

---

# Non-Parametric Calibration for Classification

---

**Jonathan Wenger**  
University of Tuebingen  
TU Munich  
KTH Royal Institute of Technology  
jonathan.wenger@uni-tuebingen.de

**Hedvig Kjellström**  
KTH Royal Institute of Technology  
**Rudolph Triebel**  
German Aerospace Center (DLR)  
TU Munich  
hedvig@kth.se  
rudolph.triebel@dlr.de

## Abstract

Many applications of classification methods not only require high accuracy but also reliable estimation of predictive uncertainty. However, while many current classification frameworks, in particular deep neural networks, achieve high accuracy, they tend to incorrectly estimate uncertainty. In this paper, we propose a method that adjusts the confidence estimates of a general classifier such that they approach the probability of classifying correctly. In contrast to existing approaches, our calibration method employs a non-parametric representation using a latent Gaussian process, and is specifically designed for multi-class classification. It can be applied to any classifier that outputs confidence estimates and is not limited to neural networks. We also provide a theoretical analysis regarding the over- and underconfidence of a classifier and its relationship to calibration, as well as an empirical outlook for calibrated active learning. In experiments we show the universally strong performance of our method across different classifiers and benchmark data sets, in particular for state-of-the-art neural network architectures.

## 1 INTRODUCTION

With the recent achievements in machine learning, in particular in the area of deep learning, the application range for learning methods has increased significantly. Especially in challenging fields such as computer vision or speech recognition, important advancements have

been made using powerful and complex network architectures, trained on very large data sets. Most of these techniques are used for classification tasks, e.g. object recognition. We also consider classification in our work. However, in addition to achieving high classification accuracy, our goal is to provide reliable prediction uncertainty estimates. This is particularly relevant in safety-critical applications, such as autonomous driving and robotics (Amodei et al., 2016). Reliable uncertainties can be used to increase a classifier’s precision by reporting only class labels that are predicted with low uncertainty or for information theoretic analyses of what was learned and what was not. The latter is especially interesting in active learning, where the model actively selects the most relevant data samples for training via a query function based on the predictive uncertainty of the model (Settles, 2010).

Unfortunately, current probabilistic classification approaches that inherently provide good uncertainty estimates, such as Gaussian processes (GP), often suffer from lower accuracy and higher computational complexity on high-dimensional classification tasks compared to state-of-the-art convolutional neural networks (CNN). It was recently observed that many modern CNNs are overconfident (Lakshminarayanan et al., 2017, Hein et al., 2019) and miscalibrated (Guo et al., 2017). Calibration refers to how well confidence estimates of a classifier match the probability of the associated prediction being correct. Originally developed in the context of forecasting (Murphy, 1973, DeGroot and Fienberg, 1983), uncertainty calibration has seen an increased interest in recent years (Naeini et al., 2015, Guo et al., 2017, Vaicenavicius et al., 2019), partly because of the popularity of CNNs which generally lack an inherent uncertainty representation. Earlier studies show that also classical methods such as decision trees, boosting, SVMs and naive Bayes classifiers tend to be miscalibrated (Zadrozny and Elkan, 2001, Niculescu-Mizil and Caruana, 2005a,b, Naeini et al., 2015). Based on these observations, we claim that training and calibrating a classifier can be two different objectives that benefit




Figure 1: *Motivating example for calibration.* We trained a neural network with one hidden layer on MNIST (LeCun et al., 1998) and computed the classification error, the negative log-likelihood (NLL) and the expected calibration error ( $ECE_1$ ) for each training epoch. While accuracy continues to improve on the test set, the  $ECE_1$  increases after 20 epochs. This differs from classical overfitting as the test error continues to decrease. This indicates that improving both accuracy and uncertainty estimation can be conflicting objectives. However, we can mitigate this post-hoc via our calibration method (red dot). The uncertainty estimation after training and calibration is improved with maintained classification accuracy.

from being considered separately, as shown in a toy example in Figure 1. Here, a simple neural network continually improves its accuracy on the test set during training, but eventually overfits in terms of NLL and calibration error. A similar phenomenon was observed by Guo et al. (2017) for more complex models.

Calibration methods approach this problem by performing a post-hoc improvement to uncertainty estimation using a small subset of the training data. Our goal in this paper is to develop a *multi-class calibration method for arbitrary classifiers*, to provide reliable predictive uncertainty estimates in addition to maintaining high accuracy. In contrast to recent approaches, which strive to improve uncertainty estimation only for neural networks, including Bayesian neural networks (MacKay, 1992, Gal, 2016) and Laplace approximations (LA) (Martens and Grosse, 2015, Ba et al., 2017), our aim is a framework that is not based on tuning a specific classification method. This has the advantage that our method operates independently of the training process.

**Contribution** In this work we develop a new multi-class and model-agnostic approach to calibration, based on a latent Gaussian process inferred using variational inference. We replicate and extend previous findings that popular classification models are generally not calibrated and demonstrate the superior performance of our method for deep neural networks. Finally, we study the relationship between active learning and calibration from a theoretical perspective and give an empirical outlook.

**Related Work** Estimation of uncertainty, in particular in deep learning (Kendall and Gal, 2017), is of considerable interest in the machine learning community. There are two main approaches in classification. The first chooses a model and a (regularized) loss function for a particular problem to inherently learn a good

representation, and the second performs post-hoc calibration by transforming the output of the underlying model. For example, Pereyra et al. (2017) propose to penalize low-entropy output distributions, Kumar et al. (2018) suggest a trainable measure of calibration as a regularizer and Maddox et al. (2019) employ an approximate Bayesian inference technique using stochastic weight averaging. Milios et al. (2018) approximate Gaussian process classifiers by GP regression on transformed labels for better scalability and Wilson et al. (2016) combine additive Gaussian processes with deep neural network architectures. Research on calibration goes back to statistical forecasting (Murphy, 1973, DeGroot and Fienberg, 1983) and approaches to provide uncertainty estimates for non-probabilistic binary classifiers (Platt, 1999, Lin et al., 2007, Zadrozny and Elkan, 2002). More recently, Bayesian binning into quantiles (Naeini et al., 2015) and beta calibration (Kull et al., 2017a) for binary classification and temperature scaling (Guo et al., 2017) for multi-class problems were proposed. A theoretical framework for evaluating calibration in classification was suggested by Vaicenavicius et al. (2019). Calibration was also previously considered in the online setting with potentially adversarial input (Kuleshov and Ermon, 2017). Calibration in a broader sense is also of interest outside of the classification setting, e.g. in regression (Kuleshov et al., 2018, Song et al., 2019), in the discovery of causal Bayesian network structure from observational data (Jabbari et al., 2017) and in the algorithmic fairness literature (Pleiss et al., 2017, Kleinberg, 2018).

