

# Post-hoc Uncertainty Calibration for Domain Drift Scenarios

Christian Tomani<sup>1</sup>, Sebastian Gruber<sup>3,4,5,\*</sup>, Muhammed Ebrar Erdem<sup>2</sup>,  
Daniel Cremers<sup>1</sup>, Florian Buettner<sup>2,3,4,5,\*</sup>

<sup>1</sup>Technical University of Munich

<sup>2</sup>Siemens AG <sup>3</sup> German Cancer Consortium

<sup>4</sup> German Cancer Research Center Heidelberg <sup>5</sup> Frankfurt University

{christian.tomani, cremers}@tum.de,

{sebastian.gruber, muhammed.erdem, buettner.florian}@siemens.com

## Abstract

We address the problem of uncertainty calibration. While standard deep neural networks typically yield uncalibrated predictions, calibrated confidence scores that are representative of the true likelihood of a prediction can be achieved using post-hoc calibration methods. However, to date, the focus of these approaches has been on in-domain calibration. Our contribution is two-fold. First, we show that existing post-hoc calibration methods yield highly over-confident predictions under domain shift. Second, we introduce a simple strategy where perturbations are applied to samples in the validation set before performing the post-hoc calibration step. In extensive experiments, we demonstrate that this perturbation step results in substantially better calibration under domain shift on a wide range of architectures and modelling tasks.

## 1. Introduction

### 1.1. Towards calibrated classifiers

Due to their high predictive power, deep neural networks are increasingly being used as part of decision making systems in real world applications. However, such systems require not only high accuracy, but also reliable and calibrated uncertainty estimates. A classifier is *calibrated*, if the confidence of predictions matches the probability of being correct for all confidence levels [4]. Especially in safety critical applications in medicine where average case performance is insufficient, but also in dynamically changing environments in industry, practitioners need to have access to reliable predictive uncertainty during the entire life-cycle of the model. This means confidence scores (or predictive uncertainty) should be well calibrated not only for in-domain

predictions, but also under gradual domain drift where the distribution of the input samples gradually changes from in-domain to truly out-of-distribution (OOD). Such domain drift scenarios commonly include changes in object backgrounds, rotations, and imaging viewpoints [1]; Fig. 1.

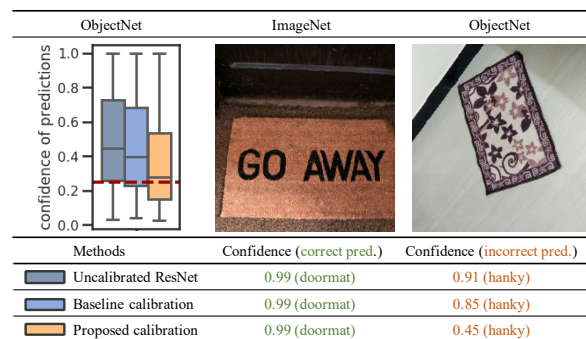


Figure 1: Re-calibrated neural networks make over-confident predictions under domain drift. Left: Only for our approach confidence scores match model accuracy (dashed red line) across all predictions; all other approaches are over-confident under domain shift. Middle: For Imagenet, all models make a correct prediction with high confidence (left). Right: Under domain drift (different viewpoint; Objectnet) all models make wrong predictions, but only our approach has a low confidence score reflecting model uncertainty.

Since deep neural networks typically only yield uncalibrated confidence scores, a variety of different post-hoc calibration approaches have been proposed [15, 4, 23, 22, 24]. These methods use the validation set to transform predictions returned by a trained neural network such that in-domain predictions are well calibrated. Such post-hoc uncertainty calibration approaches are particularly appealing since costly training of intrinsically uncertainty-aware neu-

\* Work done for Siemens AG

ral networks can be avoided.

Current efforts to systematically quantify the quality of predictive uncertainties have focused on assessing model calibration for in-domain predictions. Here, post-processing in form of temperature scaling has shown great promise and Guo *et al.* [4] illustrated that this approach yields well calibrated predictions for a wide range of model architectures. More recently, more complex combinations of parametric and non-parametric methods have been proposed [24]. However, little attention has been paid to uncertainty calibration under domain drift and no comprehensive analysis of the performance of post-hoc uncertainty calibration methods under domain drift exists.

## 1.2. Contribution

In this work we focus on the task of post-hoc uncertainty calibration under domain drift scenarios and make the following contributions:

- We first show that neural networks yield overconfident predictions under domain shift even after re-calibration using existing post-hoc calibrators.
- We generalise existing post-hoc calibration methods by transforming the validation set before performing the post-hoc calibration step.
- We demonstrate that our approach results in substantially better calibration under domain shift on a wide range of architectures and image data sets.

In addition to the contributions above, our code is made available at <https://github.com/tochris/calibration-domain-drift>.

