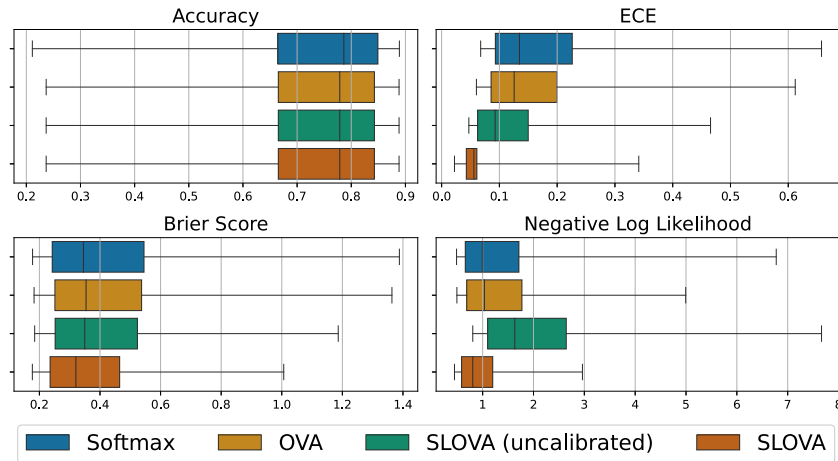


Fig. 8. Ablation study on CIFAR-10-C testsets. It is evident that each step of our method successively improves the uncertainty estimation while preserving the model accuracy.

allows us to align predictive probabilities with the accuracy of the model. To support our theses, we experimentally demonstrated that SLOVA compares favorably against state-of-the-art methods on common benchmark tasks.

### 5.2. Discussion and limitations

The problem of uncertainty estimation can be divided into at least three separate tasks: in-distribution calibration, robustness under dataset shifts, and detection of OOD samples. Although all of these tasks are closely related, most methods focus on solving only one or two of them. We showed that combining exponential calibration with appropriate treatment of the OVA classifier gives a model that presents satisfactory performance on all of these problems. In particular, SLOVA is one of the three best-performing methods on in-distribution calibration with optional dataset shifts. At the same time, it outperforms state-of-the-art baselines on OOD detection in most cases. In consequence, SLOVA is a method that performs well in diverse applications, while other methods are directly designed only for a particular task.

It is not obvious whether SLOVA can be applied to the case of extreme classification. SLOVA confidence score is defined as a product of OVA predictive probabilities. When the number of classes is high, but only one is correct, a single mistake can result in inaccurate or incorrect predictions. Such a disadvantage disappears in the multi-label classification, where more than one label is correct and we select the most probable combination of labels. A possible remedy in a single-label situation would be to project the class vector to a lower-dimensional space as it is commonly done in extreme classification [54].

One could argue that a classification neural network implemented with the use of a softmax function is the standard approach that gives state-of-the-art performance. As a consequence, most of the available pre-trained models are softmax models.

This is a considerable limitation, as even though SLOVA is a post hoc method, it requires an OVA model. However, OVA classifiers usually give performance comparable to softmax, and both models can be trained effectively. In consequence, SLOVA can be an alternative to softmax in applications where reliable uncertainty modeling is a desirable property.

Finally, observe that [Theorem 3.1](#) states that for ReLU activations, there is a  $\frac{1}{2^k}$  probability of selecting a direction for which the network is overconfident with arbitrarily high  $\alpha$ . In practice, OOD samples can be distributed closer to in-distribution samples than  $\alpha x$  with saturated confidence. Despite the limitations of these theoretical results, our experiments confirm that SLOVA performs particularly well across several OOD benchmark tasks, strongly suggesting that the proposed method is useful for real-life use cases.

### 5.3. Future works

We showed that the predictive probability returned by the OVA classifier can be used to define confidence score in a single-label classification. As discussed in [Remark 3.1](#), we can formulate an analogical confidence measure in a multi-label situation. In future work, we plan to explore this idea and verify it experimentally.

Uncertainty estimation plays a crucial role in many areas of machine learning, including active learning, semi-supervised learning, or conditional computation. It would be interesting to verify whether the application of SLOVA in these areas improves performance.

