

Improving Label Error Detection and Elimination with Uncertainty Quantification

JOHANNES JAKUBIK*, Karlsruhe Institute of Technology, IBM Research Europe, Switzerland

MICHAEL VOESSING, Karlsruhe Institute of Technology, IBM Germany, Germany

MANIL MASKEY, NASA, Marshall Space Flight Center, US

CHRISTOPHER WÖLFLE, Karlsruhe Institute of Technology, Germany

GERHARD SATZGER, Karlsruhe Institute of Technology, Germany

Identifying and handling label errors can significantly enhance the accuracy of supervised machine learning models. Recent approaches for identifying label errors demonstrate that a low self-confidence of models with respect to a certain label represents a good indicator of an erroneous label. However, latest work has built on softmax probabilities to measure self-confidence. In this paper, we argue that—as softmax probabilities do not reflect a model’s predictive uncertainty accurately—label error detection requires more sophisticated measures of model uncertainty. Therefore, we develop a range of novel, model-agnostic algorithms for Uncertainty Quantification-Based Label Error Detection (UQ-LED), which combine the techniques of confident learning (CL), Monte Carlo Dropout (MCD), model uncertainty measures (e.g., entropy), and ensemble learning to enhance label error detection. We comprehensively evaluate our algorithms on four image classification benchmark datasets in two stages. In the first stage, we demonstrate that our UQ-LED algorithms outperform state-of-the-art confident learning in identifying label errors. In the second stage, we show that removing all identified errors from the training data based on our approach results in higher accuracies than training on all available labeled data. Importantly, besides our contributions to the detection of label errors, we particularly propose a novel approach to generate realistic, class-dependent label errors synthetically. Overall, our study demonstrates that selectively cleaning datasets with UQ-LED algorithms leads to more accurate classifications than using larger, noisier datasets.

JAIR Associate Editor: Bo Li

JAIR Reference Format:

Johannes Jakubik, Michael Voessing, Manil Maskey, Christopher Wölfle, and Gerhard Satzger. 2025. Improving Label Error Detection and Elimination with Uncertainty Quantification. *Journal of Artificial Intelligence Research* 84, Article 19 (November 2025), 26 pages. DOI: [10.1613/jair.1.16158](https://doi.org/10.1613/jair.1.16158)

1 Introduction

In recent years, researchers and practitioners have mainly focused on improving the performance of machine learning models based on adjustments in the architectures and hyperparameters of the models. This so-called model-centric paradigm is increasingly complemented by the emerging data-centric paradigm, where the model—instead of the data—is considered static (Hamid 2022; Mazumder et al. 2022). In contrast to improving the model,

*Corresponding Author.

Authors’ Contact Information: Johannes Jakubik, johannes.jakubik1@ibm.com, Karlsruhe Institute of Technology, IBM Research Europe, Zurich, Switzerland; Michael Voessing, michael.voessing@kit.edu, Karlsruhe Institute of Technology, IBM Germany, Karlsruhe, Germany; Manil Maskey, manil.maskey@nasa.gov, NASA, Marshall Space Flight Center, Huntsville, US; Christopher Wölfle, ch.woelfle@icloud.com, Karlsruhe Institute of Technology, Karlsruhe, Germany; Gerhard Satzger, gerhard.satzger@kit.edu, Karlsruhe Institute of Technology, Karlsruhe, Germany.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.16158](https://doi.org/10.1613/jair.1.16158)

this approach aims at systematically enhancing the quality and quantity of data used (Aroyo et al. 2022; Haakman et al. 2021; Sambasivan et al. 2021; D. Wang et al. 2022). For supervised machine learning, an essential ingredient for strong performance is represented by high label quality (Han, Yao, Liu, et al. 2020; D. Wang et al. 2022). However, gathering data labels is not only costly but also prone to errors (Cordeiro and Carneiro 2020; Daniel et al. 2018; Han, Yao, Liu, et al. 2020). Misabeled data can then result in reduced performances of machine learning models, increased computational effort (e.g., driven by additionally required training data), and higher model complexity (Frénay and Verleysen 2014).

Recent work has demonstrated that label errors (often also called label noise) occur even in the most popular datasets (Northcutt, Athalye, et al. 2021). For example, in MNIST images depicting the digit nine are labeled as the digit *eight*, CIFAR-10 includes images of frogs, which are incorrectly labeled as cats, CIFAR-100 entails images depicting tigers, which are labeled as leopards, and in ImageNet images depicting meerkats are labeled as red pandas. Such incorrect labels affect the learning process of supervised machine learning models. Importantly, this especially applies to a vast majority of models pre-trained on standard datasets like ImageNet in the computer vision domain (Northcutt, Athalye, et al. 2021). To mitigate the undesired downstream effects of mislabeled data, several techniques have emerged to account for label noise (Song et al. 2022). Most of these approaches employ a model-centric perspective by trying to minimize the impact of label noise through specialized architectures (Goldberger and Ben-Reuven 2017; Xiao et al. 2015), regularization techniques (Hendrycks, Lee, et al. 2019; Lukasik et al. 2020), noise-robust loss functions (Ghosh et al. 2017; Zhang and Sabuncu 2018), or loss adjustment (Patrini et al. 2017; Shu et al. 2019). In contrast, only a small set of works adopt a data-centric perspective and address the root cause of the problem—identifying and excluding erroneously labeled data. We summarize this in Figure 1.




Fig. 1. Classification of state-of-the-art methods for handling label errors. Adapted from (Song et al. 2022).

The approach of identifying and avoiding mislabeled data is called *data cleansing*. One of the most prominent works in the field of data cleansing (both in research and practice), is represented by the confident learning (CL) technique (Northcutt, Jiang, et al. 2021). CL encompasses a collection of model-agnostic algorithmic approaches that automatically identify and remove label errors in datasets. CL approaches utilize deep neural networks to compute the model’s self-confidence for all labeled data, with the assumption that low self-confidence represents a good indicator for erroneous labels (Northcutt, Jiang, et al. 2021). For the computation of self-confidence, CL utilizes out-of-sample softmax probabilities. However, in this work, we argue that softmax probabilities do not reflect a model’s predictive uncertainty (Gal and Ghahramani 2016) and, therefore, are not ideal candidates for self-confidence measures in label error detection. In contrast, we leverage work on accurately measuring the predictive uncertainty of ML models based on approximations of Bayesian neural networks and apply it in the context of label error detection. In this work, we specifically focus on Monte Carlo Dropout (MCD) for measuring predictive uncertainty due to the small required computational effort (Gal and Ghahramani 2016).

Our contributions are as follows: (1) we improve CL by infusing more accurate estimates of the model uncertainty as self-confidence measures. Utilizing MCD probabilities allows us to account for the model’s

predictive uncertainty by including additional uncertainty measures, such as the entropy, in the label error detection process, thus improving the confidence of the algorithm’s decision on the presence of a label error even under high uncertainty. We develop a range of algorithms for uncertainty quantification-based label error detection based on the different concepts and measurements of predictive uncertainty. (2) We propose to use ensembles of algorithms (combining the novel algorithms with the state-of-the-art represented by the Prune-By-Noise-Rate Algorithm (CL-PBNR)) to leverage their individual strengths and, thereby, further improve the performance and generalizability of label error detection. (3) We develop a novel approach of adding synthetic label errors to common ML benchmark datasets, which accounts for the distribution of classes and permutes the labels of most similar classes instead of random classes and, thereby, generates more realistic label errors compared to previous works (Northcutt, Jiang, et al. 2021). (4) We holistically evaluate our algorithms in a two-stage setup: First, we measure the accuracy of label error detection by comparing the data samples classified as label errors with our methods with the synthetically generated ground truth. Second, we exclude detected label errors from the training dataset and measure the effect on the final model accuracy with and without label errors to account for recent debates on the importance of well-curated datasets instead of big data. Overall, our experiments demonstrate that incorporating predictive uncertainty based on approximations of Bayesian neural networks substantially improves the state-of-the-art in label error detection.

The remainder of this paper is structured as follows: First, we define label noise in the context of ML, outline different types of label errors, and highlight the resulting adverse consequences for the learning behavior. We then summarize CL as the current state-of-the-art method for identifying and excluding label errors and give a brief overview on approximations of Bayesian neural networks for uncertainty quantification. In Section 3, we describe our approach of including uncertainty in label error detection and cover the algorithm development. In the subsequent section, we propose a novel, more realistic approach for the synthetic generation of label noise, which generates asymmetric and class-dependent label errors compared to recent approaches for randomly generated synthetic label noise. We then holistically evaluate the proposed algorithms in a two-stage evaluation setup (first on the test sets of common ML benchmark datasets, second by measuring the effects of removing label errors from the entire training data). Finally, we discuss the implications and limitations of this work and provide an outlook on how future work can build on UQ-LED.