## 2 BACKGROUND

**Notation** Consider a data set  $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$  assumed to consist of independent and identically distributed realizations of the random variable  $(\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y}$  with  $K := |\mathcal{Y}|$  classes. If not stated otherwise,

any expectation is taken with respect to the law of  $(\mathbf{x}, y)$ . Let  $f : \mathcal{X} \rightarrow \mathbb{R}^K$  be a classifier with output  $\mathbf{z} = f(\mathbf{x})$ , prediction  $\hat{y} = \arg \max_i (z_i)$  and associated confidence score  $\hat{z} = \max_i (z_i)$ . Lastly,  $v : \mathbb{R}^K \rightarrow \mathbb{R}^K$  denotes a calibration method.

## 2.1 Calibration

A classifier is called *calibrated* if the confidence in its class prediction matches the probability of its prediction being correct, i.e.  $\mathbb{E}[1_{\hat{y}=y} | \hat{z}] = \hat{z}$ . In order to measure calibration, we define the *expected calibration error* following Naeini et al. (2015) for  $1 \leq p < \infty$  by

$$\text{ECE}_p = \mathbb{E}[|\hat{z} - \mathbb{E}[1_{\hat{y}=y} | \hat{z}]|^p]^{\frac{1}{p}} \quad (1)$$

and the *maximum calibration error* by  $\text{ECE}_\infty = \max_{z \in [0,1]} |\hat{z} - \mathbb{E}[1_{\hat{y}=y} | \hat{z} = z]|$ . In practice, we estimate the calibration error using a fixed binning for  $\hat{z}$  as described by Naeini et al. (2015). However, calibration alone is not sufficient for useful uncertainty estimates. A classifier on a balanced binary classification problem that always returns a confidence of 0.5 is perfectly calibrated, because this equals the probability of making a correct prediction. However, intuitively prediction confidence should be sufficiently close to 0 and 1 to be informative. This notion is known as sharpness or refinement (DeGroot and Fienberg, 1983, Murphy and Winkler, 1992, Cohen and Goldszmidt, 2004).

## 2.2 Over- and Underconfidence

We build on the notions of *over-* and *underconfidence*, as introduced previously in the context of active learning by Mund et al. (2015). The idea is to measure the average confidence of a classifier on its false predictions and the average uncertainty on its correct predictions:

$$o(f) = \mathbb{E}[\hat{z} | \hat{y} \neq y] \quad u(f) = \mathbb{E}[1 - \hat{z} | \hat{y} = y] \quad (2)$$

Over- and underconfidence are properties of a classifier independent of accuracy. They relate to query efficiency in active learning (Settles, 2010). Counter to intuition, both can be present to varying degrees simultaneously. We refer to Section S1 of the supplementary material for more details. We demonstrate that there is a direct link between calibration and these two notions.

**Theorem 1** (Calibration, Over- and Underconfidence) *Let  $1 \leq p < q \leq \infty$ , then the following relationship between overconfidence, underconfidence and the expected calibration error holds:*

$$|o(f)\mathbb{P}(\hat{y} \neq y) - u(f)\mathbb{P}(\hat{y} = y)| \leq \text{ECE}_p \leq \text{ECE}_q.$$

A proof is given in Section S1.2 of the supplementary material. We see that the expected calibration error




Figure 2: *Schematic diagram of calibration.* A fraction of the training data is split off and the remaining data is used for training. The split-off calibration data is classified by the trained model and subsequently used to fit the calibration method (blue). Confidence estimates from the classifier for new data are then adjusted by the calibration method (orange).

bounds the weighted absolute difference of over- and underconfidence. This implies that for perfect calibration, the odds of making a correct prediction equal the ratio between over- and underconfidence. Comparable statements exist in algorithmic fairness, where over- and underconfidence were termed *generalized* false positive and negative rates (Pleiss et al., 2017, Kleinberg, 2018).

## 2.3 Calibration Methods

The aim of calibrating a classifier is to transform its output to be closer to the true correctness probability. This is typically done by fitting a calibration method  $v$  on a small hold-out set called the *calibration data* (see Figure 2). In the following, we describe the most prevalent methods.

### 2.3.1 Binary Calibration

**Platt Scaling** (Platt, 1999, Lin et al., 2007) was originally introduced to provide probabilistic output for SVMs. It is a parametric calibration method, where a logistic regressor is fit to the confidence scores of the positive class such that

$$v(\mathbf{z})_2 = (1 + \exp(-a z_2 - b))^{-1}$$

for  $a, b \in \mathbb{R}$ . This parametric assumption is justified if the scores of each class are normally distributed with identical variance (Kull et al., 2017b).

**Isotonic Regression** (Zadrozny and Elkan, 2002) is a non-parametric approach. It assumes a non-decreasing relation between the model confidence  $z_2$  of the positive class and its correctness probability  $v(\mathbf{z})_2$ . A piecewise-constant isotonic function  $m$  is found by

minimizing a squared loss function, resulting in the calibration function  $v(\mathbf{z})_2 = m(z_2) + \epsilon$ .