## 2. Related work

In this section, we review existing approaches towards neural networks with calibrated predictive uncertainty. The focus of this work is on post-hoc calibration methods, which we review in detail. These approaches can broadly be divided into 2 categories: accuracy-preserving methods, where the ranking of confidence scores across classes remain unchanged and those where the ranking, and thus accuracy, can change. Other related work includes intrinsically uncertainty-aware neural networks and out-of-distribution detection methods.

### 2.1. Post-processing methods

A popular approach towards well-calibrated predictions are post-hoc calibration methods where a validation set, drawn from the generative distribution of the training data  $\pi(X, Y)$ , is used to rescale the outputs returned by a trained neural network such that in-domain predictions are well calibrated. A variety of parametric, as well as non-parametric

methods exist. We first review non-parametric methods that do not preserve accuracy. A simple but popular non-parametric post-processing approach is histogram binning [22]. In brief, all uncalibrated confidence scores  $\hat{P}_i$  are partitioned into  $M$  bins (with borders typically chosen such that either all bins are of equal size or contain the same number of samples). Next, a calibrated score  $Q_m$  that is determined by optimizing a bin-wise squared loss on the validation set, is assigned to each bin. For each new prediction, the uncalibrated confidence score  $\hat{P}_{pr}$  is then replaced by the calibrated score associated with the bin  $\hat{P}_{pr}$  falls into. Popular extensions to histogram binning are isotonic regression [23] and Bayesian Binning into Quantiles (BBQ) [12]. For isotonic regression (IR), uncalibrated confidence scores are divided into  $M$  intervals and a piecewise constant function  $f$  is fitted on the validation set to transform uncalibrated outputs to calibrated scores. BBQ is a Bayesian generalisation of histogram binning based on the concept of Bayesian model averaging.

In addition to these non-parametric approaches, also parametric alternatives for post-processing confidence scores exist. Platt scaling [15] is an approach for transforming the non-probabilistic outputs (logits)  $z_i \in \mathbb{R}$  of a binary classifier to calibrated confidence scores. More specifically, the logits are transformed to calibrated confidence scores  $\hat{Q}_i$  using logistic regression  $\hat{Q}_i = \sigma(az_i + b)$ , where  $\sigma$  is the sigmoid function and the two parameters  $a$  and  $b$  are fitted by optimising the negative log-likelihood of the validation set.

Guo *et al.* [4] have proposed Temperature Scaling (TS), a simple generalisation of Platt scaling to the multi-class case, where a single scalar parameter  $T$  is used to re-scale the logits of a trained neural network. In the case of  $C$ -class classification, the logits are a  $C$ -dimensional vector  $\mathbf{z}_i \in \mathbb{R}^C$ , which are typically transformed into confidence scores  $\hat{P}_i$  using the softmax function  $\sigma_{SM}$ . For temperature scaling, logits are rescaled with temperature  $T$  and transformed into calibrated confidence scores  $\hat{Q}_i$  using the softmax function as

$$\hat{Q}_i = \max_c \sigma_{SM}(\mathbf{z}_i/T)^{(c)} \quad (1)$$

$T$  is learned by minimizing the negative log likelihood of the validation set. In contrast to the non-parametric methods introduced above or other multi-class generalisations of Platt scaling such as vector scaling or matrix scaling, Temperature Scaling has the advantage that it does not change the accuracy of the trained neural network. Since re-scaling does not affect the ranking of the logits, also the maximum of the softmax function remains unchanged.

More recently, Zhang *et al.* [24] have proposed to combine to combine parametric methods with non-parametric methods, in particular they suggest it can be beneficial to perform IR after a TS step (TS-IR). In addition, they introduced an accuracy-preserving version of IR, termed

IRM and an ensemble version of temperature scaling, called ETS.

## 2.2. Intrinsically uncertainty-aware neural networks

A variety of approaches towards intrinsically uncertainty-aware neural networks exist, including probabilistic and non-probabilistic approaches. In recent years, a lot of research effort has been put into training Bayesian neural networks. Since exact inference is untractable, a range of approaches for approximate inference has been proposed, [3, 21]. A popular non-probabilistic alternative is Deep Ensembles [10], where an ensemble of networks is trained using adversarial examples, yielding smooth predictions with meaningful predictive uncertainty. Other non-probabilistic alternatives include, e.g., [17, 20]. While a comprehensive analysis of the performance of post-hoc calibration methods is missing for domain-drift scenarios, recently Ovadia *et al.* [13] have presented a first comprehensive evaluation of calibration under domain drift for intrinsically uncertainty-aware neural networks and have shown that the quality of predictive uncertainties, i.e. model calibration, decreases with increasing dataset shift, regardless of method.

