

Improving Calibration by Relating Focal Loss, Temperature Scaling, and Properness

Viacheslav Komisarenko^{a,*} and Meelis Kull^a

^aInstitute of Computer Science, University of Tartu, Estonia

Abstract. Proper losses such as cross-entropy incentivize classifiers to produce class probabilities that are well-calibrated on the training data. Due to the generalization gap, these classifiers tend to become overconfident on the test data, mandating calibration methods such as temperature scaling. The focal loss is not proper, but training with it has been shown to often result in classifiers that are better calibrated on test data. Our first contribution is a simple explanation about why focal loss training often leads to better calibration than cross-entropy training. For this, we prove that focal loss can be decomposed into a confidence-raising transformation and a proper loss. This is why focal loss pushes the model to provide underconfident predictions on the training data, resulting in being better calibrated on the test data, due to the generalization gap. Secondly, we reveal a strong connection between temperature scaling and focal loss through its confidence-raising transformation, which we refer to as the focal calibration map. Thirdly, we propose focal temperature scaling - a new post-hoc calibration method combining focal calibration and temperature scaling. Our experiments on three image classification datasets demonstrate that focal temperature scaling outperforms standard temperature scaling.

1 Introduction

Modern machine learning systems demonstrate remarkable accuracy, often surpassing human capabilities in classification problems within textual and image domains. However, these systems still lack reliable confidence estimation in supporting their decisions, a critical issue in applications where the cost of mistakes is substantial, such as medical diagnostics, autonomous vehicle navigation, and financial decision-making.

Neural classification models typically output a real-valued logits (scores) vector. The higher the logit is, the higher the model confidence that the corresponding class should be predicted as positive. Quantifying model confidence in probabilities, not logits, is often more convenient.

The softmax transformation is traditionally applied on top of logits to obtain predicted probabilities. Still, alternative approaches exist, including entmax and sparsemax [14], designed to control the sparseness of vector representations in natural language processing, and tempered exponential functions derived under convex duality assumption [1].

Accuracy and Area Under the Curve (AUC) assess a model's discriminative power, while the Expected Calibration Error (ECE) evaluates calibration, measuring how well model confidence aligns

with true class probabilities. ECE is essentially the average absolute difference between predicted probabilities and empirical accuracies across probability bins. Due to finite sample sizes, this expected value is approximated using binning, which inherently introduces approximation errors.

The cross-entropy loss function is a common choice for training a classification model. Originating from information theory, cross-entropy measures the dissimilarity between two probability distributions, specifically, the true distribution of class labels and the predictions provided by a model. However, a notable limitation of cross-entropy trained models is their tendency to be overconfident during test time, wherein the predicted probabilities frequently exceed the actual frequencies of the corresponding classes [5].

Post-hoc calibration, a transformation on top of predicted logits or probabilities, is usually applied to improve the calibration of trained models [16]. The calibration method typically involves selecting from a set of potential parametrized transformations, which are applied to adjust the model's output. This selection process is usually performed using a validation set, which helps in identifying the most effective transformation to enhance the model's calibration by adjusting the predicted probabilities to align unknown perfectly calibrated probabilities more accurately while keeping the trained model unchanged.

Temperature scaling [5] is a notably straightforward yet effective single-parameter calibration method. The temperature parameter T adjusts the model's confidence by dividing the logits before the softmax, i.e. the last-layer activation function. Due to its simplicity and robust performance, it is often favoured over more complex calibration methods, such as matrix and vector scaling [5], Bayesian binning [13], and Dirichlet calibration [10]. Calibration methods in deep learning have been reviewed by Wang [17].

Cross-entropy is a proper loss function, meaning that it is minimized when the predicted probabilities match the true underlying probabilities of the classes, thereby incentivizing calibrated probability forecasts on the training set. Focal loss, parametrized by positive γ , is a recent alternative to cross-entropy, with enhanced calibration on the test set, and comparable or even higher accuracy [11, 12, 18]. Focal loss is not proper, but somehow, it outperforms proper losses such as cross-entropy.

We were driven by the question of what is special in the focal loss expression that explains its calibration performance, why it could be beneficial to intentionally push predicted probabilities away from the true underlying class probabilities (because that's what focal loss does) and still receive better calibration on the test set.

In this study, our contributions are threefold:

* Corresponding Author. Email: viacheslav.komisarenko@ut.ee.

- We deconstruct the focal loss into a composition of a confidence-raising transformation (which we call the focal calibration map) and a proper loss. We show how this decomposition explains why focal loss encourages models to generate under-confident predictions on the training data, yielding better calibration on test data.
- We demonstrate a notable link between temperature scaling and focal calibration maps. More specifically, for the binary case, we show that the focal calibration map parametrized with γ could be bounded between temperature scaling maps with $\frac{1}{T} = \gamma + 1$ and $\frac{1}{T} = \gamma + 1 - \frac{\log(\gamma+1)}{2}$ for all possible logits. For the multiclass case, we show that the focal calibration map for any $\gamma > 0$ increases the model confidence, similar to temperature scaling with $T < 1$. Moreover, we conducted an experiment showing that in the binary case, the difference between focal calibration and temperature scaling is almost neglectable when their parameters are related via the expression $\frac{1}{T} \approx 0.95 \cdot \gamma + 0.85$. For three- and four-dimensional cases, the closest focal calibration and temperature scaling transformations have close to the linear relationship of parameters γ and $\frac{1}{T}$ but with higher approximation error than in binary classification.
- We introduce focal temperature scaling — a novel post-hoc calibration technique that composes focal calibration with temperature scaling. Our experiments on three image classification datasets and several learning strategies (mainly focal or cross-entropy-based and calibration-oriented) show that focal temperature scaling enhances the model’s calibration compared to standard temperature scaling.

2 Focal loss

The cross-entropy loss function is a default choice for many machine learning classification problems. It has roots in maximum likelihood estimation and is a well-studied loss function with the simple expression $L_{CE}(p, y) = -\sum_{i=1}^n y_i \cdot \log(p_i)$, where p and y denote the predicted probability vector and the true class vector, respectively, \log hereafter is the natural logarithm and n is the number of classes.

It is a proper loss function, which implies that the model will yield almost perfect calibration on a training set. However, on a validation and test set, cross-entropy results are often overconfident [5].