**Beta Calibration** (Kull et al., 2017a) was designed for probabilistic classifiers with output range  $z_2 \in [0, 1]$ . A family of calibration maps is defined based on the likelihood ratio between two Beta distributions. The calibration map is given by

$$v(\mathbf{z})_2 = (1 + \exp(-c)(1 - z_2)^b z_2^{-a})^{-1},$$

where  $a, b, c \in \mathbb{R}$  are fit on the calibration data.

**Bayesian Binning into Quantiles (BBQ)** (Naeini et al., 2015) scores multiple equal-frequency binning models and uses a score weighted average of the accuracy in each bin as a calibration map. A binning model  $\mathcal{M}$  is weighted by  $\mathbb{P}(\mathcal{M})\mathbb{P}(\mathcal{D} | \mathcal{M})$ , where  $\mathbb{P}(\mathcal{M})$  is uniform and the marginal likelihood  $\mathbb{P}(\mathcal{D} | \mathcal{M})$  can be computed in closed form given parametric assumptions on data generation.

### 2.3.2 Multi-class Calibration

**One-vs-All** In order to extend binary calibration methods to multi-class problems, Zadrozny and Elkan (2002) suggest a one-vs-all approach, training a binary classifier on each split and calibrating subsequently. As most modern classifiers are inherently multi-class, this approach is not feasible anymore. We instead use a one-vs-all approach for the output  $\mathbf{z}$  of the multi-class classifier, train a calibration method on each split and average their predictions.

**Temperature Scaling** (Guo et al., 2017) was introduced as a multi-class extension to Platt scaling for neural networks. For an output logit vector  $\mathbf{z}$  of a neural network and a *temperature* parameter  $T > 0$ , the calibrated confidence is defined as

$$v(\mathbf{z}) = \sigma\left(\frac{\mathbf{z}}{T}\right) = \frac{\exp\left(\frac{z}{T}\right)}{\sum_{j=1}^K \exp\left(\frac{z_j}{T}\right)}. \quad (3)$$

The parameter  $T$  is determined by optimizing the NLL. By construction, the accuracy of the classifier is unchanged after scaling. Variants of this method where the factor is replaced by an affine map were shown to be ineffective (Guo et al., 2017).

## 3 GAUSSIAN PROCESS CALIBRATION

In the following section, we will outline our non-parametric calibration approach. Our aim is to develop a calibration method, which is inherently *multi-class*,




Figure 3: *Toy example of multi-class GP calibration.* The top panel shows the resulting latent Gaussian process with prior mean  $\mu(z_k) = \ln(z_k)$  when applying GP calibration to a synthetic calibration data set with four classes and 100 calibration samples. At the bottom, the class confidence scores making up the calibration data are plotted. The *calibration uncertainty* of the latent GP depends on their distribution.

suitable for *arbitrary classifiers*, makes as few assumptions as possible on the shape of the calibration map and can take prior knowledge into account. These desired properties motivate the use of a latent GP.

**Definition** Assume a one-dimensional Gaussian process prior over the latent function  $g : \mathbb{R} \rightarrow \mathbb{R}$ , i.e.

$$g \sim \mathcal{GP}(\mu(\cdot), k(\cdot, \cdot | \theta))$$

with mean function  $\mu$ , kernel  $k$  and kernel parameters  $\theta$  (Rasmussen and Williams, 2005). Further, let the calibrated output be given by the *softargmax* inverse link function applied to the latent process evaluated at the model output

$$v(\mathbf{z})_k = \sigma(g(z_1), \dots, g(z_K))_k = \frac{\exp(g(z_k))}{\sum_{j=1}^K \exp(g(z_j))}.$$

Note the similarity to multi-class Gaussian process classification, which has  $K$  latent functions (Rasmussen and Williams, 2005). In contrast, we consider *one shared latent function* applied to each component  $z_k$ . We use the categorical likelihood

$$\text{Cat}(\mathbf{y} | v(\mathbf{z})) = \prod_{k=1}^K \sigma(g(z_1), \dots, g(z_K))_k^{[y=k]} \quad (4)$$

to obtain a prior on the class prediction. We make the prior assumption that the classifier is already calibrated. This corresponds to either  $\mu(z_k) = \ln(z_k)$  if the inputs are confidence estimates, or to  $\mu(z_k) = z_k$  if the inputs are logits. For specific models other choices may be beneficial, e.g. a linear prior. An example of a latent function for a synthetic data set is shown in Figure 3. If the latent function  $g$  is monotonically increasing in its domain, the accuracy of the underlying classifier is unchanged after calibration.

**Inference** In order to infer the calibration map, we need to fit the underlying GP based on the confidence scores or logits and labels in the calibration set. Given the likelihood (4), the posterior is not analytically tractable. We use variational inference to approximate the posterior (Girolami and Rogers, 2006, Paul et al., 2012). For our method to scale to large data sets we only retain a sparse representation of the inputs, making inference computationally less intensive. We extend an approach by Hensman et al. (2015) to our choice of likelihood. The joint distribution of the data  $(\mathbf{z}_n, y_n)$  and latent variables  $\mathbf{g}$  is given by

$$\begin{aligned} p(\mathbf{y}, \mathbf{g}) &= p(\mathbf{y} | \mathbf{g})p(\mathbf{g}) = \prod_{n=1}^N p(y_n | \mathbf{g}_n)p(\mathbf{g}) \\ &= \prod_{n=1}^N \text{Cat}(y_n | \sigma(\mathbf{g}_n))\mathcal{N}(\mathbf{g} | \boldsymbol{\mu}, \boldsymbol{\Sigma}_{\mathbf{g}}), \end{aligned}$$