## 2.3. Other related work

Orthogonal approaches have been proposed where trust scores and other measures for out-of-distribution detection are derived to detect truly OOD samples, often also based on trained networks and with access to a known OOD set [11, 8, 14]. However, rather than only detecting truly OOD samples, in this work, we are interested in calibrated confidence scores matching model accuracy at all stages of domain drift, from in-domain samples to truly OOD samples.

## 3. Problem setup and definitions

Let  $X \in \mathbb{R}^D$  and  $Y \in \{1, \dots, C\}$  be random variables that denote the  $D$ -dimensional input and labels in a classification task with  $C$  classes with a ground truth joint distribution  $\pi(X, Y) = \pi(Y|X)\pi(X)$ . The dataset  $\mathcal{D}$  consists of  $N$  i.i.d. samples  $D = \{(X_n, Y_n)\}_{n=1}^N$  drawn from  $\pi(X, Y)$ . Let  $h(X) = (\hat{Y}, \hat{P})$  be the output of a neural network classifier  $h$  predicting a class  $\hat{Y}$  and associated confidence  $\hat{P}$  based on  $X$ . Here, we are interested in the quality of predictive uncertainty (i.e. confidence scores  $\hat{P}$ ) not only on test data from the generative distribution of the training data  $\mathcal{D}$ ,  $\pi(X, Y)$ , but also under dataset shift, that is test data from a distribution  $\rho(X, Y) \neq \pi(X, Y)$ . More specifically, we investigate domain drift scenarios where the distribution of samples seen by a model gradually moves away from the

training distribution  $\pi(X, Y)$  (in an unknown fashion) until it reaches truly OOD levels.

We assess the quality of the confidence scores using the notion of *calibration*. We define perfect calibration such that accuracy and confidence match for all confidence levels:

$$\mathbb{P}(\hat{Y} = Y | \hat{P} = p) = p, \quad \forall p \in [0, 1] \quad (2)$$

This directly leads to a definition of miss-calibration as the difference in expectation between confidence and accuracy: Based on equation 3 it is straight-forward to define miss-calibration as the difference in expectation between confidence and accuracy:

$$\mathbb{E}_{\hat{P}} \left[ \left| \mathbb{P}(\hat{Y} = Y | \hat{P} = p) - p \right| \right] \quad (3)$$

The expected calibration error (ECE) [12] is a scalar summary measure estimating miss-calibration by approximating equation 3 based on predictions, confidence scores and ground truth labels  $\{(Y_i, \hat{Y}_i, \hat{P}_i)\}_{i=1}^L$  of a finite number of  $L$  samples. ECE is computed by first partitioning all  $L$  confidence scores  $\hat{P}_i$  into  $M$  equally sized bins of size  $1/M$  and computing accuracy and average confidence of each bin. Let  $B_m$  be the set of indices of samples whose confidence falls into its associated interval  $I_m = (\frac{m-1}{M}, \frac{m}{M}]$ .  $\text{conf}(B_m) = 1/|B_m| \sum_{i \in B_m} \hat{P}_i$  and  $\text{acc}(B_m) = 1/|B_m| \sum_{i \in B_m} \mathbf{1}(\hat{Y}_i = Y_i)$  are the average confidence and accuracy associated with  $B_m$ , respectively. The ECE is then computed as

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} \left| \text{acc}(B_m) - \text{conf}(B_m) \right| \quad (4)$$

It can be shown that ECE is directly connected to miss-calibration, as ECE using  $M$  bins converges to the  $M$ -term Riemann-Stieltjes sum of eq. 3 [4].

## 4. Uncertainty calibration under domain drift

### 4.1. Baseline methods and experimental setup

We assess the performance of the following post-hoc uncertainty calibration methods in domain drift scenarios:

- Base: Uncalibrated baseline model
- Temperature scaling (TS): Post-hoc calibration by temperature scaling [4]
- Isotonic regression (IR) [23]
- Accuracy preserving version of Isotonic regression (IRM) [24]
- Composite model combining Temperature Scaling and Isotonic Regression (TS-IR) [24]


We quantify calibration under domain shift for 28 distinct perturbation types not seen during training, including 9 affine transformations, 19 image perturbations introduced by [6] and a dedicated bias-controlled dataset [1]. Each perturbation strategy mimics a scenario where the data of a deployed model encounters stems from a distribution that gradually shifts away from the training distribution in a different manner. For each model and each perturbation, we compute the micro-averaged ECE by first perturbing each sample in the test set at 10 different levels and then calculating the overall ECE across all samples; we denote relative perturbation strength as epsilon. A common manifestation of dataset shift in real-world applications is a change in object backgrounds, rotations, and imaging viewpoints. In order to quantify the expected calibration error under those scenarios, we use Objectnet, a recently proposed large-scale bias-controlled dataset [1]. The Objectnet dataset contains 50,000 test images with a total of 313 classes, of which 113 overlap with Imagenet. Uncertainty calibration under domain drift was evaluated for CIFAR-10 based on affine transformations, and for Imagenet based on the perturbations introduced by [6] as well as the overlapping classes in Objectnet.