An interesting direction for future work would be to improve the method by combining it with approaches that are orthogonal, and thus could be applied simultaneously. For example, it is well known that constructing an ensemble of networks improves the uncertainty scores [16], yet the question of whether this fact is

also true for SLOVA remains unanswered. Another simple way to improve OOD detection performance is to expose the model to some OOD samples in the training process [35], and one could augment SLOVA with this approach.

### CRedit authorship contribution statement

**Bartosz Wójcik:** Conceptualization, Methodology, Software, Validation, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Jacek Greła:** Conceptualization, Methodology, Software, Validation, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Marek Śmieja:** Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing, Supervision, Project administration, Funding acquisition. **Krzysztof Misztal:** Software, Writing – original draft. **Jacek Tabor:** Conceptualization, Supervision, Project administration, Funding acquisition.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgments

The work of Bartosz Wójcik and Marek Śmieja was supported by the National Science Centre (Poland) grant no. 2018/31/B/ST6/00993. Jacek Tabor carried out this work within the research project “Bio-inspired artificial neural networks” (grant no. POIR.04.04.00-00-14DE/18-00) within the Team-Net program of the Foundation for Polish Science co-financed by the European Union under the European Regional Development Fund.

### Appendix A. Justification of the exponential calibration

We justify a particular form of the exponential calibration (5) whose aim is to “invert” the multiplication of many probabilities or to counteract the contractive character of this multiplication. We start off from an (uncalibrated) output of the SLOVA model given in (3) for a generic input  $x$ :

$$\hat{P} = \hat{p}_1(1 - \hat{p}_2)\dots(1 - \hat{p}_K),$$


where we introduce a succinct notation for the uncalibrated probabilities  $\hat{p}(i|x) \rightarrow \hat{p}_i, \hat{P}_{\text{SLOVA}}(1|x) \rightarrow \hat{P}$ . Without loss of generality, we pick out the first probability as the “correct” one. In general, since all  $\hat{p}_i$ 's are numbers between zero and one, the multiplicative SLOVA probability  $\hat{P}$  becomes, as  $K$  increases, more likely to reach zero. For single-label examples,  $\hat{p}_1 \approx 1$  and  $\hat{p}_{i \neq 1} \approx 0$ , which results in a (relatively small) drift away from 1. However, it quickly becomes unacceptably large when the number of classes  $K$  grows.

We consider this multiplication-induced drift in a tractable, yet artificial case of complete ignorance where each partial probability  $\hat{p}_i \sim U(0, 1)$  is drawn from an independent uniform distribution. A simple calculation using the Mellin transform results in the distribution over SLOVA probabilities:

$$\mathbb{P}_{\text{uncal}}(\hat{P}) = \frac{1}{(K-1)!} \left( \log \frac{1}{\hat{P}} \right)^{K-1} \quad (\text{A.1})$$

which becomes highly concentrated around  $\hat{P} \sim 0$  even for a moderate number of classes  $K = 10$ .

On the other hand, calibration (6) is a functional relation  $P = c(\hat{P})$  that maps between the uncalibrated  $\hat{P}$  and calibrated probabilities  $P$ . In the special case (A.1), a valid calibration  $c$



**Fig. A.9.** Calibration curve in the fully random case. The black solid line is the exact calibration curve (A.2) while the thin dashed line is its exponential approximate form. We use  $K = 3$ .

is conceptualized as a function that reverts back to a form of complete ignorance, i.e., the initial uniform probability density:

$$\mathbb{P}_{\text{cal}}(P) = \int d\hat{P} \delta(P - c(\hat{P})) \mathbb{P}_{\text{uncal}}(\hat{P}) = 1, \quad \text{for } P \in (0, 1).$$

With the usual properties of the Dirac delta function, a transform achieving this is simply the cumulative distribution function:

$$c(\hat{P}) = \int_{-\infty}^{\hat{P}} \mathbb{P}_{\text{uncal}}(\hat{P}') d\hat{P}'. \quad (\text{A.2})$$