where  $\mathbf{y} \in \{1, \dots, K\}^N$ ,  $\mathbf{g} = (\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_N)^\top \in \mathbb{R}^{NK}$  and  $\mathbf{g}_n = (g(z_{n1}), \dots, g(z_{nK}))^\top \in \mathbb{R}^K$ . The covariance matrix  $\boldsymbol{\Sigma}_{\mathbf{g}}$  has block-diagonal structure by independence of the calibration data. If performance is important, a further diagonal assumption can be made. Note that we drop the explicit dependence on  $\mathbf{z}_n$  and  $\boldsymbol{\theta}$  throughout to lighten the notation. We want to compute the posterior  $p(\mathbf{g} | \mathbf{y})$ . In order to reduce the computational complexity  $\mathcal{O}((NK)^3)$ , we define  $M$  inducing inputs  $\mathbf{w} \in \mathbb{R}^M$  and inducing variables  $\mathbf{u} \in \mathbb{R}^M$ . The joint distribution is given by

$$p(\mathbf{g}, \mathbf{u}) = \mathcal{N} \left( \begin{bmatrix} \mathbf{g} \\ \mathbf{u} \end{bmatrix} \middle| \begin{bmatrix} \boldsymbol{\mu}_{\mathbf{g}} \\ \boldsymbol{\mu}_{\mathbf{u}} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{\mathbf{g}} & \boldsymbol{\Sigma}_{\mathbf{g}, \mathbf{u}} \\ \boldsymbol{\Sigma}_{\mathbf{g}, \mathbf{u}}^\top & \boldsymbol{\Sigma}_{\mathbf{u}} \end{bmatrix} \right). \quad (5)$$

The joint distribution factorizes as  $p(\mathbf{y}, \mathbf{g}, \mathbf{u}) = p(\mathbf{y} | \mathbf{g})p(\mathbf{g} | \mathbf{u})p(\mathbf{u})$ . We aim to find a variational approximation  $q(\mathbf{u}) = \mathcal{N}(\mathbf{u} | \mathbf{m}, \mathbf{S})$  to the posterior  $p(\mathbf{u} | \mathbf{y})$ . For general treatments on variational inference we refer to Blei et al. (2017), Zhang et al. (2018). We find the variational parameters  $\mathbf{m}$  and  $\mathbf{S}$ , the locations of the inducing inputs  $\mathbf{w}$  and the kernel parameters  $\boldsymbol{\theta}$  by maximizing a lower bound to the marginal log-likelihood

$$\begin{aligned} \ln p(\mathbf{y}) &\geq \text{ELBO}(q(\mathbf{u})) \\ &= \mathbb{E}_{q(\mathbf{u})} [\ln p(\mathbf{y} | \mathbf{u})] - \text{KL} [q(\mathbf{u}) \| p(\mathbf{u})] \\ &\geq \mathbb{E}_{q(\mathbf{u})} [\mathbb{E}_{p(\mathbf{g} | \mathbf{u})} [\ln p(\mathbf{y} | \mathbf{g})] \\ &\quad - \text{KL} [q(\mathbf{u}) \| p(\mathbf{u})]] \\ &= \mathbb{E}_{q(\mathbf{g})} [\ln p(\mathbf{y} | \mathbf{g})] - \text{KL} [q(\mathbf{u}) \| p(\mathbf{u})] \\ &= \sum_{n=1}^N \mathbb{E}_{q(\mathbf{g}_n)} [\ln p(y_n | \mathbf{g}_n)] \\ &\quad - \text{KL} [q(\mathbf{u}) \| p(\mathbf{u})] \end{aligned} \quad (6)$$

where  $q(\mathbf{g}) := \int p(\mathbf{g} | \mathbf{u})q(\mathbf{u}) d\mathbf{u}$  is Gaussian and only its  $K$ -dimensional marginals  $q(\mathbf{g}_n) = \mathcal{N}(\mathbf{g}_n | \boldsymbol{\phi}_n, \mathbf{C}_n)$

are required to compute the expectation terms. To do so, we use a second order Taylor approximation for  $\ln p(y_n | \mathbf{g}_n)$  and obtain

$$\begin{aligned} \mathbb{E}_{q(\mathbf{g}_n)} [\ln p(y_n | \mathbf{g}_n)] &\approx \ln p(y_n | \boldsymbol{\phi}_n) \\ &\quad + \frac{1}{2} (\boldsymbol{\sigma}(\boldsymbol{\phi}_n)^\top \mathbf{C}_n \boldsymbol{\sigma}(\boldsymbol{\phi}_n) - \text{diag}(\mathbf{C}_n)^\top \boldsymbol{\sigma}(\boldsymbol{\phi}_n)) \end{aligned}$$

which can be computed in  $\mathcal{O}(K^2)$ . Computing the KL-divergence term is in  $\mathcal{O}(M^3)$ . Therefore, computing the objective (6) has complexity  $\mathcal{O}(NK^2 + M^3)$ . Note that this can be remedied through parallelization as all  $N$  expectation terms can be computed independently. The optimization is performed via a gradient-based optimizer and automatic differentiation. We refer to Section S2 of the supplementary material for a more detailed treatment of inference.

**Calibration** Given the approximate posterior  $p(\mathbf{g}, \mathbf{u} | \mathbf{y}) \approx p(\mathbf{g} | \mathbf{u})q(\mathbf{u})$ , calibrated predictions at new inputs  $(\mathbf{z}_1, \dots, \mathbf{z}_L)^\top \in \mathbb{R}^{LN}$  are obtained via

$$\begin{aligned} p(\mathbf{g}_* | \mathbf{y}) &= \int p(\mathbf{g}_* | \mathbf{g}, \mathbf{u})p(\mathbf{g}, \mathbf{u} | \mathbf{y}) d\mathbf{g} d\mathbf{u} \\ &\approx \int p(\mathbf{g}_* | \mathbf{u})q(\mathbf{u}) d\mathbf{u} \end{aligned}$$

which is Gaussian. Mean  $\boldsymbol{\mu}_{\mathbf{g}_*}$  and variance of a latent value  $\mathbf{g}_* \in \mathbb{R}^K$  can be computed in  $\mathcal{O}(KM^2)$ . The class predictions  $\mathbf{y}_*$  are then obtained by marginalization

$$p(\mathbf{y}_* | \mathbf{y}) = \int p(\mathbf{y}_* | \mathbf{g}_*)p(\mathbf{g}_* | \mathbf{y}) d\mathbf{g}_*$$

via Monte-Carlo integration. While inference and calibration have higher computational cost than in other methods, it is orders of magnitude less than the training cost of the classifier. Furthermore, calibration can be performed in parallel with training in the online setting. We can speed up calibration by approximating the predictive distribution via the GP mean, i.e.  $p(\mathbf{y}_* | \mathbf{y}) \approx p(\mathbf{y}_* | \boldsymbol{\mu}_{\mathbf{g}_*}) = \sigma(\boldsymbol{\mu}_{\mathbf{g}_*})$ .