Focal loss [11], which could be considered a modification on top of the cross-entropy loss, has gained widespread acceptance for classification problems. It is defined as $L_{FL}(p, y) = -\sum_{i=1}^n y_i \cdot (1 - p_i)^\gamma \log(p_i)$, where the multiplier $(1 - p)^\gamma$ serves as a scaling factor that for $\gamma > 1$ reduces the impact of easier examples (where p is close to 1) while increasing the influence of difficult examples (where p is far from 1) on the loss. This adjustment helps to focus the model’s learning on harder instances that are incorrectly classified, enhancing its predictive accuracy on more challenging cases.

While the focal loss is improper, it was shown to be classification-calibrated [2], implying that the minimizer of the focal loss always has the best possible accuracy (Bayes-optimal classifier). Moreover, the authors proved the existence of the closed-form transformation that recovers true class probability from the focal minimizer predictions. In practice, reaching perfect focal loss minimization is almost impossible due to reasons such as data noise and imperfect learning algorithms. Still, the classification-calibration gives a theoretical justification for focal loss usage when optimizing for accuracy.

Convexity is an essential property of most commonly used loss functions such as cross-entropy and Brier score. Convexity is helpful as many convex optimization methods and properties could be used to enhance the model training process. However, some methods as-

sume, in fact, convexity with respect to input features, which is very hard to achieve in deep learning settings. At the same time, losses are generally called convex if they are convex with respect to the predicted probability vector. In the following proposition, we show that the focal loss is convex with respect to the predicted probability vector.

Proposition 1. *Focal loss is convex with respect to the first (prediction) argument p for all $\gamma \geq 0$.*

Proof. The proof is based on an analysis of the sign of focal loss’s second derivative and is presented in the supplementary material [8]. \square

The hyperparameter $\gamma \geq 0$ is usually chosen on a validation set, while the authors of the original paper suggested value $\gamma = 2$ to improve accuracy. Model calibration is noticeably impacted by the choice of γ [18].

In their comprehensive study, Mukhoti et al. [12] conducted an extensive experimental comparison involving cross-entropy and focal losses, along with the Brier score and Label smoothing. Their findings indicate that models trained using focal loss have improved calibration compared to those trained with cross-entropy. They recommend an optimal γ value of 3 based on experiments conducted within $\gamma \in \{1, 2, 3\}$. Additionally, they observed that models trained with focal loss generally demonstrate noticeably enhanced performance in terms of area under the curve (AUC) and log-loss. The authors also noted that minimizing ECE on the validation set to determine optimal temperature T yielded better calibration performance than minimizing the log-loss.

Temperature scaling can further improve the calibration of models trained with focal loss. The improvement is often minor as the optimal temperature is frequently close to 1 [12]. This contrasts with models trained using cross-entropy, which significantly benefit from temperature scaling with an optimal temperature often reported to be in the range of [1.5, 3] [5].

Interestingly, optimal γ for log-loss minimization may differ from the value that minimizes the ECE score or accuracy [12, 18]. Furthermore, γ is also extremely dataset-specific, e.g. $\gamma = 10$ was the best on both log-loss and ECE in [18].

In this paper, we do not specifically address the issue of class imbalance. However, it is worth noting that focal loss has been reported to perform effectively in scenarios characterized by class imbalance also [18].

3 Proper deconstruction of the focal loss

3.1 Motivation

Proper losses are minimized by the predictions equal to ground-truth distribution. To illustrate it, let us consider a binary classification example. Let there be a point in feature space, meaning a fixed feature combination, which occurs 100 times in a training set with 80 actual positives and 20 actual negatives. What should the model predict for this feature combination? The model will output the number that minimizes the loss, meaning $\hat{p} = \operatorname{argmin}_{p \in [0,1]} (\sum_{i=1}^{20} L(p, 0) + \sum_{i=21}^{100} L(p, 1))$. The great thing about proper losses is that they will always predict $\hat{p} = 0.8$ for this situation, which is an intuitively right thing to do. It makes properness a desirable property for loss functions.

Optimizing for the improper focal loss will result in the prediction of a lower value of around 0.62 for $\gamma = 2$. Moreover, by varying

parameter $\gamma > 0$, we could receive any value between 0.8 and 0.5 (decision boundary). Despite this counterintuitive decision, the models trained with focal loss produce quite well-calibrated probabilities and high accuracy on the test set, in contrast with proper losses such as cross-entropy.

We aim to identify and locate part of the focal loss expression that stands for pushing predicted probabilities away from perfectly calibrated ones. We believe that the remaining part of the focal loss expression could be connected with properness as during test time, models trained with focal loss produces quite well-calibrated predictions.

3.2 Decomposition in the binary case

Let p be the instance's ground-truth probability of being positive, and q be the model's predicted probability for class 1. Let the loss on each instance be defined as $L(q, y)$, where $y \in \{0, 1\}$. The conditional risk, which is an expected loss under the ground-truth distribution, is defined as $R(p, q) = p \cdot L(q, 1) + (1 - p) \cdot L(1 - q, 0)$. The loss is called (strictly) proper if for any ground-truth distribution, the associated conditional risk is (uniquely) minimized by the prediction equal to the ground truth:

$$\forall p \in [0, 1] : \operatorname{argmin}_{q \in [0, 1]} R(p, q) = p$$

In the following proposition, we deconstruct the binary focal loss into a composition of a proper loss and a monotonic probability transformation (which we refer to as the focal calibration map):

Proposition 2. *Let $L(q, y)$ be a binary focal loss parametrized with some $\gamma > 0$. Then it can be decomposed into a bijective function $\hat{p} : [0, 1] \rightarrow [0, 1]$ (which could be seen as a fixed calibration map) and a proper loss $L^* : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R}^+$ such that $L^*(q, y) = L(\hat{p}^{-1}(q), y)$ and \hat{p} is defined as:*

$$\hat{p}(q) = \frac{1}{1 + \left(\frac{1-q}{q}\right)^\gamma \frac{(1-q) - \gamma q \cdot \log(q)}{q - \gamma(1-q) \cdot \log(1-q)}}$$

Proof. The full proof is presented in the supplementary material [8], first, showing that $\hat{p}(\cdot)$ is a bijection and then showing the properness of L^* . The proof is using known facts from the theory of binary composite losses [15]. \square

This decomposition suggests that training a focal loss is, in fact, equivalent to training a specific proper loss, which is applied on top of an additional fixed calibration layer in the network. The decomposition does not change the training process because neither the predicted probabilities (before the derived calibration layer) nor the loss value is changed, which implies that backward gradients will remain the same.