2 Background

Throughout this work, we focus on the classification of images as a very prominent and relevant task in the field of supervised ML (Bishop 2016). We consider an ML model (also referred to as classifier), which has been trained on a specific dataset and, during inference, is utilized to classify unseen samples. We follow recent work and focus on classification tasks with multiple classes (i.e., no binary classification) using deep neural networks with a softmax layer as classifiers. The datasets considered in this work are usually labeled by humans, e.g., via crowdsourcing (J. Cheng et al. 2020; Daniel et al. 2018; Patrini et al. 2017; W. Wang and Zhou 2015; Xiao et al. 2015).

2.1 Notation and Types of Label Errors

The notation of this work follows to large degrees the notation of (Northcutt, Jiang, et al. 2021). Let $\mathcal{K} := \{1, 2, \dots, c\}$ denote the set of c unique class labels and $X := (\mathbf{x}, \tilde{y})^n \in (\mathbb{R}^d, \mathcal{K})^n$ the training dataset consisting of n samples. Each sample $\mathbf{x} \in X$ has an observed noisy label $\tilde{y} \in \mathcal{K}$, which associates the sample to one of the c unique classes. While the observed label \tilde{y} should usually correspond to the true label y^* (Frénay and Verleysen 2014), it is possible that before observing \tilde{y} a stochastic noise process (Angluin and Laird 1988) maps $y^* \rightarrow \tilde{y}$, e.g., a human mislabels a picture of a crow as a raven. During this process, every sample \mathbf{x} with a true label $j \in \mathcal{K}$ may be independently mislabeled as class $i \in \mathcal{K}$ with a probability of $p(\tilde{y} = i \mid y^* = j, \mathbf{x})$ (Song et al.

2022). We will refer to this probability as the flipping probability. Probabilities that have been predicted by a model θ are denoted as \hat{p} . The $n \times m$ matrix of predicted probabilities is $\hat{P}_{z,i} = \hat{p}(\tilde{y} = i; \mathbf{x}_z, \theta)$, where \mathbf{x}_z is the z^{th} sample in the training dataset X . The corresponding observed noisy label of sample \mathbf{x}_z is \tilde{y}_z . Sometimes we abbreviate $\hat{p}(\tilde{y} = i; \mathbf{x}_z, \theta)$ with $\hat{p}_{\mathbf{x}_z, \tilde{y}=i}$ or just $\hat{p}_{\mathbf{x}, \tilde{y}=i}$. Lastly, let $\hat{\mathbf{p}}_{\mathbf{x}} := \{\hat{p}_{\mathbf{x}, \tilde{y}=i}, \forall i \in \mathcal{K}\}$ denote the vector of all softmax probabilities of sample $\mathbf{x} \in X$ and let $X_{\tilde{y}=i}$ denote the subset of samples in the dataset belonging to class $i \in \mathcal{K}$.

Following (Cordeiro and Carneiro 2020), we define label errors (or label noise) as samples whose observed labels do not represent the true class (Cordeiro and Carneiro 2020). We note that there exist different types of label errors, which we cover in the following and summarize in Figure 2.

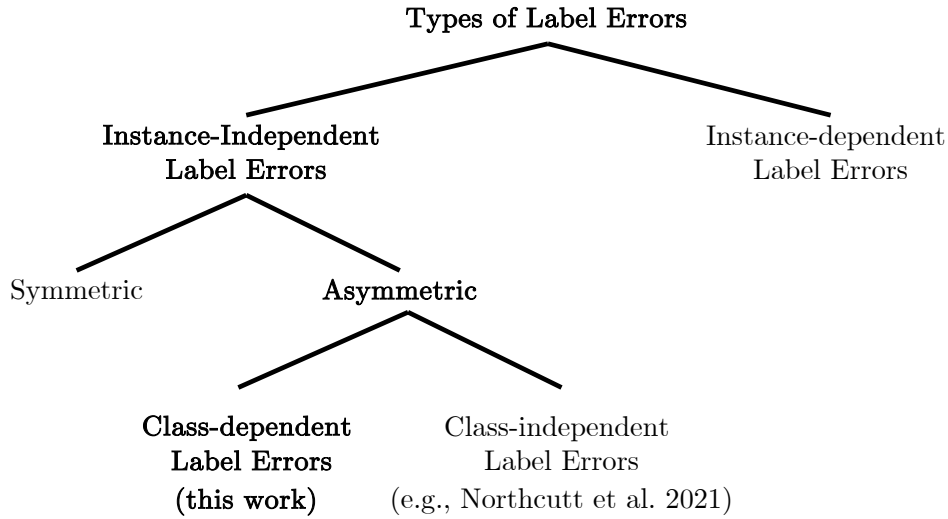


Fig. 2. Types of label errors.

In the literature, different types of label errors have been considered (Algan and Ulusoy 2020; Angluin and Laird 1988; H. Cheng et al. 2020; Cordeiro and Carneiro 2020; Frénay and Verleysen 2014; Patrini et al. 2017; Song et al. 2022). Following the taxonomy of (Song et al. 2022), the following label error types can be denoted: 1) instance-independent and 2) instance-dependent label errors. In the following, we will examine those different types and describe ways to model the respective noises.

1) Instance-Independent Label Errors: In this category, it is assumed that the noise process is independent of data features (Han, Yao, Liu, et al. 2020; Natarajan et al. 2013; Song et al. 2022; Xia et al. 2019). In general, instance-independent label errors are characterized by a noise transition matrix $T \in [0, 1]^{c \times c}$, where $\forall i, j \in \mathcal{K}$

$$T_{ij} := p(\tilde{y} = j \mid y^* = i) \quad (1)$$

denotes the flipping probability of the true label i being mislabeled as class j (Song et al. 2022). This flipping probability is instance-independent meaning that $p(\tilde{y} = j \mid y^* = i, \mathbf{x}) = p(\tilde{y} = j \mid y^* = i)$ (Han, Yao, Liu, et al. 2020). There are two specific forms of instance-independent label errors: symmetric noise and asymmetric noise.

Symmetric noise, also known as random noise or uniform noise, is a noise process where each label has an equal chance of flipping to another class with a fixed probability, independent of the true label y^* (J. Cheng et al. 2020; Cordeiro and Carneiro 2020; Song et al. 2022). Therefore, all entries in the noise transition matrix are evenly

distributed, except for the diagonal ones (Algan and Ulusoy 2020). Depending on the noise rate (ratio of noisy samples) $\tau \in [0, 1]$, the noise transition matrix can simply be modeled $\forall i, j \in \mathcal{K}$ as

$$T_{ij} = \begin{cases} 1 - \tau, & \text{if } i = j, \\ \frac{\tau}{c-1}, & \text{otherwise.} \end{cases} \quad (2)$$

Asymmetric noise, also known as label-dependent or class-conditional noise, is a noise process where the flipping probability $p(\tilde{y} = j \mid y^* = i)$ depends on the true label y^* and thus is not the same for every class (Algan and Ulusoy 2020; Natarajan et al. 2013; Patrini et al. 2017; Song et al. 2022). This type of noise is more comparable to label errors that occur in real life when classes are too similar for a human labeler regardless of the instance (Cordeiro and Carneiro 2020; Patrini et al. 2017). For example, an image of a donkey is more likely to be mislabeled as a horse than as a car. When modeling an asymmetric transition matrix for the noise rate $\tau \in [0, 1]$, it must fulfill the following criteria, $\forall_{i=j} T_{ij} = 1 - \tau \wedge \exists_{i \neq j, i \neq k, j \neq k} T_{ij} > T_{ik}$, with $i, j, k \in \mathcal{K}$, meaning that a true label has a higher chance to be mislabeled into a specific class (Song et al. 2022). (Northcutt, Jiang, et al. 2021) propose a method for modeling synthetic asymmetric label errors that randomly generates transition matrices with varying degrees of sparsity (fraction of zeros in the off-diagonals of the transition matrix). A high degree of sparsity reflects asymmetric label errors (Northcutt, Jiang, et al. 2021). However, their approach does not consider class similarities. To model an asymmetric noise process that considers class similarities, (Algan and Ulusoy 2020) propose to train a deep neural network (DNN) on the training set and then calculate its confusion matrix on the test set. The confusion matrix is then used as the noise transition matrix. (Gare and Galeotti 2021) propose a similar approach to synthetically generate class-conditional label errors, which also considers class similarities. We will use their approach as the foundation of our own synthetic noise generator.