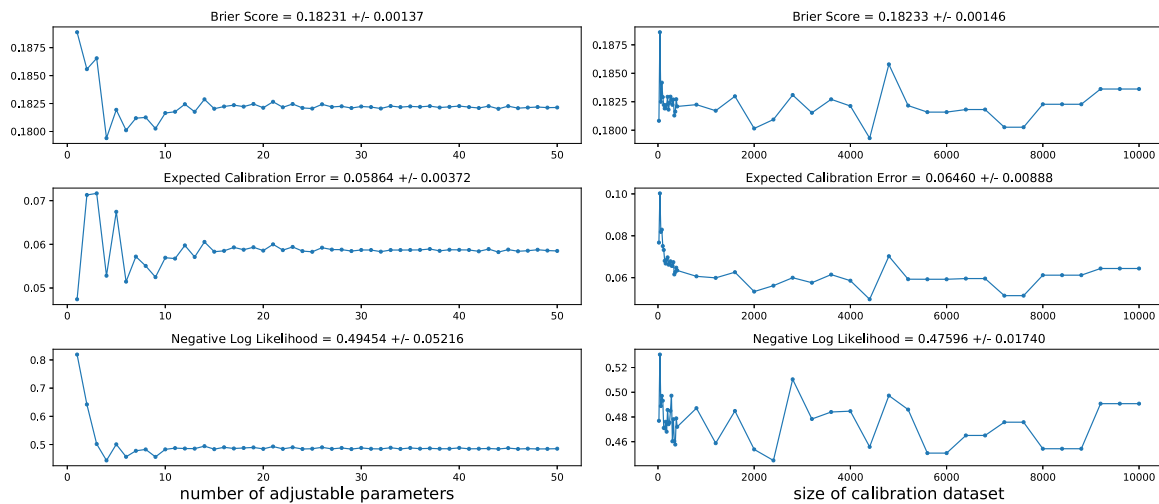
Then, in the fully random case, the calibration map is the incomplete gamma function  $c(\hat{P}) = \frac{\Gamma(K, \log 1/\hat{P})}{(K-1)!}$ . Now, we look for an approximation of this form in terms of elementary functions. The resulting calibration curve is shown in Fig. A.9. To this end, we expand the incomplete Gamma function around  $\hat{P} = 1$  resulting in  $c(\hat{P}) \sim 1 - \frac{(1-\hat{P})^K}{(K-1)!}$  and then ensure that the approximation is properly normalized  $P \in (0, 1)$ . The simplified calibration curve reads  $c_{\text{appr}}(\hat{P}) = 1 - (1 - \hat{P})^K$ . We show both exact and approximate curves in Fig. A.9.

This simple argument justifies the exponential form of the calibration curve  $c_{\text{appr}}(\hat{P}) \sim \hat{P}^\alpha$ . At the same time, it is based on a generic example and needs further generalization resulting in formula (5) where we make a linear combination of exponential terms and train both the linear coefficients  $\beta$  as well as the exponents  $\alpha$ .

### Appendix B. Calibration stability

The exponential calibration proposed in this work is an integral and novel part of the SLOVA approach. Therefore, to assess any shortcomings of the method, we inspect it under changes in the number of adjustable parameters  $M$  and the size of the calibration dataset  $n_b$ .

The number of parameters  $M$  refers to the pairs of weights  $(\alpha_i, \beta_i)$  where  $i = 1, \dots, M$  as defined in (6). The size of the calibration dataset  $n_b$  is in turn discussed in Section 3.6. We assess the robustness of the calibration algorithm in Fig. B.10 by



**Fig. B.10.** Stability of the calibration method under changes in the number of adjustable parameters  $M$  and in the size of the calibration dataset  $n_b$ . We inspect variability of the NLL, ECE, Brier score as we change  $M \in (1, 50)$  and  $n_b \in (20, 10000)$  and identify a stable region of parameters for  $M > 10$  and  $n_b > 400$ .

calculating ECE, NLL, and Brier score metrics across changes in  $M \in (1, 50)$  and  $n_b \in (20, 10000)$ .

In the case of adjustable parameters  $M$ , after a brief instability period when  $M < 10$ , we enter the regime with  $M > 10$  where the metric values do not change appreciably. Similarly, a similar stability region is found if the size of the dataset  $n_b$  is greater than 400. In conclusion, the calibration method is robust with respect to these parameter changes.