## 4 EXPERIMENTS

We experimentally evaluate our approach against the calibration methods presented in Section 2.3, applied to different classifiers on a range of binary and multi-class computer vision benchmark data sets. Besides CNNs, we are also interested in ensemble methods. All methods and experiments were implemented in Python 3.6. GPcalib was developed based on `gpflow` (Matthews et al., 2017). Any results reported used a sum kernel consisting of an RBF and a white noise kernel and a diagonal covariance matrix  $\boldsymbol{\Sigma}_{\mathbf{g}}$ .<sup>1</sup>

<sup>1</sup>An implementation of GP calibration and code replicating the experiments is available at <https://github.com/JonathanWenger/pycalib>.

#### 4.1 Calibration Results

We report the average  $ECE_1$  estimated with 100 bins over 10 Monte-Carlo cross validation runs. We chose a larger number of bins than in previous works, as too few bins typically underestimate the  $ECE_1$  (Kumar et al., 2019). See Section S3.1 of the supplementary material for details. We used the following data sets with indicated train, calibration and test splits:

- KITTI (Geiger et al., 2012, Narr et al., 2016): Stream-based urban traffic scenes with features (Himmelsbach et al., 2009) from segmented 3D point clouds. 8 or 2 classes, train: 16000, calibration: 1000, test: 8000.
- PCam (Veeling et al., 2018): Histopathologic scans of (metastatic) tissue from lymph node sections converted to grayscale. 2 classes, train: 22768, calibration: 1000, test: 9000.
- MNIST (LeCun et al., 1998): Handwritten digit recognition. 10 classes, train: 60000, calibration: 1000, test: 9000.
- CIFAR-100 (Krizhevsky, 2009): Image database of tiny color images from the web. 100 classes,, train: 50000, calibration: 1000, test: 9000.
- ImageNet 2012 (Russakovsky et al., 2015): Image database of natural objects. 1000 classes, train: 1.2 million, calibration: 1000, test: 10000.

**Binary Classification** We trained two boosting variants, AdaBoost (Freund and Schapire, 1997, Hastie et al., 2009) and XGBoost (Chen and Guestrin, 2016), two forest variants, Mondrian Forests (Lakshminarayanan et al., 2014) and Random Forests (Breiman, 2001), and a one layer neural network on the binary KITTI and PCam data sets. We report the average  $ECE_1$  in Table S3 in the supplementary material. For binary problems all calibration methods perform similarly with the exception of isotonic regression, which has particularly low calibration error on the KITTI data set. However, due to its piecewise constant calibration map the resulting confidence distribution has a set of singular peaks instead of a smooth distribution. While GPcalib is competitive across data sets and classifiers, it does not outperform isotonic regression. Hence, if exclusively binary problems are of interest a simple calibration method should be preferred. Interestingly, the 1-layer NN trained on KITTI is already well-calibrated, however all calibration methods except isotonic regression and GPcalib increase the  $ECE_1$ .

**Multi-class Classification** Besides the aforementioned classification models, which were trained on MNIST, we also calibrated pre-trained CNN ar-

chitectures<sup>2</sup> on CIFAR-100 and ImageNet. The following CNNs were used: AlexNet (Krizhevsky et al., 2012), VGG19 (Simonyan and Zisserman, 2014, Liu and Deng, 2015), ResNet-50, ResNet-152 (He et al., 2016), WideResNet (Zagoruyko and Komodakis, 2016), DenseNet-121, DenseNet-BC-190, DenseNet-201 (Huang et al., 2017), InceptionV4 (Szegedy et al., 2016), ResNeXt-29, SE-ResNeXt-50, SE-ResNeXt-101 (Xie et al., 2017, Hu et al., 2018), PolyNet (Zhang et al., 2017), SENet-154 (Hu et al., 2018), PNASNet-5-Large (Liu et al., 2018), NASNet-A-Large (Zoph et al., 2018). All binary calibration methods were extended to the multi-class setting in a one-vs-all manner. Temperature scaling and GPcalib were applied to logits for all CNNs and otherwise directly to probability scores. The average  $ECE_1$  is shown in Table 1. While binary methods still perform reasonably well for 10 classes in the case of MNIST and CIFAR-100, they worsen calibration considerably in the case of 1000 classes on ImageNet. Moreover, they also skew the posterior distribution so much that accuracy is heavily affected, disqualifying them from use. Temperature scaling preserves the underlying accuracy of the classifier by definition. Even though GP calibration has no such guarantees, our experiments show little effect on accuracy (see Table S4 in the supplementary material). GP calibration performs comparably to other methods on MNIST except for the simple NN and AdaBoost. It does not improve upon calibration for the simple NN, but it is the only method able to handle the large  $ECE_1$  of AdaBoost. GPcalib calibrates significantly better on CIFAR-100 than all other methods for all CNNs, except AlexNet. On ImageNet GPcalib demonstrates low  $ECE_1$  within one to two standard deviations of temperature scaling on four CNNs, but outperforms all other calibration methods on the remaining nine evaluated architectures. In particular on higher accuracy CNNs (see Table S4), GPcalib calibrates better. For CNNs which already demonstrate low  $ECE_1$ , such as InceptionV4 and SE-ResNeXt-50, most methods worsen calibration, whereas GPcalib does not. We attribute this desirable behavior, also seen in the binary case, to its prior assumption that the underlying classifier is already calibrated. The increased flexibility of the non-parametric latent map and its prior assumptions allow GPcalib to adapt to various classifiers and data sets.