The proper part of the binary focal loss decomposition could be presented via the inverse (which exists because $\hat{p}_\gamma(q)$ was shown to be a bijection in the proof of Proposition 2) of the focal calibration map:

$$L^*(q, y = 1) = -(1 - \hat{p}^{-1}(q))^\gamma \cdot \log(\hat{p}^{-1}(q))$$

Due to the complex composition of exponential, logarithm and power functions in the focal calibration expression $\hat{p}_\gamma(q)$, its inverse, and consequently, the "properized" binary focal loss L^* , cannot be represented as a closed-form expression. Still, it is possible to approximate and tabulate these values numerically.

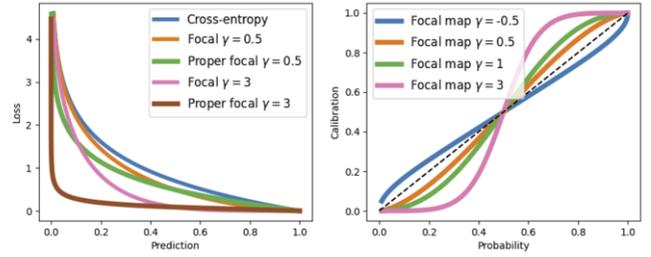


Figure 1. Left: proper part of the focal loss compared to standard focal loss with parameters 0.5 and 3 and cross-entropy on the true positive instance. Right: Binary focal calibration for different γ parameters for all possible predicted probabilities.

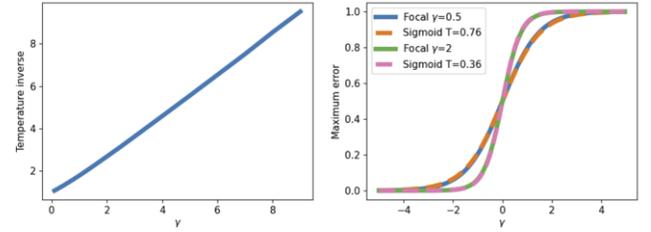


Figure 2. Left: relationship between focal calibration with parameter γ and temperature scaling parameter $\frac{1}{T}$ fitted to minimize the maximum deviation (over all logits) of these two transformations. Right: focal calibration with parameters 0.5 and 3 visualized in logit scale together with closest temperature scaling maps (with $T = 0.76$ and $T = 0.36$ correspondingly).

The binary focal calibration map and the proper part of the focal loss (on the actual positive instance) are shown in Fig. 1. It could be seen that the focal calibration has a traditional sigmoid-like shape for $\gamma > 0$ similarly to temperature scaling. Technically, negative γ could also be plugged into the calibration map, yielding an "inverse"-sigmoid shape on the opposite side of the diagonal shown on the right side in Fig. 1. The proper part of the focal loss compared with a regular focal loss with the same γ has lower loss values for predictions lower than 0.5 and higher loss otherwise, which could also be inferred from the calibration map shape.

As the focal calibration map has a similar to temperature scaling sigmoid-like shape, we approximated the temperature scaling with the focal calibration map with the parameter γ , minimising the maximum absolute difference between these calibration maps. Essentially, we ran a Linear Regression to fit these two transformations. It can be seen in Fig. 2 with the linear fit $\frac{1}{T} \approx 0.95 \cdot \gamma + 0.85$ between the inverse temperature and γ parameter. Moreover, the maximum approximation error is marginal, lower than $1e^{-3}$ in the probability scale.

Given a surprisingly small approximation error of the linear fit, we tried to find lower and upper bounds for the focal calibration map using temperature scaling transformations, see the following Proposition 3:

Proposition 3. *Let $FC(s) = \hat{p}_\gamma\left(\frac{1}{1+e^{-s}}\right)$ be a focal calibration function applied on top of sigmoid with a logit s . Then, the focal calibration could be bounded between two temperature scaling maps with $T = \frac{1}{\gamma+1}$ and $T = \frac{1}{\gamma+1 - \frac{\log(\gamma+1)}{2}}$ such that*

$$\forall s < 0 : \frac{1}{1 + e^{-\frac{s}{\gamma+1}}} > FC(s) > \frac{1}{1 + e^{-\frac{s}{\gamma+1 - \frac{\log(\gamma+1)}{2}}}}$$

$$\forall s \geq 0 : \frac{1}{1 + e^{-\frac{s}{\gamma+1}}} < FC(s) < \frac{1}{1 + e^{-\frac{s}{\gamma+1 - \frac{\log(\gamma+1)}{2}}}}$$

Proof. The proof is provided in the supplementary material [8]. \square

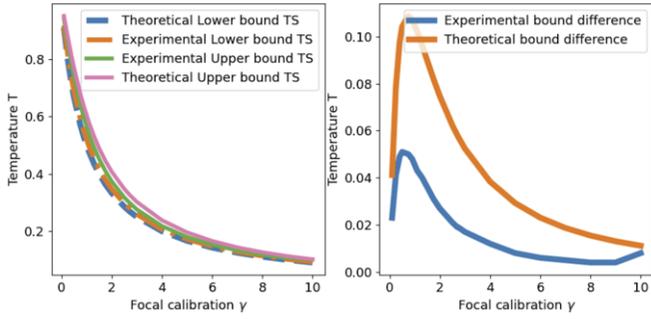


Figure 3. Left: theoretical and experimental bounds for focal calibration with temperature scaling maps. Right: theoretical and experimental width of the bounds for different γ

To evaluate the tightness of the derived bounds, we conduct an experimental assessment. For each value of γ , we systematically explore all possible temperatures $T > 0$ in increments of 0.001. We assess the focal calibration for all logits within the range $(-20, 20)$, using a step size of 0.001, to determine if it resides between two temperature scaling maps. Through this method, we identify and select the maximal lower bound and the minimal upper bound, thereby establishing the tightest experimental bounds achievable.

The theoretical and experimental lower and upper temperature scaling bounds for focal calibration are shown in Fig. 3. It could be seen that the experimental bounds are considerably tighter than the theoretical, implying that the theoretical assumptions could be strengthened to suggest even tighter theoretical bounds, at least for logits in the range $(-20, 20)$. For example, the theoretical result suggests the focal calibration with $\gamma = 4$ for all possible logits are bounded between temperature scaling with $T = 0.2$ and $T = 0.238$. If we double-check this result numerically, we could suggest even tighter bounds $T = 0.206$ and $T = 0.218$.

3.3 Decomposition in the multiclass case

Let $p = (p_1, \dots, p_n)$ be the ground truth vector and $q = (q_1, \dots, q_n)$ the prediction vector. We could consider a conditional risk $R(p, q)$, which is an expected loss under the ground-truth labels distribution p : $R(p, q) = \sum_{i=1}^n p_i \cdot L(q, y_i = 1)$. The properness condition could be written as

$$\operatorname{argmin}_{q \in \Delta^n} R(p, q) = p$$

Here Δ^n is a n -dimensional probability simplex, meaning that q should be a valid probability distribution.