2) Instance-Dependent Label Errors: Unlike instance-independent label errors, not all instances of a particular class have the same flipping probability. Instance-dependent or feature-dependent label errors is a more general noise process where the flipping probability $p(\tilde{y} = i \mid y^* = j, \mathbf{x})$ depends on the true class label and the sample's features (J. Cheng et al. 2020; Patrini et al. 2017). For example, a sample that has less information or is of lower quality is more likely to be mislabeled by a human (J. Cheng et al. 2020). Instance-dependent noise is harder to model than the aforementioned noise types since similarities between all samples need to be calculated (Algan and Ulusoy 2020; Song et al. 2022). For the complete notation of instance-dependent label errors, we refer to the works of (Xiao et al. 2015) and (H. Cheng et al. 2020) since this type of noise is not in the scope of this work.

We will describe the differences between class-dependent and class-independent label errors throughout Section 4, where we introduce the generation of label noise for this work.

2.2 Confident Learning

Confident Learning (CL) represents the state-of-the-art in sample selection-based label error detection (see Figure 1). In general, CL is a model-agnostic approach for characterizing, identifying, and learning with label errors, which makes use of the out-of-sample predicted probabilities to estimate the joint distribution between noisy and correct labels (Northcutt, Jiang, et al. 2021). The joint distribution is then consulted to identify potential label errors. These potential label errors are then excluded from the training data.

Due to the relevance of confident learning for the remainder of this work, we describe the approach by (Northcutt, Jiang, et al. 2021) in greater detail in the following: The general assumption of CL is that label noise is class-conditional, meaning that $p(\tilde{y} \mid y^*, \mathbf{x}) = p(\tilde{y} \mid y^*)$. To characterize this class-conditional noise, the latent (unknown) prior distribution of correct labels $p(\tilde{y} \mid y^*)$ and $p(y^*)$ needs to be estimated. Instead of estimating $p(\tilde{y} \mid y^*)$ and $p(y^*)$ separately, CL estimates both directly by estimating the joint distribution of the label errors $p(\tilde{y}, y^*)^1$ as matrix $\mathbf{Q}_{\tilde{y}, y^*}$.

¹Product Rule: $p(a, b) = p(a \mid b) \cdot p(b)$

CL requires two inputs for detecting label errors. That is, the out-of-sample predicted softmax probabilities $\hat{P}_{z,i}$ and the given labels of all samples \tilde{y}_z . The inputs are associated with each other via index z , $\forall x_z \in X$. The out-of-sample predicted (softmax) probabilities $\hat{P}_{z,i}$ can be obtained, for example, via cross-validation. Note that, in this work, we will change the input of CL and employ uncertainty measures instead of softmax probabilities as the input to calculate self-confidence.

CL then follows two major steps: First, the confident joint $C_{\tilde{y},y^*}$ is determined. This probability estimates $X_{\tilde{y}=i,y^*=j}$, the set of samples x labeled $\tilde{y} = i$ with the probability $\hat{p}(\tilde{y} = j; x, \theta)$ being high enough to be confidently counted as class $y^* = j$. In other words, it is the set of samples that likely belong to class j although their given noisy label is $\tilde{y} = i$. This decision is based on a class-specific threshold t_j , which represents the expected average self-confidence $\forall k \in \mathcal{K}$. (Northcutt, Jiang, et al. 2021) define self-confidence as the predicted probability that a sample x belongs to its given label \tilde{y} , which can be expressed as $\hat{p}(\tilde{y} = i; x \in X_{\tilde{y}=i}, \theta)$. Low self-confidence is then considered an indicator of a label error. After the confident joint has been calculated for every sample $x \in X$ the joint distribution matrix $Q_{\tilde{y},y^*}$ can be estimated (see (Northcutt, Jiang, et al. 2021) for details). Second, after the confident joint $C_{\tilde{y},y^*}$ and the joint distribution $Q_{\tilde{y},y^*}$ have been estimated, various rank and prune algorithms can be utilized to determine which samples should be excluded from the data. Recent work has demonstrated that the so-called *Prune by Noise Rate (PBNR)* approach achieves particularly high accuracy in identifying label errors (Northcutt, Jiang, et al. 2021). PBNR essentially measures the margin of how certain the model is that the given label $\tilde{y} = i$ is false, assuming that the true label is $y^* = j$. For each off-diagonal entry in $\hat{Q}_{\tilde{y}=i,y^*=j}$, where $i \neq j$, remove the $n \cdot \hat{Q}_{\tilde{y}=i,y^*=j}$ samples with the largest margin $\hat{p}_{x,y^*=j} - \hat{p}_{x,\tilde{y}=i}$.

2.3 Uncertainty in Deep Learning

Deep neural networks (DNN) do not capture the model's predictive uncertainty during a classification task. In contrast, DNNs calculate predictive probabilities obtained as output during inference by a DNN with a softmax function (Gal and Ghahramani 2016). These softmax probabilities $\hat{p}_{x,\tilde{y}=k}$ indicate how likely a sample x belongs to one of the $k \in \mathcal{K}$ classes, by approximating the relative probabilities between the different classes (Gal and Ghahramani 2016; Kendall et al. 2015). Contrary to common belief, those probabilities do not reflect the model's predictive uncertainty (Gal and Ghahramani 2016; Hendrycks and Gimpel 2016; Northcutt, Jiang, et al. 2021), meaning that even with a high softmax output, a model can be uncertain (Gal and Ghahramani 2016).

One approach to model predictive uncertainty is to use Bayesian neural networks (BNN). The Bayesian probability theory provides well-established methods like BNNs and Gaussian processes for obtaining model uncertainty in addition to the predictive probabilities. BNNs are stochastic neural networks that place distributions over the model parameters that express how likely the different parameter values are (Jospin et al. 2022; Neal 1996). This approach differs from conventional deep neural networks, which assign each parameter a single value obtained by minimizing a loss function (Jospin et al. 2022). So unlike deep neural networks, BNNs provide predictive distributions as outputs, which can be used to obtain model uncertainty (Alarab et al. 2021; Gal and Ghahramani 2016). However, BNNs are computationally expensive. To mitigate the computational costs of BNNs, (Gal and Ghahramani 2016) have proposed lightweight approximations of BNNs based on the so-called Monte Carlo Dropout (MCD).

Dropout is typically applied as a regularization technique at training time to reduce overfitting (Srivastava et al. 2014). During inference, dropout is usually not employed (Srivastava et al. 2014). However, recent work has demonstrated that drawing Monte Carlo samples from a model based on enabled dropout inference approximates BNN. During MCD inference, a probability distribution is generated by drawing multiple predictive samples from the model over various stochastic forward passes. During each forward pass, different neurons are switched off, leading to slightly different softmax outputs (Gal and Ghahramani 2016). Formally, let $\hat{p}_{x,\tilde{y}=k,f} = \hat{p}_f(\tilde{y} = k; x, \theta)$

denote the softmax output of model θ of forward pass f for class $k \in \mathcal{K}$. To get the final predictive output (per-class MCD probabilities), the mean over the outcomes of all F forward passes is calculated (Gal and Ghahramani 2016) as

$$\hat{p}^{\text{MCD}}_{\mathbf{x}, \tilde{y}=k} = \frac{1}{F} \sum_{f=1}^F \hat{p}_{\mathbf{x}, \tilde{y}=k, f} \quad \forall k \in \mathcal{K}, \quad (3)$$

while the final classification \hat{y} is obtained as

$$\hat{y} = \arg \max_{k \in \mathcal{K}} \left(\hat{p}^{\text{MCD}}_{\mathbf{x}, \tilde{y}=k} \right). \quad (4)$$

Since MCD inference provides a probability distribution that reflects an estimation of the model's uncertainty (Alarab et al. 2021), this probability distribution can be used to calculate uncertainty measures, such as the entropy (Gal and Ghahramani 2016), which is defined as

$$H(\hat{\mathbf{p}}^{\text{MCD}}_{\mathbf{x}}) = - \sum_{k=1}^K \hat{p}^{\text{MCD}}_{\mathbf{x}, \tilde{y}=k} \cdot \log \left(\hat{p}^{\text{MCD}}_{\mathbf{x}, \tilde{y}=k} \right). \quad (5)$$

For a uniform distribution of probabilities across classes, the entropy reaches its maximum. In contrast, if one class $i \in \mathcal{K}$ has a probability of $\hat{p}^{\text{MCD}}_{\mathbf{x}, \tilde{y}=i} = 1$, while the other classes $j \in \mathcal{K} \wedge j \neq i$ have a probability of $\hat{p}^{\text{MCD}}_{\mathbf{x}, \tilde{y}=j} = 0$, the entropy is minimal (Treiss et al. 2021).