In addition, we quantify the quality of predictive uncertainty for truly OOD scenarios by computing the predictive entropy and distribution of confidence scores. We use complete OOD datasets as well as data perturbed at the highest level. In these scenarios we expect entropy to reach maximum levels, since the model should transparently communicate it "does not know" via low and unbiased confidence scores.

#### 4.2. Improving calibration under domain drift

Existing methods for post-hoc uncertainty calibration are based on a validation set, which is drawn from the same generative distribution  $\pi(X, Y)$  as the training set and the test set. Using these data to optimize a post-hoc calibration method results in low calibration errors for data drawn from  $\pi(X, Y)$ . If we would like to generalise this approach to calibration under domain drift, we need access to samples from the generative distribution along the axis of domain drift. However, such robustness under domain drift is a challenging requirement since in practice for a  $D$ -dimensional input domain drift can occur in any of the  $2^D$  directions in  $\{-1, 1\}^D$ , and to any degree. Manifestations of such domain shifts include for example changes in viewpoint (Fig. 1), lighting condition, object rotation or background.

To obtain a transformed validation set representing a generic domain shift, we therefore sample domain drift scenarios by randomly choosing direction and magnitude of the domain drift. We use these scenarios to perturb the validation set and, taken together, simulate a generic domain



(a) Accuracy and entropy for (b) ECE for CIFAR-10 data

Figure 2: Model performance in terms of accuracy, entropy and expected calibration error for CIFAR-10 data for perturbation shear. (a) As expected accuracy degrades with increasing perturbation to almost random levels for all models. (b) While entropy increases with increasing perturbation strength Epsilon for all models, ECE also increases for all models, indicating a mis-match between confidence and accuracy.

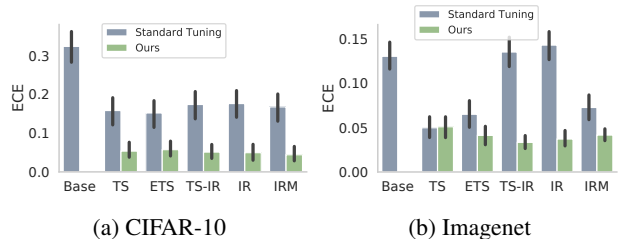


Figure 3: Mean expected calibration error, averaged over all test scenarios and levels of domain drift. Using our proposed calibration approach improves the overall calibration error for all post-hoc calibrators.

drift. More specifically, we first choose a random direction  $d_t \in \{-1, 1\}^D$ . Next, we sample from a set of 10 noise levels  $\epsilon$  covering the entire spectrum from in-domain to truly out-of domain. Each noise level corresponds to the variance of a Gaussian which in turn is used to sample the magnitude of domain drift. Since level and direction of domain shift are not known a priori, we argue that an image transformation using such Gaussian noise results in a generic validation set in the spirit of the central limit theorem: we emulate complex domain shifts in the real world by performing additive random image transformations, which in turn can be approximated by a Gaussian distribution.

We optimise  $\epsilon$  in a dataset-specific manner such that the accuracy of the pre-trained model decreases linearly in 10 steps to random levels (See Appendix for detailed algorithm).

In summary, we sample a domain shift scenario using Gaussian noise for each sample in the validation set, thereby generating a perturbed validation set. We then tune a given

	Base	TS	ETS	TS-IR	IR	IRM	TS-P	ETS-P	TS-IR-P	IR-P	IRM-P
CIFAR VGG19	0.323	0.158	0.152	0.173	0.176	0.167	0.053	0.057	0.051	0.049	<b>0.044</b>
CIFAR ResNet50	0.202	0.176	0.171	0.191	0.190	0.179	0.083	0.090	0.092	0.093	<b>0.076</b>
CIFAR DenseNet121	0.206	0.151	0.145	0.166	0.168	0.152	0.135	0.122	0.103	<b>0.088</b>	0.120
CIFAR MobileNetv2	0.159	0.150	0.141	0.165	0.165	0.147	0.107	0.125	0.094	<b>0.079</b>	0.108
ImgNet ResNet50	0.130	0.049	0.064	0.134	0.142	0.072	0.050	0.041	<b>0.033</b>	0.037	0.041
ImgNet ResNet152	0.129	0.043	0.049	0.127	0.135	0.062	0.037	0.034	<b>0.028</b>	0.039	0.045
ImgNet VGG19	0.057	0.045	0.047	0.120	0.122	0.051	0.093	0.075	0.064	<b>0.029</b>	0.047
ImgNet Den.Net169	0.117	0.044	0.040	0.127	0.133	0.057	0.024	<b>0.023</b>	0.026	0.045	0.050
ImgNet Eff.NetB7	0.092	0.135	0.085	0.131	0.132	0.074	0.074	0.047	<b>0.038</b>	0.049	0.058
ImgNet Xception	0.205	0.068	0.042	0.109	0.130	0.076	0.060	0.031	<b>0.031</b>	0.101	0.101
ImgNet Mob.Netv2	0.063	0.143	0.114	0.186	0.181	0.107	0.099	0.074	0.066	<b>0.046</b>	0.069