We decompose the multiclass focal loss into a monotonic probability transformation (focal calibration map) and a proper loss. The derived focal calibration map is observed to equal with the transformation that maps the outputs from focal minimizers to the true class probabilities [2].

Proposition 4. Let $L(q, y)$ be a multiclass focal loss parametrized with some $\gamma > 0$. Then, it can be deconstructed into a composition of a bijective function $\hat{p}(q)$ and a proper loss $L^*(q, y)$ such that:

$$\hat{p}_j(q_1, \dots, q_n) = \frac{\frac{1}{(1-q_j)^\gamma \cdot \left(\frac{\gamma \cdot \log(q_j)}{1-q_j} - \frac{1}{q_j}\right)}}{\sum_{k=1}^n \frac{1}{(1-q_k)^\gamma \cdot \left(\frac{\gamma \cdot \log(q_k)}{1-q_k} - \frac{1}{q_k}\right)}} \quad \forall j = 1..n \quad (1)$$

$$L^*(q, y) = L(\hat{p}_1^{-1}(q_1, \dots, q_n), \dots, \hat{p}_n^{-1}(q_1, \dots, q_n))$$

Proof. The proof is written in the supplementary material [8]. \square

The proper part of the multiclass focal loss is also impossible to write as a closed-form expression. However, we could tabulate and visualize this function using focal calibration and standard focal loss expression: $L^*(q_1, \dots, q_n) = L(\hat{p}^{-1}(q_1, \dots, q_n))$.

Let us consider a three-dimensional probability simplex projected onto a two-dimensional triangle to understand the derived focal calibration map and the proper part of the focal loss. Each vertex of this triangle represents a scenario in which the entire probability mass is allocated to a specific class (see Fig. 4).

To get more insights about the geometric shape of the proper part of the multiclass focal loss, we visualized isolines for three selected percentiles over all possible focal proper part (for $\gamma = 1$ and $\gamma = 3$) loss values compared with the Brier score and cross-entropy in Fig. 4. We presented the conditional risk associated with a specific ground-truth vector $p = (0.55, 0.3, 0.15)$. Due to properness, conditional risks for all considered losses are minimized at $q = (0.55, 0.3, 0.15)$. Notably, the Brier score displays circular isolines, reflective of its quadratic loss expression. The isolines for cross-entropy exhibit a more oval-like configuration. In contrast, the isolines for the proper part of the focal loss assume a shape more akin to a curved triangle, a trend that becomes more pronounced with increasing values of γ .

The focal calibration and corresponding closest temperature scaling maps visualized as a directional arrows map projected into a two-dimensional triangle for two different parameters γ is shown in Fig. 5. Each arrow originates from a point representing an initial probability input and ends at a point corresponding to the output of the focal calibration (temperature scaling) transformation. It could be seen that arrow directions are all from the centre towards corners, which implies that the predicted distribution becomes sharper, meaning that the highest predicted probability becomes higher while all others - become lower. Moreover, the higher γ is, the sharper the effect becomes.

This confidence-raising effect is similar to a temperature scaling calibration with a temperature parameter lower than one $T < 1$, meaning the highest predicted class probability is mapped into an even higher probability.

Next we prove this as a proposition:

Proposition 5. Let the model's output for a test instance be a predicted probability vector $q = (q_1, \dots, q_n)$ such that the predicted class is j : $j = \operatorname{argmax}_{1 \leq i \leq n} (q_i)$. Then, if we apply focal calibration to the prediction q , the j -th coordinate of the resulting vector will be higher than or equal to all coordinates of the initial prediction q : $\hat{p}(q)_j \geq \max_{1 \leq i \leq n} (q_i) = q_j$.

Proof. The proof is listed in the supplementary material [8]. \square

Similarly to the binary case, we quantify the similarity between multiclass focal calibration and temperature scaling by finding the temperature's inverse $\frac{1}{T}$ for each γ (considering the step of 0.001 for both T and γ over the range $(0, 10)$ and step of 0.01 for each logit in range $(-5, 5)$) that minimizes the largest difference between focal calibration and temperature scaling transformations over a three-dimensional probability simplex.

The relationship between matching temperature inverse $\frac{1}{T}$ and γ is shown in Fig. 6. We could see a close to a linear trend that could be approximated with a function $\frac{1}{T} \approx 0.64 \cdot \gamma + 0.91$. For all positive γ , the temperature is lower than one, which verifies that the focal calibration has a sharpness effect. Moreover, lower temperature, which means higher sharpness, corresponds to higher γ . We should note that temperature scaling is not equivalent to focal calibration: the approximation error in a three-dimensional case rises from 0 (identity

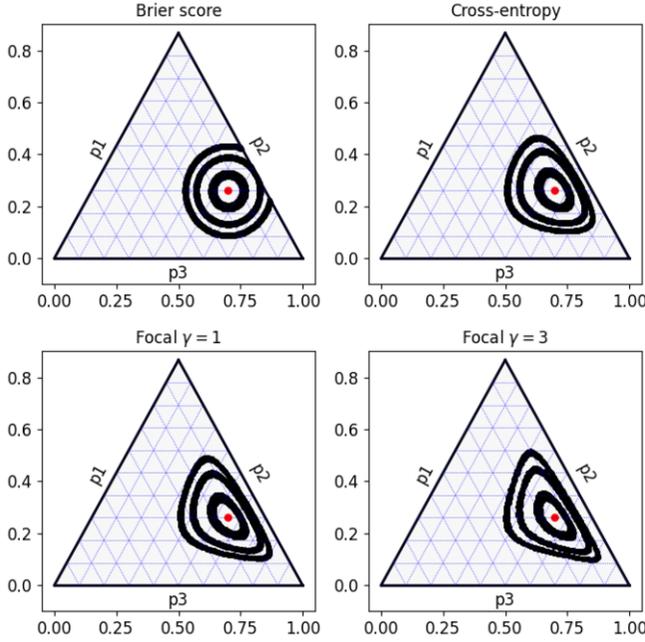


Figure 4. Brier score, cross-entropy and properized focal loss ($\gamma = 1, 3$) conditional risk isolines (defined by loss percentiles 3%, 12%, 20%) for ground truth probability $p = (0.55, 0.3, 0.15)$

transformation for $\gamma = 0$) to 8% (of the probability scale) for $\gamma = 9$. This means that the focal calibration is more complex and cannot be reliably approximated with a single simple scaling parameter of temperature scaling.