**Latent Function Visualization** In order to illustrate the benefit of a non-linear latent function when calibrating, we show some latent functions from our experiments on ImageNet. We compare GPcalib with temperature scaling and no calibration corresponding to the

<sup>2</sup>Pre-trained CNNs were obtained from <https://github.com/bearpaw/pytorch-classification> and <https://github.com/Cadene/pretrained-models.pytorch>.

Table 1: *Multi-class calibration experiments.* Average  $ECE_1$  of 10 Monte-Carlo cross validation folds on multi-class benchmark data sets. Calibration errors ( $ECE_1$ ) within one standard deviation of lowest per data set and model are printed in bold.

Data Set	Model	Uncal.	one-vs-all					Temp.	GPcalib
			Platt	Isotonic	Beta	BBQ			
MNIST	AdaBoost	.6121	.2267	.1319	.2222	.1384	.1567	<b>.0414</b>	
MNIST	XGBoost	.0740	.0449	<b>.0176</b>	<b>.0184</b>	.0207	.0222	<b>.0180</b>	
MNIST	Mondrian Forest	.2163	.0357	.0282	.0383	.0762	<b>.0208</b>	<b>.0213</b>	
MNIST	Random Forest	.1178	.0273	.0207	.0259	.1233	<b>.0121</b>	.0148	
MNIST	1 layer NN	.0262	<b>.0126</b>	<b>.0140</b>	.0168	.0186	.0195	.0239	
CIFAR-100	AlexNet	.2751	.0720	.1232	.0784	.0478	<b>.0365</b>	<b>.0369</b>	
CIFAR-100	WideResNet	.0664	.0838	.0661	.0539	.0384	.0444	<b>.0283</b>	
CIFAR-100	ResNeXt-29 (8x64)	.0495	.0882	.0599	.0492	.0392	.0424	<b>.0251</b>	
CIFAR-100	ResNeXt-29 (16x64)	.0527	.0900	.0620	.0520	.0365	.0465	<b>.0266</b>	
CIFAR-100	DenseNet-BC-190	.0717	.0801	.0665	.0543	.0376	.0377	<b>.0237</b>	
ImageNet	AlexNet	<b>.0353</b>	.1132	.2937	.2290	.1307	<b>.0342</b>	<b>.0357</b>	
ImageNet	VGG19	.0377	.0965	.2810	.2416	.1617	<b>.0342</b>	<b>.0364</b>	
ImageNet	ResNet-50	.0441	.0875	.2724	.2250	.1635	<b>.0341</b>	<b>.0335</b>	
ImageNet	ResNet-152	.0545	.0879	.2761	.2201	.1675	.0323	<b>.0283</b>	
ImageNet	DenseNet-121	.0380	.0949	.2682	.2297	.1512	<b>.0329</b>	.0357	
ImageNet	DenseNet-201	.0410	.0898	.2706	.2189	.1614	<b>.0324</b>	.0367	
ImageNet	InceptionV4	.0318	.0865	.2900	.1653	.1593	.0462	<b>.0269</b>	
ImageNet	SE-ResNeXt-50	.0440	.0889	.2684	.1789	.1990	.0482	<b>.0279</b>	
ImageNet	SE-ResNeXt-101	.0574	.0853	.2844	.1631	.1496	.0415	<b>.0250</b>	
ImageNet	PolyNet	.0823	.0806	.2590	.2006	.1787	.0369	<b>.0283</b>	
ImageNet	SENet-154	.0612	.0809	.3003	.1582	.1502	.0497	<b>.0309</b>	
ImageNet	PNASNet-5-Large	.0702	.0796	.3063	.1430	.1355	.0486	<b>.0270</b>	
ImageNet	NASNet-A-Large	.0530	.0826	.3265	.1437	.1268	.0516	<b>.0255</b>	

identity map. Figure 4 illustrates in logit space how a non-linear latent function allows for lower  $ECE_1$ , when the calibration data necessitates it. When comparing latent functions, note that  $\sigma(g(\mathbf{z}) + \text{const}) = \sigma(g(\mathbf{z}))$ , i.e. an arbitrary shift of the latent functions on the y-axis corresponds to the same uncertainty estimates. We can also see how the latent GP gives information via its covariance on where the latent function’s shape is more certain based on the seen calibration data. For more examples, also for probability scores, we refer to Section S3.7 of the supplementary material.

**Computation Time** We give a complexity analysis for temperature scaling and GPcalib in Table 2. The cost of evaluation of the optimization objective and calibration for different variants of GP calibration are shown. For a diagonal covariance matrix, evaluating the optimization objective is of similar complexity to temperature scaling, since in general  $NK$  dominates the cubed number of inducing points. However, in our experiments the optimizer converged more slowly for GPcalib than for temperature scaling. We provide wall-clock inference and prediction runtime averaged across models per benchmark data set in Section S3.6 of the supplementary material. In practice, for most classifiers the training time is orders of magnitude larger than the time for inferring the latent GP in our calibration method. Calibration is computationally more expensive

for GPcalib compared to other methods in part due to the marginalization of the *calibration uncertainty*. This can be reduced by one order of magnitude for data sets with a large number of classes via the mean approximation with practically no effect on the  $ECE_1$ . The resulting time taken for calibration is about one-order of magnitude more than temperature scaling. In our experiments this was at least two orders of magnitude less than the prediction time of the classifier.

Table 2: *Computational complexity.* Complexity analysis of evaluation of the optimization objective for parameter inference of the calibration methods and complexity of calibration.  $N$  denotes the size of the calibration data,  $K$  the number of classes,  $M$  the number of inducing points and  $Q$  the number of MC samples. For the mean approximation  $a = 1$  with a diagonal covariance and  $a = 2$  for the full covariance. Implementation choices are  $M = 10$  and  $Q = 100$ .