Table 1: Mean expected calibration error across all test domain drift scenarios (affine transformations for CIFAR-10 and perturbations proposed in [6] for Imagenet). For all architectures our approach of using a perturbed validation set outperformed baseline post-hoc calibrators

post-hoc uncertainty calibration method based on this perturbed validation set and obtain confidence scores that are calibrated under domain drift. All in all, we simulate domain drift scenarios and use the resulting perturbed validation set to tune existing post-hoc uncertainty calibration methods. We hypothesize that this facilitates calibrated predictions of neural networks under domain drift.

We refer to tuning a post-hoc calibrator using the perturbed validation set by a suffix "-P", e.g. IR-P stands for Isotonic Regression tuned on the perturbed validation set.

## 5. Experiments and results

We first illustrate limitations of post-hoc uncertainty calibration methods in domain drift scenarios using CIFAR-10. We show that while excellent in-domain calibration can be achieved using standard baselines, the quality of uncertainty decreases with increasing domain shift for all methods, resulting in highly overconfident predictions for images far away from the training domain.

Next, we show on a variety of architectures and datasets that replacing the validation set by a transformed validation set as outlined in section 4.2, substantially improves calibration under domain shift. We further assess the effect of our new approach on in-domain calibration and demonstrate that for selected post-hoc calibration methods, in-domain calibration can be maintained at competitive levels. Finally, we show that our tuning approach results in better uncertainty awareness in truly OOD settings.

### 5.1. Post-hoc calibration results in overconfident predictions under domain drift

We tuned all post-hoc calibration baseline methods on the CIFAR-10 validation set and first assessed in-domain calibration on the test set. As expected, calibration im-

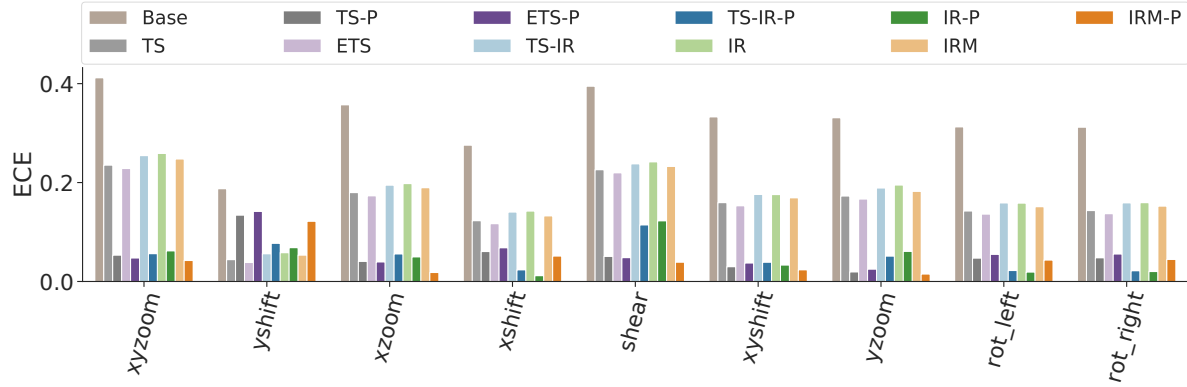
proves for all baselines over the uncalibrated network predictions in this in-domain setting. Next, we assessed calibration under domain drift by generating a perturbed test where we apply different perturbations (e.g. rotation, shear, zoom) to the images. We increased perturbation strength (i.e. shear) in 9 steps until reaching random accuracy. Figure 2 illustrates, that while entropy increases with intensifying shear for all models, ECE also increases for the entire set of models. This reveals a mis-match between confidence and accuracy that increases with increasing domain drift.

### 5.2. Perturbed validation sets improve calibration under domain drift


Next, we systematically assessed whether calibration under domain drift can be improved by tuning post-hoc calibration methods on a transformed validation set. To this end, we calibrate various neural network architectures on CIFAR-10 and Imagenet. For CIFAR-10, we first train VGG19 [18], ResNet50 [5], DenseNet121 [7] and MobileNetv2 [16] models. For Imagenet we used 7 pre-trained models provided as part of tensorflow, namely ResNet50, ResNet152, VGG19, DenseNet169, EfficientNetB7 [19], Xception [2] and MobileNetv2.