We considered the difference between focal calibration for $\gamma = [0.5, 3]$ and the correspondent closest temperature scaling mapping for all points of the probability simplex shown in Fig. 7. The highest difference between these two transformations could be seen near the corners, while it non-uniformly decreases towards the centre.

We also considered this approximation in a four-dimensional case and received a similar linear dependency trend: $\frac{1}{T} \approx 0.47 \cdot \gamma + 0.85$ with similar maximum approximation error of up to 8% on probability scale for larger γ as in the three-dimensional scale.

4 Proper decomposition implications

4.1 Generalization from the training set to the test set

The built-in focal calibration properties could hint towards the calibration performance of the focal loss function. For example, proper losses, such as cross-entropy, have almost perfect calibration performance on the training set but are generally overconfident on the validation set [5]. A typical level of overconfidence could approximately be quantified with a temperature of 2-2.5 [5], which is an approximate temperature needed to achieve calibration on a validation set via temperature scaling.

It hints towards an idea of the dataset-specific "generalization compensation", meaning no matter what loss function was used for training, a post-hoc calibration on a validation set is crucial for compensating the generalization effect from training to validation set.

If we assume that the focal loss has a similar "generalization compensation" as the cross-entropy and recall that the focal calibration map is close to temperature scaling with a temperature of about 0.4-0.5 for commonly used γ parameters, then the composition of temperature during training (around 0.4-0.5 for focal calibration) and "generalization compensation" (roughly quantified with temperature

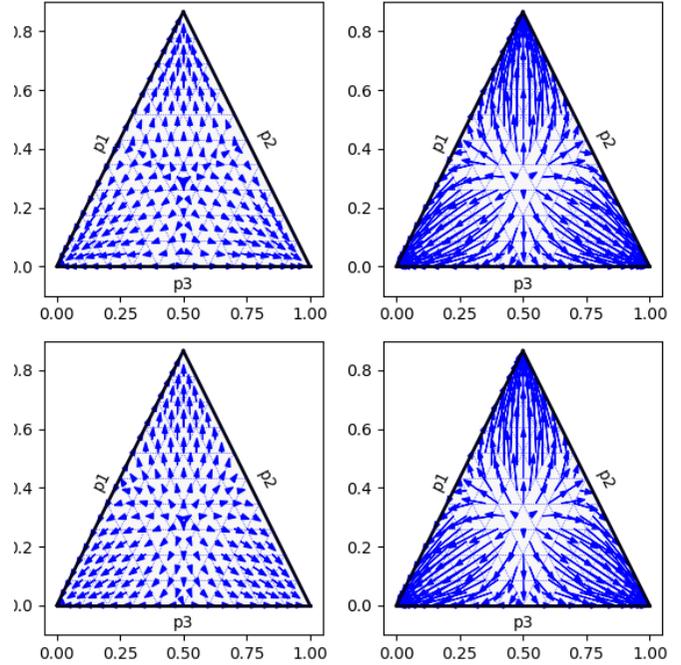


Figure 5. Focal calibration for Upper Left: $\gamma = 1$, Upper Right: $\gamma = 3$ parameters and temperature scaling maps for corresponding closest parameters: Bottom Left: $T = 0.81$, Bottom Right: $T = 0.46$ visualized as directional arrows over a uniform grid of three-dimensional probability simplex points. The focal calibration and temperature scaling maps the start of each arrow to its end.

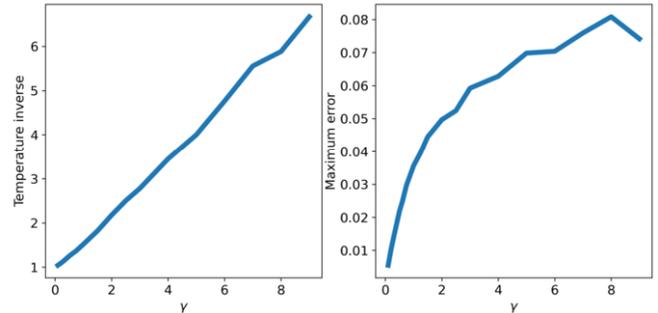


Figure 6. Left: relationship between focal calibration γ and temperature scaling parameter $\frac{1}{T}$ chosen to minimize absolute deviation of these calibration maps for three-dimensional case. Right: dependency of maximum approximation error of these calibration maps on a focal parameter γ .

2-2.5) gives an overall temperature on the validation set about 1 (as multiplication of temperature 0.4-0.5 and 2-2.5), which is precisely what the temperature is for focal loss post-hoc calibration. This explains why focal loss tends to be quite well calibrated on the test set.

4.2 Focal temperature scaling

We suggest a new post-hoc calibration method called focal temperature scaling, which consists of composing focal calibration with parameter γ_{ev} and temperature scaling with temperature T . More specifically, the method optimizes the pair of parameters (γ_{ev}, T) on a validation set. First, instance class scores s_1, \dots, s_n are converted into temperature-scaled softmax probabilities q_1, \dots, q_n . These probabilities are then transformed with a focal calibration map as defined by Eq. (1) in Proposition 4 parameterized by γ_{ev} , resulting in predicted probabilities $\hat{p}_1, \dots, \hat{p}_n$. The goal is to produce predictions that

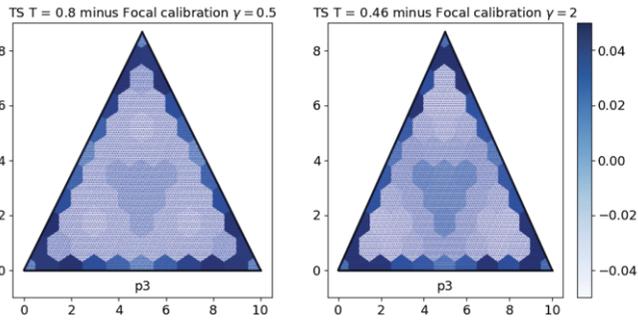


Figure 7. Temperature scaling minus focal calibration $FC(s) - TS(s)$ for the selected closest γ, T pairs $\{(0.5, 0.8), (2, 0.46)\}$ heatmap for all points of the three-dimensional probability simplex.

minimize the metric of interest (e.g. ECE). As both transformations directly impact the model’s confidence and are sufficiently different in a multiclass case, the focal temperature scaling should improve the model’s calibration to a larger extent than standard temperature scaling. We do not limit the method to the models trained with a focal loss with parameter γ_{tr} and, in case the training was performed with focal loss, do not require training and post-hoc calibration parameters to be equal: $\gamma_{ev} \neq \gamma_{tr}$.