### Appendix C. Table of metrics for the dataset shift experiment

The tables below report quartiles of Brier score, negative log-likelihood, and ECE in the dataset shift experiment considered in Section 4.1 where quartiles are computed over all corrupted variants of the dataset.

Method	SLOVA	Dropout	Ensemble	LL Dropout	LL SVI
Brier score (25th)	0.215	0.217	0.160	0.258	0.241
Brier score (50th)	0.299	0.365	0.294	0.412	0.405
Brier score (75th)	0.553	0.657	0.556	0.746	0.735
ECE (25th)	0.021	0.012	0.018	0.089	0.075
ECE (50th)	0.027	0.039	0.024	0.151	0.137
ECE (75th)	0.103	0.211	0.107	0.313	0.291
NLL (25th)	0.507	0.452	0.334	0.619	0.572
NLL (50th)	0.709	0.811	0.636	1.070	1.062
NLL (75th)	1.383	1.743	1.529	2.377	2.368

Method	SVI	Temp scaling	Softmax	OVA DM
Brier score (25th)	0.245	0.228	0.242	0.300
Brier score (50th)	0.371	0.382	0.412	0.398
Brier score (75th)	0.635	0.649	0.733	0.702
ECE (25th)	0.034	0.014	0.074	0.063
ECE (50th)	0.074	0.049	0.133	0.079
ECE (75th)	0.214	0.140	0.284	0.166
NLL (25th)	0.520	0.475	0.574	0.534
NLL (50th)	0.839	0.850	1.012	0.731
NLL (75th)	1.714	1.589	2.219	1.294

### References

- [1] R. Miotto, L. Li, B.A. Kidd, J.T. Dudley, Deep patient: an unsupervised representation to predict the future of patients from the electronic health records, *Sci. Rep.* 6 (1) (2016) 1–10.
- [2] J. Levinson, J. Askeland, J. Becker, J. Dolson, D. Held, S. Kammel, J.Z. Kolter, D. Langer, O. Pink, V. Pratt, et al., Towards fully autonomous driving: Systems and algorithms, in: 2011 IEEE Intelligent Vehicles Symposium (IV), IEEE, 2011, pp. 163–168.
- [3] R. Berk, An impact assessment of machine learning risk forecasts on parole board decisions and recidivism, *J. Exp. Criminol.* 13 (2) (2017) 193–216.
- [4] Z.-M. Chen, X.-S. Wei, P. Wang, Y. Guo, Multi-label image recognition with graph convolutional networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 5177–5186.
- [5] S. Padhy, Z. Nado, J. Ren, J. Liu, J. Snoek, B. Lakshminarayanan, Revisiting one-vs-all classifiers for predictive uncertainty and out-of-distribution detection in neural networks, 2020, arXiv preprint [arXiv:2007.05134](https://arxiv.org/abs/2007.05134).
- [6] G. Franchi, A. Bursuc, E. Aldea, S. Dubuisson, I. Bloch, One versus all for deep neural network incertitude (ovnni) quantification, 2020, arXiv preprint [arXiv:2006.00954](https://arxiv.org/abs/2006.00954).
- [7] D.W. Kammler, Least squares approximation of completely monotonic functions by sums of exponentials, *SIAM J. Numer. Anal.* 16 (5) (1979) 801–818.
- [8] Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. Dillon, B. Lakshminarayanan, J. Snoek, Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Vol. 32, Curran Associates, Inc., 2019.
- [9] M. Hein, M. Andriushchenko, J. Bitterwolf, Why ReLU networks yield high-confidence predictions far away from the training data and how to mitigate the problem, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 41–50.
- [10] J. Gawlikowski, C.R.N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, et al., A survey of uncertainty in deep neural networks, 2021, arXiv preprint [arXiv:2107.03342](https://arxiv.org/abs/2107.03342).
- [11] C. Blundell, J. Cornebise, K. Kavukcuoglu, D. Wierstra, Weight uncertainty in neural network, in: International Conference on Machine Learning, PMLR, 2015, pp. 1613–1622.
- [12] J. Yao, W. Pan, S. Ghosh, F. Doshi-Velez, Quality of uncertainty quantification for Bayesian neural network inference, 2019, arXiv preprint [arXiv:1906.09686](https://arxiv.org/abs/1906.09686).
- [13] J.M. Hernández-Lobato, R. Adams, Probabilistic backpropagation for scalable learning of bayesian neural networks, in: International Conference on Machine Learning, PMLR, 2015, pp. 1861–1869.
- [14] Y. Gal, Z. Ghahramani, Dropout as a bayesian approximation: Representing model uncertainty in deep learning, in: International Conference on Machine Learning, PMLR, 2016, pp. 1050–1059.
- [15] W.J. Maddox, P. Izmailov, T. Garipov, D.P. Vetrov, A.G. Wilson, A simple baseline for bayesian uncertainty in deep learning, *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [16] B. Lakshminarayanan, A. Pritzel, C. Blundell, Simple and scalable predictive uncertainty estimation using deep ensembles, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, in: NIPS'17, 2017, pp. 6405–6416.