3 Infusing Model Uncertainty in Label Error Detection

In this work, we develop novel, model-agnostic, uncertainty-based algorithms for label error detection. For this, we follow (Northcutt, Jiang, et al. 2021) and train an ML model on labeled data, including erroneous labels. We then utilize our developed algorithms to identify the label errors and measure both the label detection accuracy and the effect of excluding erroneous labeled data from training on the final model accuracy.

For our algorithms, we build on the previously introduced CL-PBNR algorithm by (Northcutt, Jiang, et al. 2021). However, instead of using out-of-sample predicted softmax probabilities, we leverage MCD probabilities, which allow us to employ uncertainty measures based on the underlying approximation of Bayesian neural networks resulting from multiple stochastic forward passes. On top of adding MCD probabilities, we then incorporate uncertainty measures, e.g., the entropy H or margin-based techniques, to improve the detection performance based on higher precision in assessing the confidence. In the following, we introduce the five developed algorithms for UQ-LED.

CL-MCD In order to measure the effect of using MCD as an approximation of Bayesian neural networks in label error detection, we replace the softmax probabilities with MCD probabilities as input to the CL-PBNR algorithm by (Northcutt, Jiang, et al. 2021).

CL-MCD + Entropy On top of MCD probabilities, the second algorithm leverages the entropy H of the MCD probabilities as input to the CL-PBNR algorithm. Following the recommendations for future research of (Northcutt, Jiang, et al. 2021) to try other thresholds, we use the entropy as an additional per-class threshold during the computation of the confident joint $C_{\tilde{y}, y^*}$. The entropy threshold te_j for each class $j \in \mathcal{K}$ is defined as the average entropy of all samples where the given label is j :

$$te_j = \frac{1}{|X_{\tilde{y}=j}|} \sum_{\mathbf{x} \in X_{\tilde{y}=j}} H \left(\hat{\mathbf{p}}^{\text{MCD}}_{\mathbf{x}} \right). \quad (6)$$

Since the entropy converges to zero with increasing model confidence about the prediction, the threshold is fulfilled if

$$H \left(\hat{\mathbf{p}}^{\text{MCD}}_{\mathbf{x} \in X_{\tilde{y}=j}} \right) \leq te_j. \quad (7)$$

A sample is now counted as confident prediction inside the confident joint if the default threshold t_j (average self-confidence for each class) defined in Section 2.2, and the new entropy threshold te_j is fulfilled. In other words, a sample is now only counted if the model is confident that the given label is wrong and if the model is certain with its overall prediction for the given sample. In the following, we will abbreviate this algorithm as CL-MCD-E.

CL-MCD-Ensemble When performing MCD inference, the results of F forward passes are combined into one single probability (see 2.3). Each forward pass f leads to a (slightly) different probability. By using the principles of ensemble learning, we feed each of the F predictions into the CL-PBNR algorithm and thus create a homogenous ensemble. Each of the F ensemble members receives the MCD probabilities from a different forward pass f as input and then returns a list of possible label error candidates. Based on the principle of majority vote, a sample's label is considered corrupted if more than 50% of the ensemble members marked the sample as a potential label error. Let m denote the required minimum number of ensemble members that must have identified a sample as erroneous in order to actually mark the sample as a label error. So if the number of forward passes is $F = 5$, then m must equal three in order to satisfy the majority vote condition. An overview of the CL-MCD-Ensemble can be found in Figure 3.

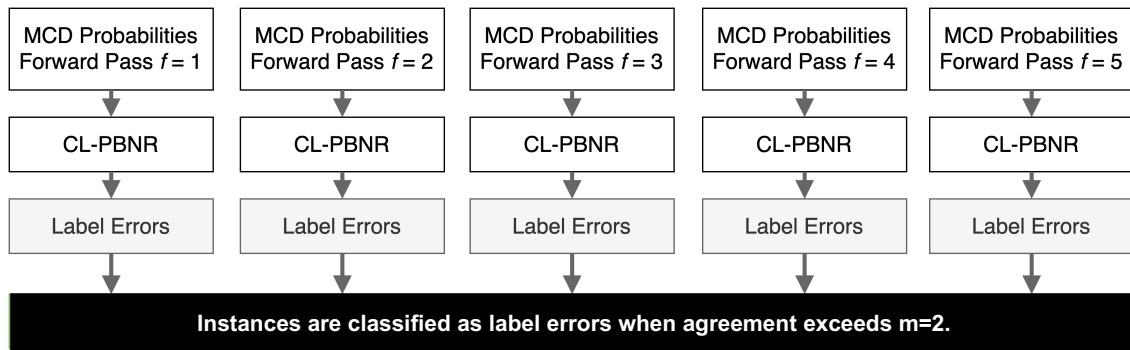


Fig. 3. Overview of the CL-MCD-Ensemble which we refer to as *CL-MCD-Ens*.

Algorithm Ensemble Having a homogenous ensemble at hand, we also want to explore a heterogenous ensemble consisting of several different algorithms. Therefore, we use our new algorithms 1-3 plus the CL-PBNR algorithm, resulting in an ensemble with four members. We included the CL-PBNR method as we noticed during the development that it has a very good recall compared to our proposed algorithms, which excel in their precision. The CL-PBNR algorithm receives softmax probabilities, and our other uncertainty-based algorithms receive MCD probabilities and, if required, the entropy as input. For this ensemble, we also apply the principle of majority vote. So, a sample's label is considered corrupted if more than 50% of the four ensemble members ($m = 3$) marked the sample as a potential label error. As this ensemble only consists of four members, we will additionally examine the label error detection performance for $m = 2$. A visual overview of the ensemble can be found in 4. In the following, we will abbreviate this ensemble with CL-MCD-Ens.

4 Generating Realistic Label Noise

For the evaluation of the label error detection, synthetic label errors need to be added to the datasets as, to our knowledge, there is no ground-truth dataset for label error detection. We follow the work on confident learning by considering asymmetric label noise, where some label errors are more likely than others. However, we deviate from (Northcutt, Jiang, et al. 2021) by considering class similarities in asymmetric label noise—e.g., a cat is more

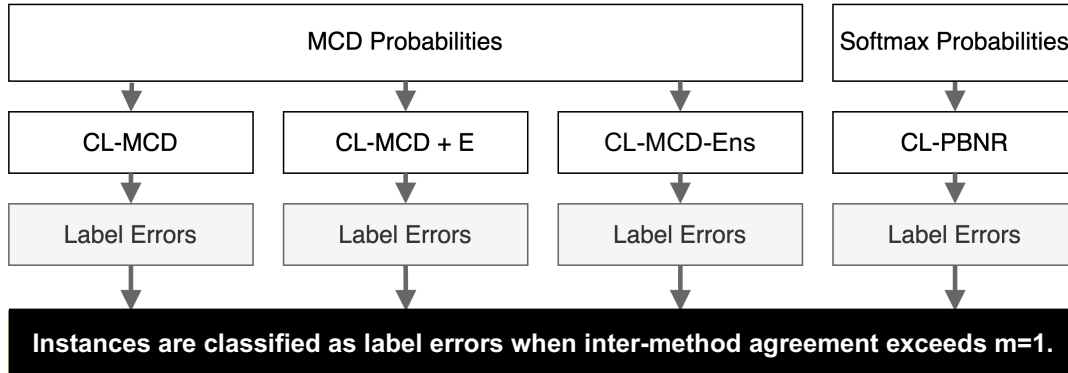


Fig. 4. Overview of the algorithm ensemble. We refer to confident learning based on prune-by-noise-ratio and softmax probabilities as *CL-PBNR*, confident learning based on Monte-Carlo Dropout as *CL-MCD*, confident learning based on Monte-Carlo Dropout and Entropy as *CL-MCD + E*, and the ensemble of CL-MCDs as *CL-MCD-Ens*.

likely to be confused with a dog than a car. We follow the literature on class-dependent label noise and train a model on the training set, and then use the predictions on the test set to generate similarity scores (Gare and Galeotti 2021). These similarity scores indicate how similar one class is to another. Following (Gare and Galeotti 2021), we define the similarity score between class k and class l , for all $\forall k, l \wedge k \neq l$ as

$$S_{k,l} = \frac{1}{|X_{\tilde{y}=k}|} \sum_{x \in X_{\tilde{y}=k}} \hat{p}(\tilde{y} = l; x, \theta), \quad (8)$$

i.e., the average of all probabilities of class l where the given label \tilde{y} equals k . The set of similarity scores for class k is defined as

$$S_k[l] := \{S_{k,l}, \forall l \in \mathcal{K} \wedge l \neq k\} \quad (9)$$

To better understand the noise generation process, we provide an example by adding label noise to the CIFAR-10 dataset based on the *airplane* class. In Table 1, we report the similarity scores on images from the test data of class $k=\text{airplane}$.