Method	Optim. obj.	Calibration
Temp. scal.	$\mathcal{O}(NK)$	$\mathcal{O}(K)$
GPcalib		
diag. cov.	$\mathcal{O}(NK + M^3)$	$\mathcal{O}(K(M^2 + Q))$
full cov.	$\mathcal{O}(NK^2 + M^3)$	$\mathcal{O}(K^2(M^2 + Q))$
mean appr.	$\mathcal{O}(NK^a + M^3)$	$\mathcal{O}(K^a M^2)$




Figure 4: *Non-linear calibration maps in logit-space.* The plot shows latent functions of temperature scaling and GPcalib from a single CV run of our experiments on ImageNet. For PolyNet and PNASNet GPcalib shows a significant decrease in  $ECE_1$  in Table 1, corresponding to a higher degree of non-linearity in the latent GP.

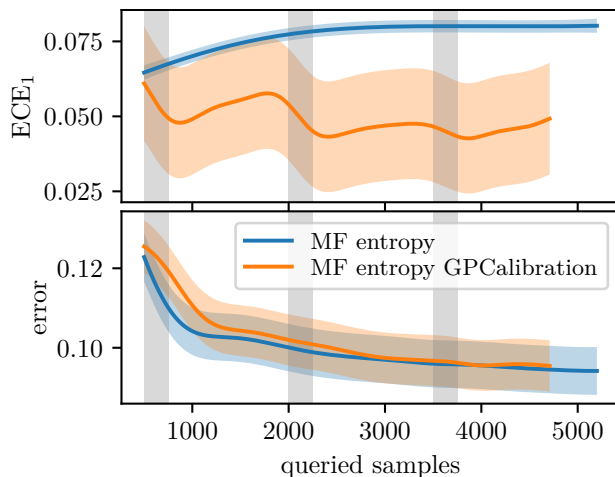


Figure 5: *Active learning and calibration.*  $ECE_1$  and classification error for two Mondrian forests trained online on labels obtained via an entropy query strategy on the KITTI data set. One forest is calibrated in regularly spaced intervals with GPcalib (gray). A GP regression up to the average number of queried samples across folds is shown. The calibrated forest queries  $\sim 10\%$  less labels, while reaching comparable accuracy.

## 5 CONCLUSION

In this paper we proposed a novel multi-class calibration method for arbitrary classifiers based on a latent Gaussian process, inferred via variational inference. We evaluated different calibration methods for a range of classifiers often employed in computer vision and robotics on a collection of benchmark data sets. Our method demonstrated strong performance across different model and data set combinations and performed particularly well for large-scale neural networks. In Theorem 1 we linked calibration to concepts from active learning. We conclude with a motivating example for the possible impact of calibration on active learning and outline future research directions.

**Active Learning** We hypothesize that calibration could improve active learning when querying based on uncertainty. We trained two Mondrian forests on the multi-class KITTI data set. These are well-suited for the online setting as they have the same distribution whether trained online or in batch. We randomly shuffled the data 10 times and requested samples based on an entropy query strategy with a threshold of 0.25. Any samples above the threshold are used for training or calibration. Both forests are trained for 500 samples and subsequently one uses 250 samples exclusively for calibration in regularly spaced intervals. We report the  $ECE_1$  and classification error in Figure 5. The calibration initially incurs a penalty on accuracy, as fewer samples are used for training. This is remedied over time through more efficient querying. The same accuracy is reached after a pass through the entire data set while *querying less samples overall*. This can be explained by calibration adjusting over- and underconfidence to the ratio determined by Theorem 1. Here, underconfidence is reduced, leading to less querying of uninformative samples. This is a promising result, but further research questions arise, regarding the size of the pre-training batch, the condition when calibration should be done, and the number of samples used for calibration. We believe that there is a trade-off, similar to an explore-exploit strategy, between classifier training and uncertainty calibration.

**Future Directions** Our proposed calibration approach is worth extending in the following directions: (a) forcing a monotone latent Gaussian process (Riihimäki and Vehtari, 2010, Agrell, 2019) provably preserves the accuracy of the underlying classifier; (b) extending our method to the online setting (Bui et al., 2017) allows for continuous calibration; (c) using our method for what we call *active calibration*, the concept of using an active learning query strategy, which switches between requesting samples for model training and uncertainty calibration based on the uncertainty of the latent Gaussian process.



## Acknowledgements

JW gratefully acknowledges financial support by the European Research Council through ERC StG Action 757275 / PANAMA; the DFG Cluster of Excellence “Machine Learning - New Perspectives for Science”, EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg. JW is grateful to the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. RT gratefully acknowledges the Helmholtz Artificial Intelligence Cooperation Unit (HAICU), which partly supported this work.

The authors thank the anonymous reviewers for helpful comments on an earlier version of this manuscript. In particular we appreciate the suggestions on inference using the mean estimate of the latent Gaussian process and regarding the definition of over- and underconfidence.

## References

- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul F. Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.
- Burr Settles. Active learning literature survey. Technical Report 55-66, University of Wisconsin, Madison, 2010.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, volume 86/11, pages 2278–2324, 1998.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6402–6413, 2017.
- Matthias Hein, Maksym Andriushchenko, and Julian 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 (CVPR)*, June 2019.
- 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 (ICML)*, 2017.
- Allan H. Murphy. A new vector partition of the probability score. *Journal of Applied Meteorology*, 12(4): 595–600, 1973.
- Morris H. DeGroot and Stephen E. Fienberg. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 32:12–22, 1983.
- Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In Blai Bonet and Sven Koenig, editors, *Proceedings of the Conference on Artificial Intelligence (AAAI)*, pages 2901–2907, 2015.
- Juozas Vaicenavicius, David Widmann, Carl Andersson, Fredrik Lindsten, Jacob Roll, and Thomas Schön. Evaluating model calibration in classification. In *Proceedings of the 22th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 3459–3467, 2019.
- Bianca Zadrozny and Charles Elkan. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *Proceedings of the 18th International Conference on Machine Learning (ICML)*, pages 609–616, 2001.
- Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, pages 625–632, 2005a.
- Alexandru Niculescu-Mizil and Rich Caruana. Obtaining calibrated probabilities from boosting. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, page 413, 2005b.
- David J. C. MacKay. A practical bayesian framework for backpropagation networks. *Neural Computation*, 4(3):448–472, 1992.
- Yarin Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016.
- James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2408–2417, 2015.
- Jimmy Ba, Roger Grosse, and James Martens. Distributed second-order optimization using kronecker-factored approximations. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5574–5584, 2017.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Lukasz Kaiser, and Geoffrey E. Hinton. Regularizing neural networks by penalizing confident output distributions. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.