- [17] T. Pearce, F. Leibfried, A. Brintrup, Uncertainty in neural networks: Approximately bayesian ensembling, in: International Conference on Artificial Intelligence and Statistics, PMLR, 2020, pp. 234–244.
- [18] F.K. Gustafsson, M. Danelljan, T.B. Schon, Evaluating scalable bayesian deep learning methods for robust computer vision, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2020, pp. 318–319.
- [19] C. Guo, G. Pleiss, Y. Sun, K.Q. Weinberger, On calibration of modern neural networks, in: International Conference on Machine Learning, PMLR, 2017, pp. 1321–1330.
- [20] M. Kull, M. Perello Nieto, M. Kängsepp, T. Silva Filho, H. Song, P. Flach, Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with Dirichlet calibration, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, R. Garnett (Eds.), Advances in Neural Information Processing Systems, Vol. 32, Curran Associates, Inc., 2019.
- [21] J. Zhang, B. Kailkhura, T.Y.-J. Han, Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning, in: International Conference on Machine Learning, PMLR, 2020, pp. 1117–1128.
- [22] A. Rahimi, A. Shaban, C.-A. Cheng, R. Hartley, B. Boots, Intra order-preserving functions for calibration of multi-class neural networks, in: H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, Vol. 33, Curran Associates, Inc., 2020, pp. 13456–13467.
- [23] D. Milius, R. Camoriano, P. Michiardi, L. Rosasco, M. Filippone, Dirichlet-based Gaussian processes for large-scale calibrated classification, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), Advances in Neural Information Processing Systems, Vol. 31, Curran Associates, Inc., 2018.
- [24] J. Wenger, H. Kjellström, R. Triebel, Non-parametric calibration for classification, in: International Conference on Artificial Intelligence and Statistics, PMLR, 2020, pp. 178–190.
- [25] J. Mukhoti, V. Kulharia, A. Sanyal, S. Golodetz, P. Torr, P. Dokania, Calibrating deep neural networks using focal loss, in: H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, Vol. 33, Curran Associates, Inc., 2020, pp. 15288–15299.
- [26] S. Seo, P.H. Seo, B. Han, Learning for single-shot confidence calibration in deep neural networks through stochastic inferences, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 9030–9038.
- [27] A. Kumar, S. Sarawagi, U. Jain, Trainable calibration measures for neural networks from kernel mean embeddings, in: International Conference on Machine Learning, PMLR, 2018, pp. 2805–2814.
- [28] C. Leibig, V. Allken, M.S. Ayhan, P. Berens, S. Wahl, Leveraging uncertainty information from deep neural networks for disease detection, *Sci. Rep.* 7 (1) (2017) 1–14.
- [29] D. Hendrycks, K. Gimpel, A baseline for detecting misclassified and out-of-distribution examples in neural networks, 2016, arXiv preprint arXiv:1610.02136.
- [30] T. DeVries, G.W. Taylor, Learning confidence for out-of-distribution detection in neural networks, 2018, arXiv preprint arXiv:1802.04865.
- [31] S. Liang, Y. Li, R. Srikant, Enhancing the reliability of out-of-distribution image detection in neural networks, in: International Conference on Learning Representations, 2018.
- [32] K. Lee, H. Lee, K. Lee, J. Shin, Training confidence-calibrated classifiers for detecting out-of-distribution samples, in: International Conference on Learning Representations, 2018.