Table 1. Example of class similarities represented by softmax scores of deep learning models for class $k=\text{airplane}$ of CIFAR-10 dataset.

l	automobile	bird	cat	deer	dog	frog	horse	ship	truck
$S_{\text{airplane},l}$	0.004	0.019	0.007	0.002	0.000	0.001	0.002	0.027	0.010

After the similarity scores have been computed, a threshold

$$ts_k = \overline{S_k} + \sigma(S_k) \quad (10)$$

is calculated, and the indexes l of all similarity score set entries $S_k[l]$ that are higher than the threshold ts_k are grouped in the similarity group \mathcal{G}_k , which is defined as

$$\mathcal{G}_k := \{l : S_k[l] \geq ts_k, \forall l \in \mathcal{K} \wedge l \neq k\}. \quad (11)$$

For the CIFAR-10 example, the threshold $ts_{\text{airplane}} = 0.016$ and since only $S_{\text{airplane}, \text{bird}}$ and $S_{\text{airplane}, \text{ship}}$ are greater than the threshold the similarity group for class $k = \text{airplane}$ equals

$$\mathcal{G}_{\text{airplane}} := \{\text{bird}, \text{ship}\}. \quad (12)$$

After the similarity group \mathcal{G}_k is available, the synthetic flipping probabilities FP_k for class k are calculated as

$$FP_k[i] := \begin{cases} \text{softmax}(S_k[l]), & \text{if } l \in \mathcal{G}_k \\ 0, & \text{otherwise,} \end{cases} \quad (13)$$

with $i \in \mathcal{K}$. All the aforementioned steps are repeated for all classes $k \in \mathcal{K}$.

For the CIFAR-10 example, the synthetic flipping probabilities FP_{airplane} are reported in Table 2.

Table 2. Class-dependent flipping probabilities for the generation of synthetic label errors for all instances with given label $\tilde{y} = k = \text{airplane}$.

i	airplane	automobile	bird	cat	deer	dog	frog	horse	ship	truck
FP_{airplane}	0.0%	0.0%	49.8%	0.0%	0.0%	0.0%	0.0%	0.0%	50.2%	0.0%

Finally, to generate label noise for different noise rates τ , $\tau \cdot |\mathbf{X}|$ instances are randomly selected whose label will be flipped. The label of each selected instance $x \in \mathbf{X}$ with given label $\tilde{y} = k$ is flipped based on the previously calculated probabilities FP_k .

In Table 3, the noise transition matrix generated by our asymmetric noise generator for the CIFAR-10 dataset with a noise rate of $\tau = 0.2$ is shown. For our CIFAR-10 example, an instance with given label $\tilde{y} = k = \text{airplane}$ is misclassified as *bird* or *ship* with the same probability. The misclassification as *bird* seems plausible. The misclassification as a *ship* may come from the fact that the CIFAR-10 dataset also contains images of seaplanes that can easily be mistaken for a ship.

Table 3. Transition matrix $T_{i,j}$ for synthetic label errors for the CIFAR-10 dataset with a noise rate of $\tau = 0.2$, e.g., instances with given label $i = \text{automobile}$ are mislabeled with a probability of 20.1% as $j = \text{truck}$.

T_{ij}	air-plane	auto-mobile	bird	cat	deer	dog	frog	horse	ship	truck
airplane	80.3%	0.0%	9.9%	0.0%	0.0%	0.0%	0.0%	0.0%	9.9%	0.0%
automobile	0.0%	79.9%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	20.1%
bird	0.0%	0.0%	80.4%	10.0%	9.7%	0.0%	0.0%	0.0%	0.0%	0.0%
cat	0.0%	0.0%	0.0%	78.9%	0.0%	21.1%	0.0%	0.0%	0.0%	0.0%
deer	0.0%	0.0%	0.0%	19.3%	80.7%	0.0%	0.0%	0.0%	0.0%	0.0%
dog	0.0%	0.0%	0.0%	19.0%	0.0%	81.0%	0.0%	0.0%	0.0%	0.0%
frog	0.0%	0.0%	0.0%	20.2%	0.0%	0.0%	79.8%	0.0%	0.0%	0.0%
horse	0.0%	0.0%	0.0%	6.7%	7.0%	6.6%	0.0%	79.7%	0.0%	0.0%
ship	20.5%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	79.5%	0.0%
truck	0.0%	20.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	79.9%

5 Experimental Setup

In our experiments, we focus on image classification due to its comparability with existing approaches, its practical relevance, and computational feasibility. Specifically we make use of the prominent datasets MNIST (LeCun et al. 2010), CIFAR-10 (Krizhevsky 2009), CIFAR-100 (Krizhevsky 2009), and tiny-ImageNet (Le and Yang 2015). We evaluate each algorithm with regard to its label error detection performance as part of stage 1 of our evaluation. In stage 2 of the evaluation, we then measure the effect of excluding label errors from the training data on the final model accuracy based on different algorithms. We present a simplified overview of this procedure in Figure 5. During both stages, the state-of-the-art CL-PBNR algorithm is employed as a baseline (Northcutt, Jiang, et al. 2021).

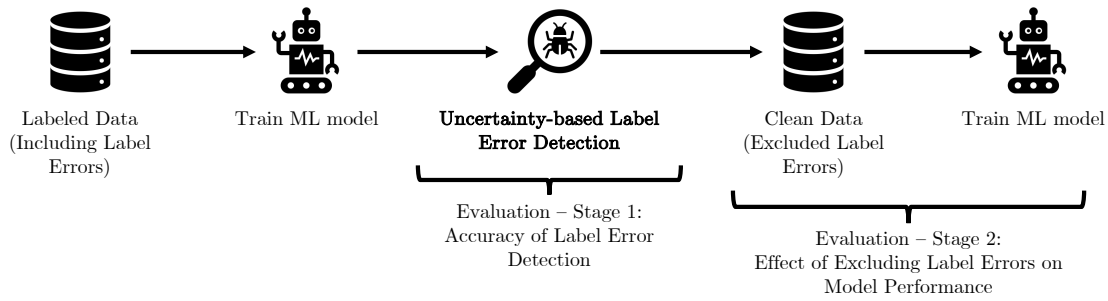


Fig. 5. Simplified overview of the procedure of uncertainty quantification-based label error detection.

In the following, we will examine our two-stage evaluation procedure in greater detail. First, we explain how we evaluate our algorithms with regard to the two evaluation criteria, label error detection performance and the effect on model accuracy, in a controlled setting. Subsequently, we will introduce the datasets for evaluation.

5.1 Performance Metrics

To evaluate our algorithms with regard to label error detection performance and the effect on model accuracy, synthetic label noise must be generated and added to the datasets before training. This is required because we can only measure the label error detection performance of our algorithms and compare the novel algorithms to the state-of-the-art if we have ground truth data on whether a label is correct or wrong. Furthermore, the noise rate τ must be controllable to investigate the effects of different noise rates on the algorithms' performance and the achieved accuracy after cleaning.

Contrary to other studies that deal with very high synthetic noise rates (Gare and Galeotti 2021; Han, Yao, Yu, et al. 2018; Northcutt, Jiang, et al. 2021), in this work, we want to use realistic noise rates as they occur in real-world datasets. (Northcutt, Athalye, et al. 2021) state that the test sets of common benchmark datasets, such as ImageNet, or Quickdraw, contain up to about 10% of label errors. Therefore, we use noise rates $\tau_1 = 0.05$, $\tau_2 = 0.1$, and $\tau_3 = 0.2$ for the evaluation of our algorithms.

In stage 1, we measure the label error detection performance of the different algorithms with the F1 score. Note that we additionally evaluated the area under precision-recall curve (AUPR) and the F0.5 score and found consistent results. For reasons of brevity, we follow recent work and report F1 score, precision, and recall. The F1

score equals the harmonic mean of recall and precision and is defined in the following equation.