5 Experiments

In the experiments, we aim to verify whether applying focal temperature scaling leads to enhanced calibration compared to standard temperature scaling and suggest guidelines for practitioners.

Several hypotheses will be considered as the main focus of experiments:

- Does the focal temperature scaling lead to improved calibration of the model trained with focal loss?
- Does the focal temperature scaling perform well when applied for other than focal losses and still receive an improved calibration compared to standard calibration methods?
- Is there any experimentally observed relationship between temperature scaling optimal parameter T and focal loss train and calibration parameters γ_{tr}, γ_{ev} ?

5.1 Experiment settings

We used an NVIDIA Tesla V100 GPU with 16 GB of VRAM for our experiments.

We consider CIFAR-10, CIFAR-100 [9] and TinyImageNet [3] datasets with Resnet-50 architecture [6]. We used 45000 / 5000 / 10000 images split as train, validation and test sets for CIFAR-10 and CIFAR-100 datasets. For TinyImageNet, we used 90000 / 10000 image split for train and validation sets and the TinyImageNet validation set as our test set. Models were trained on the training set, and the validation set was used to find an optimal pair of temperature scaling and focal calibration parameters (T, γ_{ev}) . A test set was used to measure the final model performance. The main evaluation metrics of interest were ECE, NLL and error rate.

The training setup was chosen to be mainly consistent with the studies [12, 4]. For CIFAR-10 and CIFAR-100 datasets, we started training with the SGD optimizer with a momentum of 0.9, weight decay $5e^{-4}$ and a learning rate of 0.1, which decreased to 0.01 after 150 epochs and 0.001 after 250 epochs until epoch 350. For TinyImageNet, we used the same optimizer but with a learning rate of 0.1 for

the first 40 epochs, 0.01 for the next 20 epochs and 0.001 for the last 40 epochs. The batch size was 128 for CIFAR datasets and 64 for TinyImageNet. Results were evaluated on the last epoch following [12, 4].

We trained a model with cross-entropy loss, and more models with the focal loss with $\gamma_{train} \in \{1, 2, 3, 5, 7\}$, respectively. Also, we trained sample-dependent focal loss (FLSD-53) [12] and AdaFocal [4] for comparison with recent studies.

During evaluation, we applied focal temperature scaling calibration on top of softmax predicted probabilities for all trained models. We iterated over a range of $[-0.5, -0.25, 0.05, 0.25, 0.37, 0.5, 0.75, 1, 5]$ for focal calibration parameter γ_{ev} and over a range from 0.01 to 5 with a step of 0.01 for temperature parameter T . When presenting results for ECE, we selected the optimal parameters γ_{ev}, T based on the validation set ECE minimization. When considering results for log-loss, the hyperparameter choice criterion was log-loss. The results were compared with standard temperature scaling when only the parameter T is selected over the same values grid and same choice criteria. An equal mass ECE measured on 15 bins was used for calibration estimate.

All code used for the implementation and experiments can be accessed in our GitHub repository [7].

5.2 Results of experiments

Results for all datasets are shown in Table 1. Results are grouped per dataset. Next to each baseline approach row (cross-entropy, focal loss with different training parameters, FLSD-53, Adafocal), we put a row with the corresponding optimal parameter γ_{ev} of the focal temperature scaling that attempts to improve the baseline approach. For each evaluation metric (Accuracy, Log-loss, ECE), we report the mean score over 5 random seeds and the standard deviation, separated with \pm sign. In the brackets, we report the corresponding optimal temperature.

The results show that focal temperature scaling could noticeably improve ECE among all baseline models while keeping the same accuracy and occasionally marginally losing on log-loss compared to standard temperature scaling. The optimal temperature generally decreases with the growth of the focal training parameter γ_{tr} . Models trained with higher γ_{tr} have generally better ECE but worse accuracy. Also, despite our results being mainly in line with the reported previously results [12, 4], we could not replicate the similar ECE score for models trained with AdaFocal despite our efforts to follow the authors’ experiment guidelines closely [4] through the description provided in the original papers, supplementary materials and communication with the authors. We suspect that the reason may be in minor experiment settings, such as the exact value of the weight decay parameter. Still, the ECE scores of these methods could be improved further by applying focal temperature scaling instead of standard temperature scaling.

Moreover, we found the temperature choice criteria crucial for the experiment’s results. When evaluating ECE, we could gain an extra 20% relative improvement in average when optimizing for ECE on the validation set for the hyperparameter choice. For log-loss, the relative improvement is rather minor but consistent: around 1%.

The relationship between optimal temperature (chosen by minimizing ECE on a validation set) and focal calibration γ_{ev} parameter for a different focal loss γ_{tr} trained models is shown in Fig 8. It could be seen that the relationship is close to linear for all trained models. It implies that in practice, instead of fitting focal temperature scal-

ing parameters γ_{ev}, T over a 2-dimensional grid, we could iterate over fewer parameters by considering only parameters located on the line. To determine line slope and intercept, we still need to find at least two points of the line, which could be done by fixing two arbitrary γ_{ev} and finding corresponding optimal T . If the standard temperature scaling will iterate over a grid of m parameters, the focal temperature scaling requires $2 \cdot m$ iterations to find line coefficients and additional m iterations over the line - resulting in a total of $3 \cdot m$ considered parameters, which is not that much overhead compared to standard temperature scaling. Still, in our experiments, we performed a full 2-dimensional grid search to find the best parameters and did not compare the performance of the full and optimized parameters search. The experiments with the optimized parameter search and a full understanding of this phenomenon remain for future works.