For all neural networks we then tuned 2 sets of post-hoc calibrators: one set was tuned in a standard manner based on the validation set, the second set was tuned with the proposed method based on the perturbed validation set. We then evaluate both sets of calibrators under various domain drift scenarios that were not seen during training, as well as in terms of in-domain calibration.

We observed that for all post-hoc calibrators tuning on a perturbed validation set resulted in an overall lower calibration error when testing across all domain drift scenarios, Table 1, Fig. 3). Figure 4 illustrates for VGG19 trained on



(a) CIFAR-10



(b) Imagenet

Figure 4: Micro-averaged calibration error for individual test domain shift scenarios

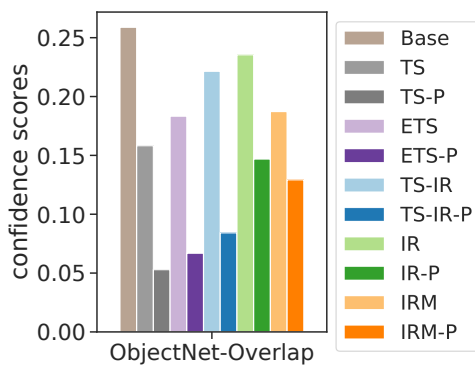


Figure 5: ECE for all overlapping classes contained in Objectnet dataset. Our tuning approach improves calibration for all post-hoc calibrators (Resnet50 trained on Imagenet dataset).

CIFAR-10 and Resnet50 trained on Imagenet that this improvement was consistent for individual domain drift scenarios not seen during training.

**Real-world domain shift** We next assessed the effect of our proposed tuning strategy on a real-world domain-shift. To this end, we computed ECE on the Objectnet test dataset. This comprises a set of images showing objects also present in Imagenet with different viewpoints, on new backgrounds and different rotation angles. As for artificial image perturbations, we found that our tuning strategy resulted in better calibration under domain drift compared to standard tuning, for all post-hoc calibration algorithms (Fig. 5).

**Dependency on magnitude of domain drift** We next focused on the Imagenet to assess how calibration depends on the amount of domain shift. We observe that while standard post-hoc calibrators yield very low in-domain calibration errors ((Fig. 7, ECE at epsilon 0), predictions become

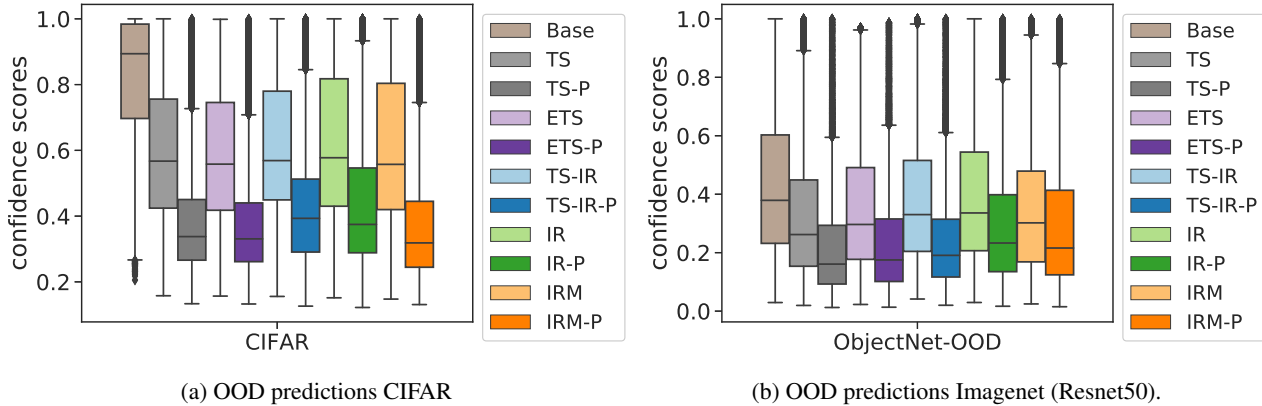


Figure 6: Distribution of confidence scores for out-of-domain predictions. (a) Confidence scores for our tuning strategy (-P) are substantially lower for the highest perturbation level compared to standard tuning, reflecting that our approach yields well calibrated predictions also for truly OOD samples (b) Confidence scores for OOD dataset Objectnet (non-overlapping classes) reveals that our approach results in substantially more uncertainty aware predictions.

increasingly overconfident with increasing domain shift. In contrast, all methods tuned using our approach had a low calibration error for large domain shifts. Notably, we observed two types of behaviour for in-domain calibration and small domain shift. One set of methods - TS-P, ETS-P and TS-IR-P - had a substantially increased calibration error for in-domain settings compared to standard calibration methods TS, ETS and TS-IR respectively. For these methods ECE at low levels of domain shift was comparable to that of uncalibrated models, and *decreased* substantially with domain shift (in contrast to uncalibrated models or standard calibrators where ECE *increases* with domain shift). However, IR-P did not show a worse in-domain calibration compared to standard post-hoc calibrators when tuned using our approach. Notably, it yielded a calibration error comparable to state-of-the-art calibrators for in-domain settings, while substantially improving calibration for increasing levels of domain shift. We further observed that IRM-P, an accuracy preserving version of IR-P, had a small but consistently worse in-domain classification error than IR-P, but performed substantially better than TS-based methods for small domain shifts.