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (14)$$

The precision reflects how many of the selected samples, i.e., true positives (TP) and false positives (FP), are actually relevant. It is defined as

$$\text{Precision} = \frac{TP}{TP + FP} \quad (15)$$

The recall reflects how many relevant items, i.e., TP and FN, are selected. It is defined as

$$\text{Recall} = \frac{TP}{TP + FN} \quad (16)$$

Apart from measuring the label error detection performance of our algorithms, we also want to assess their practical implications in terms of the effect on model accuracy after label error removal. Consistent with the approach of (Northcutt, Jiang, et al. 2021), we train a model on the noisy training set and another model on the cleaned training set before evaluating both on the unaltered test set. This allows us to determine the effect of cleaned data on model accuracy for each algorithm. We will refer to the accuracy before cleaning as *noisy accuracy*, and the accuracy after cleaning as *clean accuracy*. We describe our evaluation approach based on seven subsequent phases, which are explained in Figure 6.

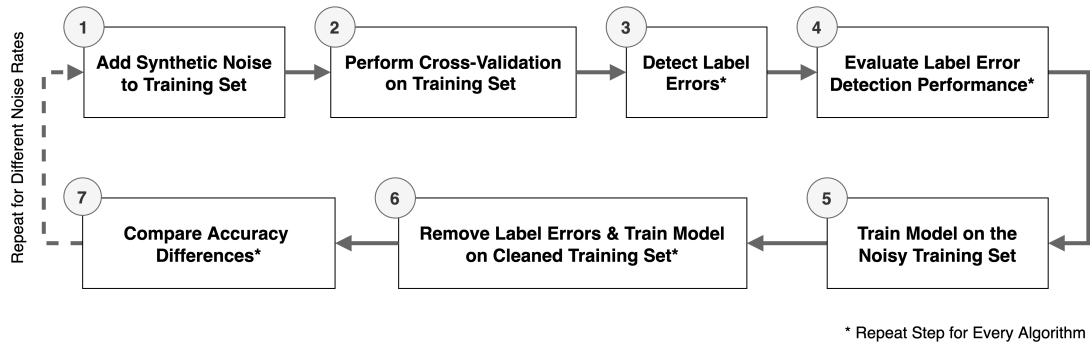


Fig. 6. In-depth evaluation procedure of the proposed algorithms for different noise rates.

First, label errors are synthetically generated and added to the training set. Second, out-of-sample predictions are computed from the noisy training set via k-fold cross-validation. Following (Northcutt, Jiang, et al. 2021), we conduct a stratified four-fold cross-validation. The predictions on each hold-out set are then combined to get the out-of-sample predictions for the entire noisy training set. Third, label errors are detected by the various algorithms. Fourth, each algorithm's label error detection performance is evaluated regarding recall, precision, and F1 score. Fifth, a model is trained on the noisy training set without cross-validation and evaluated on the unaltered test set to obtain the noisy accuracy. Sixth, for each algorithm, the label errors identified are automatically removed from the noisy training set, and a model is trained on the cleaned training set without cross-validation. The model is then evaluated on the unaltered test set to obtain the clean accuracy. Seventh, the clean accuracies achieved by our proposed algorithms are compared to the clean accuracy of the baseline (CL-PBNR) before then comparing the clean accuracies of all algorithms to the noisy accuracy. All stages are repeated for the noise rates $\tau_1 = 0.05$, $\tau_2 = 0.1$, and $\tau_3 = 0.2$.

Measuring Accuracy Differences: In stage 2, we also measure the differences in test set accuracy before and after cleaning. To calculate the noisy accuracy before label error removal, we train the dataset-specific model from scratch (without cross-validation) on the noisy training set and calculate the model’s accuracy on the unaltered test set. For each algorithm, we then remove the identified label errors from the noisy training set and train the model again from scratch on the cleaned training set. Afterward, we evaluate the model’s accuracy on the unaltered test set to obtain the clean accuracy for the respective algorithm. Lastly, we calculate the difference in accuracy between the model trained on the noisy and on the cleaned training set to measure the practical benefits of each algorithm. We use the same dataset-specific training settings as described above in the *noise generation* paragraph to train the model from scratch on the noisy and cleaned training sets. Note that we conduct a pilot study, where we evaluate all algorithms on the respective test sets of the datasets in order to determine which algorithms outperform CL-PBNR. These algorithms are then selected for an in-depth evaluation in order to reduce overall computation.

5.2 Datasets

The MNIST dataset published by (LeCun et al. 2010) is a prominent ML benchmark dataset of handwritten black and white digits ranging from zero to nine. It consists of 60,000 images for training and 10,000 images for testing. Examples of the MNIST dataset are depicted in Figure 7a. The CIFAR-10 dataset published by (Krizhevsky 2009) is another commonly used ML benchmark dataset of colored images from ten different classes and consists of 50,000 training and 10,000 testing samples. Examples of the CIFAR-10 dataset can be seen in Figure 7b. The CIFAR-100 dataset, also published by (Krizhevsky 2009), consists of colored images from 100 different classes and contains 50,000 training samples and 10,000 testing samples. Thus, for each class, there are only 500 training samples. Examples of the CIFAR-100 dataset can be seen in Figure 7c. The Tiny-ImageNet dataset published by (Le and Yang 2015) is a reduced version of the popular ImageNet classification challenge. The dataset consists of colored images from 200 different classes and has 100,000 training and 10,000 testing samples. With just 500 training images per class, this dataset has, together with the CIFAR-100 dataset, the lowest number of training images per class and the highest number of classes. In Figure 7d, we depict examples of the Tiny-ImageNet dataset.

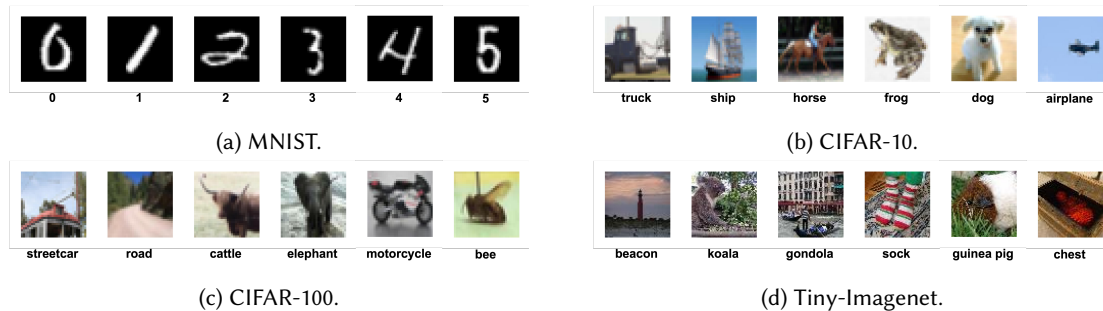


Fig. 7. Example visualizations from each considered dataset.

5.3 Implementation Details

In the following, we describe the pre-training, the noise generation, cross-validation, as well as hardware and frameworks.

Pre-Training For pre-training on the original training sets, we use a modified version of the VGG (Simonyan and Zisserman 2014) network as well as ResNet (He et al. 2016) architecture. Specifically, we use a VGG-11 model

for the MNIST dataset and ResNet-50 models for CIFAR-10, CIFAR-100, and Tiny-ImageNet. We adapt both model architectures towards uncertainty quantification by adding five dropout layers after the five most inner encoder layers, following (Kendall et al. 2015). The dropout probability of those layers is set to the standard value of 0.5 (Srivastava et al. 2014). For the pre-training on the MNIST training set, we train the modified VGG-11 model over 15 epochs with a batch size of 256 and a learning rate of 0.01. For the pre-training on the CIFAR-10, CIFAR-100, and Tiny-ImageNet datasets, we run three modified ResNet-50 models over 350 epochs with a batch size of 256. The learning rate scheduler reduces the initial learning rate of 0.1 after 150 epochs to 0.01 and, finally, after 250 epochs to 0.001. For all datasets, we split the non-test data into 80% training and 20% validation data. During the training of all models, we drop the last partial batch of each epoch. According to (Northcutt, Jiang, et al. 2021), this improves stability by avoiding weight updates from just a few noisy samples.