- Aviral Kumar, Sunita Sarawagi, and Ujjwal Jain. Trainable calibration measures for neural networks from kernel mean embeddings. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 2810–2819, 2018.
- Wesley Maddox, Timur Garipov, Pavel Izmailov, Dmitry Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. *arXiv preprint arXiv:1902.02476*, 2019.
- Dimitrios Miliotis, Raffaello Camoriano, Pietro Michiardi, Lorenzo Rosasco, and Maurizio Filippone. Dirichlet-based gaussian processes for large-scale calibrated classification. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6008–6018, 2018.
- Andrew G Wilson, Zhiting Hu, Ruslan R Salakhutdinov, and Eric P Xing. Stochastic variational deep kernel learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2586–2594, 2016.
- 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.
- Hsuan-Tien Lin, Chih-Jen Lin, and Ruby C Weng. A note on platt’s probabilistic outputs for support vector machines. *Machine learning*, 68(3):267–276, 2007.
- Bianca Zadrozny and Charles Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the 8th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pages 694–699, 2002.
- Meelis Kull, Telmo Silva Filho, and Peter Flach. Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pages 623–631, 2017a.
- Volodymyr Kuleshov and Stefano Ermon. Estimating uncertainty online against an adversary. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*, pages 2110–2116, 2017.
- Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pages 2796–2804, 2018.
- Hao Song, Tom Diethe, Meelis Kull, and Peter Flach. Distribution calibration for regression. *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- Fattaneh Jabbari, Mahdi Pakdaman Naeni, and Gregory Cooper. Obtaining accurate probabilistic causal inference by post-processing calibration. In *Workshop on Causal Inference and Machine Learning (NeurIPS)*, 2017.
- Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5680–5689, 2017.
- Jon Kleinberg. Inherent trade-offs in algorithmic fairness. *ACM SIGMETRICS Performance Evaluation Review*, 46:40–40, 2018.
- Allan H. Murphy and Robert L. Winkler. Diagnostic verification of probability forecasts. *International Journal of Forecasting*, 7:435–455, 1992.
- Ira Cohen and Moises Goldszmidt. Properties and benefits of calibrated classifiers. In *Proceedings of the 8th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD)*, pages 125–136. Springer, 2004.
- D. Mund, R. Triebel, and D. Cremers. Active online confidence boosting for efficient object classification. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1367–1373, 2015.
- Meelis Kull, Telmo M Silva Filho, Peter Flach, et al. Beyond sigmoids: How to obtain well-calibrated probabilities from binary classifiers with beta calibration. *Electronic Journal of Statistics*, 11(2):5052–5080, 2017b.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2005.
- M. Girolami and S. Rogers. Variational bayesian multinomial probit regression with gaussian process priors. *Neural Computation*, 18(8):1790–1817, 2006.
- R. Paul, R. Triebel, D. Rus, and P. Newman. Semantic categorization of outdoor scenes with uncertainty estimates using multi-class gaussian process classification. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- James Hensman, Alexander G. de G. Matthews, and Zoubin Ghahramani. Scalable variational gaussian process classification. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

- Cheng Zhang, Judith Bütetpage, Hedvig Kjellström, and Stephan Mandt. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018.
- Alexander G. de G. Matthews, Mark van der Wilk, Tom Nickson, Keisuke Fujii, Alexis Boukouvalas, Pablo Le'on-Villagr'a, Zoubin Ghahramani, and James Hensman. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40):1–6, 2017.
- Ananya Kumar, Percy Liang, and Tengyu Ma. Variance reduced uncertainty calibration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- Alexander Narr, Rudolph Triebel, and Daniel Cremers. Stream-based active learning for efficient and adaptive classification of 3d objects. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 227–233. IEEE, 2016.
- Michael Himmelsbach, Thorsten Luettel, and H-J Wuensche. Real-time object classification in 3d point clouds using point feature histograms. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 994–1000. IEEE, 2009.
- Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant CNNs for digital pathology. In *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 210–218. Springer, 2018.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Trevor Hastie, Saharon Rosset, Ji Zhu, and Hui Zou. Multi-class adaboost. *Statistics and its Interface*, 2(3):349–360, 2009.
- Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, KDD '16, pages 785–794, 2016.
- Balaji Lakshminarayanan, Daniel M. Roy, and Yee Whye Teh. Mondrian forests: Efficient online random forests. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3140–3148, 2014.
- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1097–1105, 2012.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- S. Liu and W. Deng. Very deep convolutional neural network based image classification using small training sample size. In *3rd IAPR Asian Conference on Pattern Recognition (ACPR)*, pages 730–734, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *Proceedings of the British Machine Vision Conference (BMVC)*, 2016.
- 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 (CVPR)*, 2017.
- Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*, 2016.
- Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5987–5995, 2017.
- Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

- X. Zhang, Z. Li, C. C. Loy, and D. Lin. Polynet: A pursuit of structural diversity in very deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3900–3908, July 2017.
- Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018.
- Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8697–8710, 2018.
- Jaakko Riihimäki and Aki Vehtari. Gaussian processes with monotonicity information. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 645–652, 2010.
- Christian Agrell. Gaussian processes with linear operator inequality constraints. *Journal of Machine Learning Research*, 20(135):1–36, 2019.
- Thang D Bui, Cuong Nguyen, and Richard E Turner. Streaming sparse gaussian process approximations. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3299–3307, 2017.