One key difference between IR and TS-based methods, is that the latter methods are accuracy preserving, while IR-P (and IR) can potentially change the ranking of predictions and thus accuracy. However, for Imagenet we only observed minor changes in accuracy for IR-based methods, irrespective of tuning strategy (Fig. 4 (b), Fig. 7 (b)). We hypothesize that the systematic difference in calibration behaviour is due to the limitations of expressive power of TS-based accuracy-preserving methods: first, by definition the ranking between classes in a multi-class prediction does not change after applying this family of calibrators and second only one parameter (for TS) or 4 parameters (ETS) have

to capture the potentially complex behaviour of the calibrator. While this has important advantages for practitioners, it can also result in a lack of expressive power which may be needed to achieve calibration for all levels of domain drift, ranging from in-domain to truly out-of-domain. Finally, we observed that a variance reduction via temperature scaling [9] before IR (TS-IR) was also not beneficial for calibration under small levels of domain shift.

### 5.3. Out-of-distribution scenarios

To further investigate the behaviour of post-hoc calibrators for truly OOD scenarios, we analysed the distribution of confidence scores for a completely OOD dataset for Imagenet, using the 200 non-overlapping classes in Objectnet. This revealed that all post-hoc calibrators resulted in significantly less overconfident predictions when tuned using our approach. We observed a similar behaviour for CIFAR-10 with VGG19, when assessing the distribution of confidence scores at the highest perturbation level across all perturbations (Fig. 6).

## 6. Discussion and conclusion

We present a simple and versatile approach for tuning post-hoc uncertainty calibration methods. We demonstrate that our new approach, when used in conjunction with isotonic regression-based methods (IR or IRM), yields well-calibrated predictions in the case of any level of domain drift, from in-domain to truly out-of-domain scenarios. Notably, IR-P and IRM-P maintain their calibration performance for in-domain scenarios compared to standard isotonic regression (IR and IRM). In other words, our experiments suggest that when using our IR(M)-P approach, there is only a minor trade-off between choos-