Noise Generation Before obtaining out-of-sample predicted probabilities from the MNIST, CIFAR-10, CIFAR-100, and Tiny-ImageNet datasets, we add synthetic asymmetric label noise to each dataset based on the method proposed above. For the noise generation, we leverage the previously pretrained models. After pre-training, we let each model predict on its test set and use the softmax probabilities to calculate the class-specific similarity scores and flipping probabilities per dataset. Finally, using the calculated similarity scores, we generate the noise transition matrices for the three different noise rates $\tau_1 = 0.05$, $\tau_2 = 0.1$, and $\tau_3 = 0.2$ for each dataset.

Cross-Validation After the label noise has been added to each training set, we conduct a four-fold cross-validation to obtain out-of-sample probabilities. For the cross-validation, we use the same model and training settings per dataset as described above. To obtain out-of-sample predicted MCD probabilities during the cross-validation, we conduct $F = 5$ forward passes. After the out-of-sample predicted softmax and MCD probabilities are obtained via the cross-fold validation, we identify label errors in the training set using the different algorithms and measure their label error detection performance.

Hardware and Frameworks Every model is trained on an NVIDIA Tesla V100. All code is developed in *PyTorch* (Paszke et al. 2019), and *PyTorch Lightning* (Falcon 2019), which is a high-level interface for *PyTorch* simplifying many repetitive tasks, like monitoring.

6 Experimental Results

In this section, we report the results and the key findings from our experiments. As introduced in Figure 5, we holistically evaluate the proposed algorithms in two stages. In the first stage, we assess the accuracy of detecting label errors within popular image classification datasets. In the second stage, we then examine the effect of excluding these potential label errors from the training datasets on the final model performance. Across both stages, we compare all proposed algorithms with the state-of-the-art label error detection represented by the CL-PBNR algorithm. We report our results across different levels of label errors in the datasets to gain additional insights into the effects of different amounts of label errors on performance.

6.1 Evaluation Stage 1: Accuracy of Label Error Detection

We report the accuracy in detecting label errors in terms of F1 scores, precision, and recall of all proposed algorithms alongside the state-of-the-art algorithm of CL-PBNR in Table 4. The experiments consistently demonstrate that all proposed uncertainty-based algorithms outperform the current state-of-the-art in label error detection for all considered datasets and across noise rates. The mean performance over CL-PBNR increases by 2.1pp for MNIST, 14.3pp for CIFAR-10, 8.2pp for CIFAR-100, and 6.6pp for Tiny-ImageNet. Additionally, we make several important observations.

First, our experiments suggest that it is beneficial to leverage MCD probabilities instead of softmax probabilities in general. Utilizing MCD probabilities improves the performances even when we do not calculate additional uncertainty measures, like entropy, on top of the probabilities (see the performance of CL-MCD). For example,

Table 4. Evaluation stage 1: Accuracy in identifying label errors for uncertainty quantification-based label error detection compared to the baseline, *prune by noise rate* algorithm CL-PBNR (Northcutt, Jiang, et al. 2021).

Measure Noise Rate τ	F1				Precision			Recall		
	0.05	0.1	0.2	Mean	0.05	0.1	0.2	0.05	0.1	0.2
MNIST										
Baseline (CL-PBNR)	93.9%	94.2%	94.4%	94.2%	92.6%	94.4%	94.9%	95.3%	93.9%	93.9%
CL-MCD	95.4%	95.2%	95.0%	95.2%	96.2%	97.4%	97.5%	94.7%	93.2%	92.7%
CL-MCD-E	95.2%	94.6%	92.3%	94.0%	94.9%	97.1%	97.4%	95.5%	92.2%	87.7%
CL-MCD-Ens	95.4%	95.5%	95.0%	95.3%	95.3%	96.7%	97.2%	95.5%	94.2%	93.0%
Alg. Ens (Agree.=2)	96.1%	96.2%	95.9%	96.1%	95.8%	97.4%	97.1%	96.4%	95.0%	94.8%
Alg. Ens (Agree.=3)	95.8%	94.9%	93.4%	94.7%	97.2%	98.4%	98.7%	94.4%	91.7%	88.7%
CIFAR-10										
Baseline (CL-PBNR)	49.3%	56.1%	60.3%	55.2%	33.7%	41.1%	47.7%	92.0%	88.4%	82.0%
CL-MCD	61.2%	63.9%	63.4%	62.8%	46.7%	51.0%	52.6%	88.8%	85.7%	80.0%
CL-MCD-E	67.6%	71.5%	69.4%	69.5%	55.7%	62.5%	62.0%	86.0%	83.5%	78.7%
CL-MCD-Ens	53.0%	59.2%	62.3%	58.2%	37.1%	44.1%	49.7%	92.8%	89.7%	83.3%
Alg. Ens (Agree.=2)	61.2%	64.4%	64.4%	63.3%	46.5%	51.3%	53.1%	89.5%	86.7%	81.8%
Alg. Ens (Agree.=3)	64.0%	67.1%	66.1%	65.7%	50.8%	56.2%	57.4%	86.5%	83.4%	77.9%
CIFAR-100										
Baseline (CL-PBNR)	24.4%	38.6%	53.4%	38.8%	14.1%	24.5%	37.9%	89.6%	90.8%	90.6%
CL-MCD	28.5%	42.7%	56.7%	42.6%	17.1%	28.3%	41.9%	85.8%	87.1%	87.3%
CL-MCD-E	32.5%	46.4%	59.0%	46.0%	20.5%	32.7%	46.7%	77.4%	79.9%	80.3%
CL-MCD-Ens	25.1%	39.4%	54.2%	39.6%	14.6%	25.0%	38.2%	92.2%	93.6%	93.7%
Alg. Ens (Agree.=2)	29.3%	43.5%	57.4%	43.4%	17.6%	29.0%	42.5%	85.7%	87.2%	88.2%
Alg. Ens (Agree.=3)	33.7%	47.8%	59.6%	47.0%	21.9%	34.8%	48.9%	73.4%	76.2%	76.2%
Tiny-ImageNet										
Baseline (CL-PBNR)	22.5%	37.4%	55.1%	38.3%	13.0%	24.0%	40.7%	82.5%	84.2%	85.2%
CL-MCD	24.6%	39.5%	57.5%	40.5%	14.5%	26.0%	43.7%	81.2%	81.9%	84.2%
CL-MCD-E	27.2%	43.0%	61.3%	43.8%	16.5%	29.6%	49.1%	76.6%	78.7%	81.7%
CL-MCD-Ens	23.8%	38.8%	57.1%	39.9%	13.8%	25.1%	42.3%	85.3%	85.8%	87.8%
Alg. Ens (Agree.=2)	25.6%	40.9%	59.0%	41.8%	15.2%	27.2%	45.2%	81.1%	82.5%	85.0%
Alg. Ens (Agree.=3)	29.0%	44.5%	61.1%	44.9%	18.4%	32.5%	52.6%	67.9%	70.4%	72.9%

the mean performance over CL-PBNR increases by 1.0pp for MNIST, 7.6pp for CIFAR-10, 3.8pp for CIFAR-100, and 2.2pp for Tiny-ImageNet. Second, the performance consistently increases when we leverage entropy as a measure of uncertainty on top of MCD probabilities. This underlines the importance of quantifying model uncertainty, especially for lower noise rates where the relative performance improvement over the state-of-the-art is particularly pronounced. We present these relative improvements in Figure 8. Third, combining different

algorithms in an ensemble can additionally improve the performances by leveraging different capabilities. This behavior is particularly pronounced for more complex datasets like CIFAR-100 and Tiny-ImageNet.