6 Discussion and conclusions

We derived a way to look into focal loss training as proper loss training with a confidence-raising transformation applied on top of softmax probabilities. The presented decomposition into a proper loss

Table 1. Test set performance for focal loss (defined by γ_{tr}), sample-dependent focal loss FLSD-53 [12], AdaFocal with default parameters as in [4] and cross-entropy trained models with temperature scaling versus focal temperature scaling (defined by γ_{ev} and temperature). The CIFAR-100, CIFAR-10 and TinyImageNet datasets were used, and the results were averaged over 5 random seeds after applying temperature scaling. The mean result is reported together with the standard deviation after the \pm sign. The optimal temperature is reported in brackets; temperature choice criteria were log-loss for log-loss evaluation and ECE for ECE evaluation. The best result for each metric and dataset is highlighted in bold formatting.

| CIFAR-100 DATASET | | | | |
|-------------------------|-----------------------|-------------------------------|-------------------------------|--|
| APPROACH | ACCURACY | LOG-LOSS | ECE | |
| CROSS-ENTROPY | 77.6 \pm 0.6 | 0.88 \pm 0.02 (1.31) | 3.01 \pm 0.43 (1.45) | |
| + $\gamma_{ev} = -0.5$ | 77.6 \pm 0.6 | 0.86 \pm 0.02 (1.17) | 2.13 \pm 0.42 (1.35) | |
| FOCAL $\gamma_{tr} = 1$ | 77.7 \pm 0.3 | 0.83 \pm 0.01 (1.05) | 1.66 \pm 0.23 (1.15) | |
| + $\gamma_{ev} = -0.25$ | 77.7 \pm 0.3 | 0.82 \pm 0.01 (1.00) | 1.34 \pm 0.17 (1.10) | |
| FOCAL $\gamma_{tr} = 3$ | 77.3 \pm 0.5 | 0.81 \pm 0.02 (0.87) | 1.28 \pm 0.15 (0.91) | |
| + $\gamma_{ev} = 0.05$ | 77.3 \pm 0.5 | 0.82 \pm 0.02 (0.87) | 1.23 \pm 0.14 (0.95) | |
| FOCAL $\gamma_{tr} = 7$ | 76.3 \pm 0.5 | 0.83 \pm 0.01 (0.70) | 1.83 \pm 0.20 (0.65) | |
| + $\gamma_{ev} = 0.5$ | 76.3 \pm 0.5 | 0.83 \pm 0.01 (0.75) | 0.99 \pm 0.07 (0.75) | |
| FLSD-53 | 77.5 \pm 0.5 | 0.88 \pm 0.01 (1.20) | 1.89 \pm 0.18 (1.27) | |
| + $\gamma_{ev} = 0.25$ | 77.5 \pm 0.3 | 0.88 \pm 0.02 (1.05) | 1.68 \pm 0.19 (1.15) | |
| ADAFOCAL | 77.6 \pm 0.2 | 0.91 \pm 0.03 (1.40) | 2.96 \pm 0.22 (1.52) | |
| + $\gamma_{ev} = 0.25$ | 77.6 \pm 0.2 | 0.93 \pm 0.03 (1.48) | 2.71 \pm 0.17 (1.60) | |

| CIFAR-10 DATASET | | | | |
|-------------------------|-----------------------|-------------------------------|-------------------------------|--|
| APPROACH | ACCURACY | LOG-LOSS | ECE | |
| CROSS-ENTROPY | 95.0 \pm 0.1 | 0.16 \pm 0.00 (1.59) | 1.03 \pm 0.17 (1.72) | |
| + $\gamma_{ev} = 1$ | 95.0 \pm 0.1 | 0.17 \pm 0.01 (2.20) | 0.71 \pm 0.16 (2.36) | |
| FOCAL $\gamma_{tr} = 1$ | 95.0 \pm 0.1 | 0.17 \pm 0.01 (1.05) | 1.05 \pm 0.25 (1.13) | |
| + $\gamma_{ev} = 0.5$ | 95.0 \pm 0.1 | 0.17 \pm 0.01 (1.30) | 0.82 \pm 0.35 (1.37) | |
| FOCAL $\gamma_{tr} = 3$ | 94.3 \pm 0.3 | 0.19 \pm 0.01 (0.75) | 1.48 \pm 0.25 (0.77) | |
| + $\gamma_{ev} = 5$ | 94.3 \pm 0.3 | 0.20 \pm 0.01 (1.83) | 0.93 \pm 0.16 (1.87) | |
| FOCAL $\gamma_{tr} = 7$ | 93.1 \pm 0.1 | 0.23 \pm 0.01 (0.49) | 0.66 \pm 0.07 (0.44) | |
| + $\gamma_{ev} = 0.37$ | 93.1 \pm 0.1 | 0.22 \pm 0.01 (0.55) | 0.61 \pm 0.10 (0.52) | |
| FLSD-53 | 94.6 \pm 0.1 | 0.18 \pm 0.01 (1.40) | 1.30 \pm 0.13 (1.40) | |
| + $\gamma_{ev} = 1$ | 94.6 \pm 0.1 | 0.17 \pm 0.01 (1.30) | 1.23 \pm 0.20 (1.33) | |
| ADAFOCAL | 94.9 \pm 0.2 | 0.18 \pm 0.01 (1.58) | 1.80 \pm 0.17 (1.63) | |
| + $\gamma_{ev} = 5$ | 94.9 \pm 0.2 | 0.20 \pm 0.02 (3.70) | 0.95 \pm 0.10 (3.70) | |

| TINYIMAGENET DATASET | | | | |
|-------------------------|-----------------------|-------------------------------|-------------------------------|--|
| APPROACH | ACCURACY | LOG-LOSS | ECE | |
| CROSS-ENTROPY | 49.9 \pm 0.1 | 2.21 \pm 0.00 (1.35) | 5.57 \pm 0.31 (1.40) | |
| + $\gamma_{ev} = -0.5$ | 49.9 \pm 0.1 | 2.19 \pm 0.00 (1.30) | 3.66 \pm 0.18 (1.40) | |
| FOCAL $\gamma_{tr} = 1$ | 50.6 \pm 0.2 | 2.11 \pm 0.01 (1.10) | 3.27 \pm 0.13 (1.20) | |
| + $\gamma_{ev} = -0.5$ | 50.6 \pm 0.2 | 2.10 \pm 0.01 (1.10) | 1.86 \pm 0.12 (1.18) | |
| FOCAL $\gamma_{tr} = 3$ | 51.6 \pm 0.1 | 2.04 \pm 0.01 (0.95) | 2.21 \pm 0.14 (0.98) | |
| + $\gamma_{ev} = -0.25$ | 51.6 \pm 0.1 | 2.03 \pm 0.01 (0.95) | 1.63 \pm 0.05 (0.97) | |
| FOCAL $\gamma_{tr} = 7$ | 50.9 \pm 0.3 | 2.01 \pm 0.02 (0.85) | 1.01 \pm 0.02 (0.85) | |
| + $\gamma_{ev} = 0.05$ | 50.9 \pm 0.3 | 2.01 \pm 0.02 (0.85) | 0.96 \pm 0.00 (0.85) | |
| FLSD-53 | 52.1 \pm 0.1 | 2.02 \pm 0.01 (0.95) | 2.06 \pm 0.17 (0.98) | |
| + $\gamma_{ev} = -0.25$ | 52.1 \pm 0.1 | 2.02 \pm 0.01 (0.95) | 1.48 \pm 0.22 (0.95) | |
| ADAFOCAL | 51.6 \pm 0.3 | 2.07 \pm 0.03 (1.05) | 2.97 \pm 0.67 (1.10) | |
| + $\gamma_{ev} = -0.5$ | 51.6 \pm 0.3 | 2.07 \pm 0.03 (1.05) | 1.89 \pm 0.24 (1.09) | |