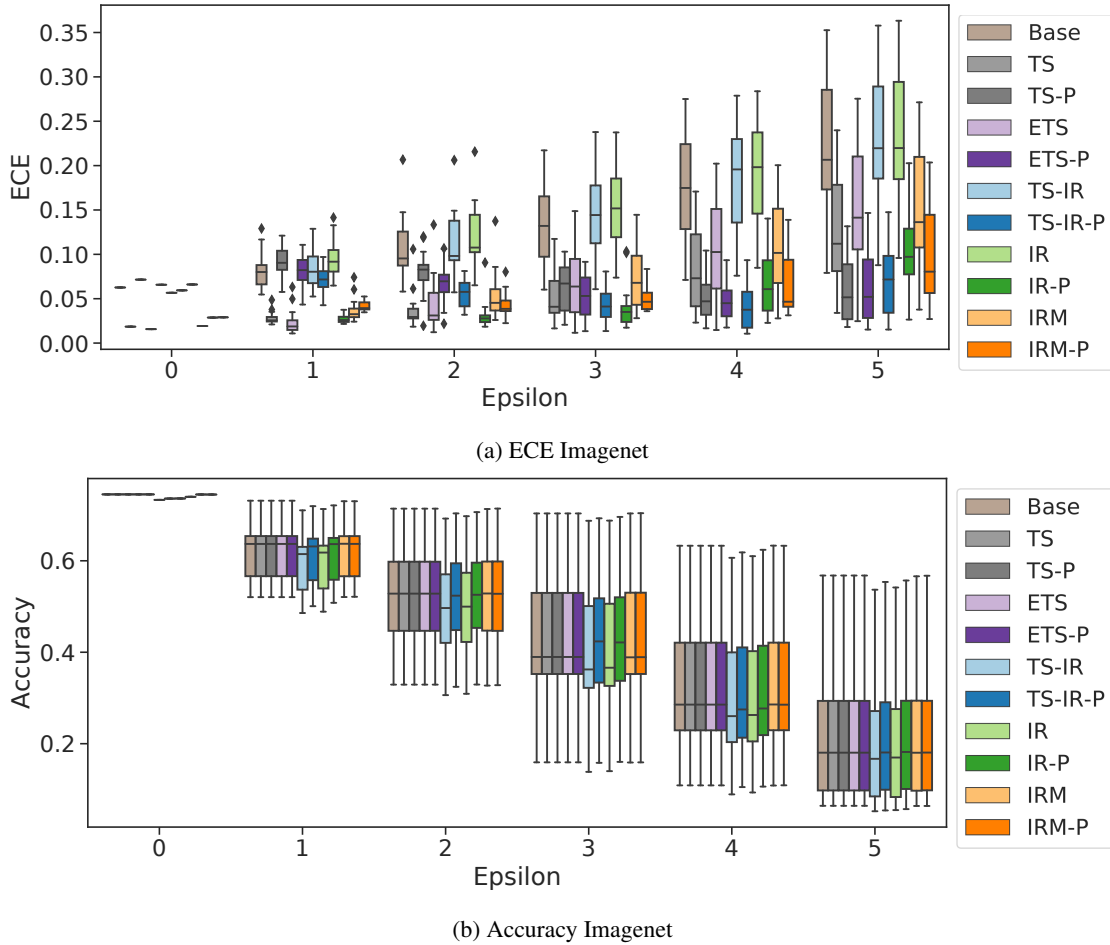


Figure 7: Dependency of accuracy and ECE on level of domain shift. Accuracy decreases for all methods with increasing levels of domain shift. For standard post-hoc calibrators ECE also increases, but our approach of tuning on a perturbed validation set results in good calibration throughout all levels, especially for IR-P.

ing a model that is either well calibrated in-domain vs. out-of-domain. In contrast, methods based on temperature scaling may not have enough expressive power to achieve good calibration across this range: standard tuning results in highly overconfident OOD predictions and perturbation-based tuning results in calibration errors comparable to uncalibrated models for in-domain predictions. However, when averaging across all domain drift scenarios, overall calibration for TS-P and ETS-P still improves substantially over standard TS and ETS. Consequently, for use-cases requiring an accuracy preserving method and reliable uncertainty estimates especially for larger levels of domain shift, TS-P and ETS-P are good options.

We further observe this trade-off between accuracy preserving properties and calibration error for IR-based methods. While IRM-P has accuracy-preserving properties, overall calibration errors are higher than for IR-P in particular for small domain shifts.

Our perturbation-based tuning can be readily applied to any post-hoc calibration method. When used in combination with an expressive non-parametric method such as IR, this results in well calibrated predictions not only for in-domain and small domain shifts, but also for truly OOD samples. This is in stark contrast to existing methods with standard tuning, where performance in terms of calibration and uncertainty-awareness degrades with increasing domain drift.

## 7. Acknowledgements

This work was supported by the Munich Center for Machine Learning and has been funded by the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036B.



## References

- [1] Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In *Advances in Neural Information Processing Systems*, pages 9448–9458, 2019.
- [2] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017.
- [3] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [4] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1321–1330. JMLR. org, 2017.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [6] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.
- [7] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [8] Heinrich Jiang, Been Kim, Melody Guan, and Maya Gupta. To trust or not to trust a classifier. In *Advances in Neural Information Processing Systems*, pages 5541–5552, 2018.
- [9] Ananya Kumar, Percy S Liang, and Tengyu Ma. Verified uncertainty calibration. In *Advances in Neural Information Processing Systems*, pages 3792–3803, 2019.
- [10] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pages 6402–6413, 2017.
- [11] Shiyu Liang, Yixuan Li, and R. Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. 2018.
- [12] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [13] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems*, pages 13991–14002, 2019.
- [14] Nicolas Papernot and Patrick McDaniel. Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning. *arXiv preprint arXiv:1803.04765*, 2018.
- [15] John C. Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In *Advances in large margin classifiers*, pages 61–74. MIT Press, 1999.
- [16] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [17] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. In *Advances in Neural Information Processing Systems*, pages 3179–3189, 2018.
- [18] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [19] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- [20] Christian Tomani and Florian Buettner. Towards trustworthy predictions from deep neural networks with fast adversarial calibration. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*, 2021.
- [21] Yeming Wen, Paul Vicol, Jimmy Ba, Dustin Tran, and Roger Grosse. Flipout: Efficient pseudo-independent weight perturbations on mini-batches. *arXiv preprint arXiv:1803.04386*, 2018.
- [22] Bianca Zadrozny and Charles Elkan. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *Icml*, volume 1, pages 609–616. Citeseer, 2001.
- [23] Bianca Zadrozny and Charles Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 694–699, 2002.
- [24] Jize Zhang, Bhavya Kailkhura, and T Han. Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning. In *International Conference on Machine Learning (ICML)*, 2020.