- [33] W. Wang, A. Wang, A. Tamar, X. Chen, P. Abbeel, Safer classification by synthesis, 2017, arXiv preprint arXiv:1711.08534.
- [34] E. Nalisnick, A. Matsukawa, Y.W. Teh, D. Gorur, B. Lakshminarayanan, Do deep generative models know what they don't know? in: International Conference on Learning Representations, 2019, URL <https://openreview.net/forum?id=H1xwNhCcYm>.
- [35] D. Hendrycks, M. Mazeika, T. Dietterich, Deep anomaly detection with outlier exposure, in: International Conference on Learning Representations, 2019, URL <https://openreview.net/forum?id=HycCxhRcY7>.
- [36] A. Tewari, P.L. Bartlett, On the consistency of multiclass classification methods, *J. Mach. Learn. Res.* 8 (5) (2007).
- [37] N. Carlini, D. Wagner, Adversarial examples are not easily detected: Bypassing ten detection methods, in: Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security, 2017, pp. 3–14.
- [38] J. Yang, K. Zhou, Y. Li, Z. Liu, Generalized out-of-distribution detection: A survey, 2021, arXiv preprint arXiv:2110.11334.
- [39] B. Zadrozny, C. Elkan, Transforming classifier scores into accurate multiclass probability estimates, in: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2002, pp. 694–699.
- [40] L. Shu, H. Xu, B. Liu, DOC: Deep open classification of text documents, in: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Copenhagen, Denmark, 2017.
- [41] D. Macêdo, T.I. Ren, C. Zanchettin, A.L.I. Oliveira, T. Luderemir, Entropic out-of-distribution detection, in: 2021 International Joint Conference on Neural Networks (IJCNN), 2021, pp. 1–8.
- [42] R.J. McGlenn, Uniform approximation of completely monotone functions by exponential sums, *J. Math. Anal. Appl.* 65 (1) (1978) 211–218.
- [43] D. Hendrycks, T. Dietterich, Benchmarking neural network robustness to common corruptions and perturbations, in: International Conference on Learning Representations, 2019.
- [44] A. Graves, Practical variational inference for neural networks, in: Advances in Neural Information Processing Systems, Citeseer, 2011, pp. 2348–2356.
- [45] C. Riquelme, G. Tucker, J. Snoek, Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling, in: International Conference on Learning Representations, 2018.
- [46] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [47] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: ICLR (Poster), 2015, URL <http://arxiv.org/abs/1412.6980>.
- [48] M.P. Naeni, G. Cooper, M. Hauskrecht, Obtaining well calibrated probabilities using bayesian binning, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 29, 2015.
- [49] G.W. Brier, Verification of forecasts expressed in terms of probability, *Mon. Weather Rev.* 78 (1) (1950) 1–3.
- [50] T. Gneiting, A.E. Raftery, Strictly proper scoring rules, prediction, and estimation, *J. Amer. Statist. Assoc.* 102 (477) (2007) 359–378.
- [51] J. Quinero-Candela, C.E. Rasmussen, F. Sinz, O. Bousquet, B. Schölkopf, Evaluating predictive uncertainty challenge, in: Machine Learning Challenges Workshop, Springer, 2005, pp. 1–27.
- [52] J. Demšar, Statistical comparisons of classifiers over multiple data sets, *J. Mach. Learn. Res.* 7 (2006) 1–30.
- [53] Y. Sun, C. Guo, Y. Li, React: Out-of-distribution detection with rectified activations, *Adv. Neural Inf. Process. Syst.* 34 (2021).
- [54] C.-K. Yeh, W.-C. Wu, W.-J. Ko, Y.-C.F. Wang, Learning deep latent space for multi-label classification, in: Thirty-First AAAI Conference on Artificial Intelligence, 2017.