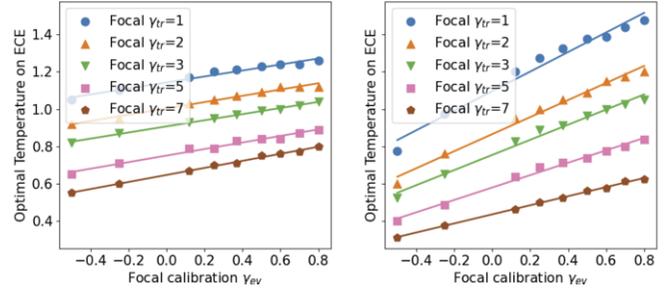


Figure 8. Relationship between optimal temperature T and the focal calibration parameter γ_{ev} chosen on ECE metric for focal loss trained models with different training parameters γ_{tr} . Left: CIFAR-100, right: CIFAR-10 dataset.

and a focal calibration could explain the recent successes of focal loss in calibration. The proper part of the decomposition applied on top of confidence-raised predictions pushes the probabilities to be slightly underconfident on the training set, which, due to generalization gap, leads to better calibration on the test set.

We discovered a surprising connection between the focal calibration and temperature scaling; specifically, focal calibration could be bounded by two temperature scaling transformations with parameters $T = \frac{1}{\gamma+1}$ and $T = \frac{1}{\gamma+1 - \frac{\log(\gamma+1)}{2}}$ for all logit values in the binary case. Moreover, for the multiclass case, we showed that focal calibration always increases the model's confidence, behaving similarly to temperature scaling with $T < 1$. We also conducted experiments showing a close to linear dependency of parameters $\frac{1}{T}$ and γ when minimizing deviation between these two transformations in three- and four-dimensional cases. While for the binary case, the difference between focal calibration and temperature scaling transformations for correspondent parameters is tiny, for the multiclass case, the deviation is larger and non-trivially varies across logit vectors. This implies that, for the multiclass case, composing focal and temperature scaling calibration into a single calibration method could be beneficial. We called this method focal temperature scaling, which is parametrized by focal calibration parameter γ_{ev} and temperature T .

We applied the suggested focal temperature scaling method for three multiclass image classification datasets trained with cross-entropy, focal loss with different parameters and recent calibration-oriented approaches FLSD-53 and AdaFocal. The experiment results suggest a consistent improvement in calibration over all trained models compared to standard temperature scaling while keeping the accuracy unchanged by design and log-loss being approximately the same.

Moreover, the experiments showed a linear trend between optimal γ_{ev} and T parameters over different training models and datasets, suggesting that instead of fitting hyperparameters over a 2-dimensional values grid, we could iterate γ_{ev}, T values only within the line (individual for each dataset and training method) without much performance loss. This implies that the focal temperature scaling could be applied when computational resources are limited. The detailed analysis of this dependency remains for the future work.

Acknowledgements

This work was supported by the Estonian Research Council under grant PRG1604, and by the Estonian Centre of Excellence in Artificial Intelligence (EXAI), funded by the Estonian Ministry of Education and Research.

References

- [1] E. Amid, M. K. Warmuth, R. Anil, and T. Koren. Robust Bi-Tempered Logistic Loss Based on Bregman Divergences. *Advances in Neural Information Processing Systems*, 32, 2019.
- [2] N. Charoenphakdee, J. Vongkulbhisal, N. Chairatanakul, and M. Sugiyama. On Focal Loss for Class-Posterior Probability Estimation: A Theoretical Perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5202–5211, 2021.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. IEEE, 2009.
- [4] A. Ghosh, T. Schaaf, and M. Gormley. AdaFocal: Calibration-aware Adaptive Focal Loss. *Advances in Neural Information Processing Systems*, 35:1583–1595, 2022.
- [5] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On Calibration of Modern Neural Networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR, 2017.
- [6] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [7] V. Komisarenko. Code Repository for Improving Calibration by Relating Focal Loss, Temperature Scaling, and Properness, 2024. URL https://github.com/slavikkom/focal_temperature_scaling. Accessed: 20.08.2024.
- [8] V. Komisarenko and M. Kull. Improving Calibration by Relating Focal Loss, Temperature Scaling, and Properness. *arXiv preprint arXiv:2408.11598*, 2024. URL <http://arxiv.org/abs/2408.11598>.
- [9] A. Krizhevsky, V. Nair, and G. Hinton. CIFAR-10 and CIFAR-100 datasets. URL: <https://www.cs.toronto.edu/kriz/cifar.html>, 6, 2009.
- [10] M. Kull, M. Perello Nieto, M. Kängsepp, T. Silva Filho, H. Song, and P. Flach. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with Dirichlet calibration. *Advances in Neural Information Processing Systems*, 32, 2019.
- [11] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal Loss for Dense Object Detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017.
- [12] J. Mukhoti, V. Kulharia, A. Sanyal, S. Golodetz, P. Torr, and P. Dokania. Calibrating Deep Neural Networks using Focal Loss. *Advances in Neural Information Processing Systems*, 33:15288–15299, 2020.
- [13] M. P. Naeini, G. Cooper, and M. Hauskrecht. Obtaining Well Calibrated Probabilities Using Bayesian Binning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- [14] B. Peters, V. Niculae, and A. F. Martins. Sparse Sequence-to-Sequence Models. *arXiv preprint arXiv:1905.05702*, 2019.
- [15] M. D. Reid and R. C. Williamson. Composite Binary Losses. *The Journal of Machine Learning Research*, 11:2387–2422, 2010.
- [16] T. Silva Filho, H. Song, M. Perello-Nieto, R. Santos-Rodriguez, M. Kull, and P. Flach. Classifier calibration: a survey on how to assess and improve predicted class probabilities. *Machine Learning*, pages 1–50, 2023.
- [17] C. Wang. Calibration in Deep Learning: A Survey of the State-of-the-Art. *arXiv preprint arXiv:2308.01222*, 2023.
- [18] C. Wang, J. Balazs, G. Szarvas, P. Ernst, L. Poddar, and P. Danchenko. Calibrating Imbalanced Classifiers with Focal Loss: An Empirical Study. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 145–153, 2022.