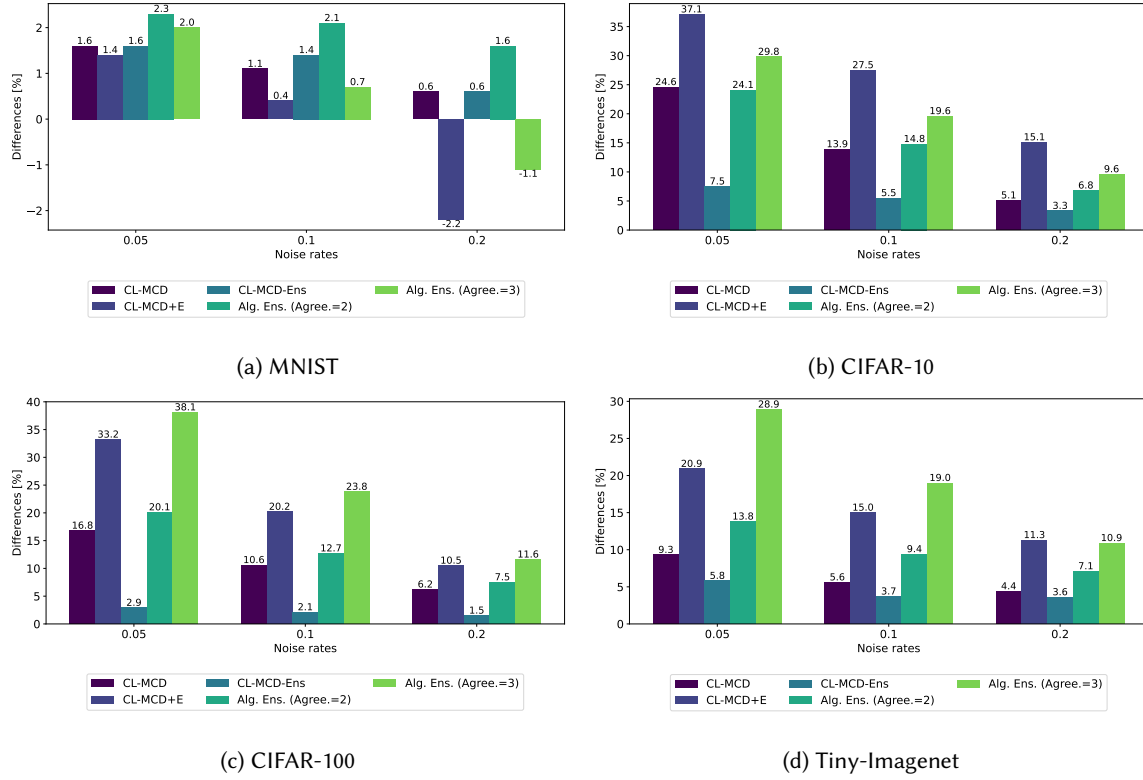


Fig. 8. Relative comparison between the proposed algorithms and the CL-PBNR algorithm in terms of increases or decreases in the F1 score.

The experiments reveal that the performance improvements over the state-of-the-art are especially large for more complex datasets. For CIFAR-10, CIFAR-100, and Tiny-ImageNet, our best-performing algorithms CL-MCD-E and the Algorithm Ensemble (Agreement = 3) improve over the state-of-the-art by at least 10% across all considered noise rates (see Figure 8). The best-performing algorithms outperform the state-of-the-art by up to 9.3pp in Table 4 (or 38.1%, see CIFAR-100, $\tau = 0.05$ in Figure 8). Although we observe smaller performance gains on simpler datasets, like MNIST, we acknowledge that it is more challenging for a model to further improve over a baseline, when the baseline performance is already very high.

We further note that our uncertainty-based algorithms perform particularly strong compared to the state-of-the-art when the relative share of label errors on the full dataset is small (e.g., $\tau = 0.05$). This observation is especially relevant since recent work has found that the share of label errors on real-world datasets ranges between 1% and 10% of the dataset size (Northcutt, Athalye, et al. 2021). Thus, the advantages of our algorithms enlarge when noise rates become more realistic.

The improvement in the detection of label errors (i.e., increasing F1 scores) of the proposed algorithms can be attributed to significant increases in the precision scores accompanied by smaller improvements in the recall. In

Table 4, we observe that the shortcomings of CL-PBNR lie in the precision with scores below 50% for CIFAR-10, CIFAR-100, and Tiny-ImageNet, while the recall of CL-PBNR is already consistently over 80%. All our algorithms increase both precision and recall. The precision improves by up to 22pp (or 65.28%, see CIFAR-10 results at $\tau = 0.05$) while the recall increases slightly by up to 2.8 (or 3.39%, see Tiny-ImageNet results at $\tau = 0.05$).

6.2 Evaluation Stage 2: Effect of Excluding Label Errors on Model Performance

In the second stage of our experiments, we assess the effect of excluding the previously identified label errors from the training datasets on the performance of the ML models. For these experiments, potential label errors are only removed from the training data while the test datasets remain unchanged. Overall, for all datasets besides Tiny-Imagenet, excluding label errors from the training dataset improves the model performance with all proposed algorithms and CL-PBNR baseline. The mean performance increases based on excluding label errors with the proposed algorithms compared to training on the original training datasets, amounting to up to 0.2pp for MNIST, 7.1pp for CIFAR-10, and 3.5pp for CIFAR-100 as reported in Table 5.

Importantly, the proposed algorithms again consistently outperform the state-of-the-art CL-PBNR algorithm in terms of the resulting model performance after detecting and removing label errors with the algorithms. In our experiments, removing label errors based on the proposed algorithms improved the mean performance (i.e., averaged over all considered noise rates) over the CL-PBNR algorithm by up to 0.3pp for MNIST, 0.9pp for CIFAR-10, and 2.6pp for CIFAR-100. For Tiny-Imagenet, excluding label errors did not improve the performance for any algorithm—however, our proposed algorithms still outperformed CL-PBNR by up to 2.0pp. Thus, overall, the proposed algorithms are not only more accurate in identifying label errors (see Evaluation Stage 1), but these improvements also translate to an increased performance of ML models based on excluding label errors. We present the relative improvements of the developed algorithms in Figure 9.

Besides the performance improvements of the proposed algorithm, our experiments demonstrate that these algorithms exclude significantly less data as potential label errors than CL-PBNR. Together with the strong increases in precision scores we observed in the first stage of the evaluation, we can infer that the performance improvements compared to CL-PBNR can be attributed to a significant reduction of false positives. For example, at $\tau_3 = 0.2$, the CL-MCD-E algorithm achieves a 2.1% better accuracy after cleaning than the baseline and removes 26.0% fewer samples on the CIFAR-10 dataset. On the CIFAR-100 dataset, the Algorithm Ensemble (Agreement=3) outperforms the baseline by 3.5% while removing 47.0% fewer samples at $\tau_1 = 0.05$. For Tiny-ImageNet, at $\tau_1 = 0.05$, the Algorithm Ensemble (Agreement=3) achieves a 4.3% increase in accuracy over the baseline while removing 41.8% fewer samples.

6.3 Sensitivity Analyses

Throughout the following sensitivity analyses, we will investigate three important relations in label error detection. The first sensitivity analysis studies the relationship between the relative share of label errors in the dataset and the accuracy of label error detection. We hypothesize that the relative share of label errors in the dataset may influence the accuracy of detecting label errors since the distributions of correctly labeled and mislabeled data change, which might influence confident learning in general and UQLED in specific. The second relationship is represented by the initial accuracy of the deep learning model and the accuracy in detecting label errors. We expect that a higher initial model accuracy translates to increased accuracy in label error detection as the model introduces fewer classification errors. Finally, the third relationship deals with the influence of the accuracy in label error detection on the final model performance trained on the cleaned datasets. In line with previous work, we expect higher performance in label error detection to translate to higher final model performances based on high-quality input data.

Table 5. Evaluation stage 2: Final model accuracy after removing label errors from the training datasets. Test sets remain unchanged.²

Metric Noise Rate τ	Accuracy				# of potential label errors		
	0.05	0.1	0.2	Mean	0.05	0.1	0.2
MNIST							
Original dataset	99.0%	98.9%	98.5%	98.8%	0	0	0
CL-PBNR	98.9%	98.8%	98.5%	98.7%	3,089	5,970	11,879
CL-MCD	99.3%	98.9%	98.7%	99.0%	2,954	5,744	11,405
CL-MCD-E	99.2%	98.8%	99.0%	99.0%	3,019	5,700	10,808
CL-MCD-Ensemble	99.0%	99.0%	98.1%	98.7%	3,004	5,847	11,475
Alg. Ens. (Agreement=2)	99.2%	99.0%	98.6%	98.9%	3,020	5,850	11,718
Alg. Ens. (Agreement=3)	99.1%	98.9%	98.7%	98.9%	2,916	5,592	10,785
CIFAR-10							
Original dataset	89.2%	84.3%	76.6%	83.4%	0	0	0
CL-PBNR	91.9%	90.2%	86.0%	89.4%	6,820	10,754	17,182
CL-MCD	91.9%	90.7%	86.2%	89.6%	4,755	8,399	15,221
CL-MCD-E	92.6%	91.1%	87.8%	90.5%	3,862	6,680	12,707
CL-MCD-Ensemble	92.1%	90.8%	86.4%	89.8%	6,253	10,167	16,758
Alg. Ens. (Agreement=2)	92.4%	90.8%	86.6%	89.9%	4,806	8,451	15,412
Alg. Ens. (Agreement=3)	91.9%	90.8%	86.4%	89.7%	4,259	7,426	13,579
CIFAR-100							
Original dataset	69.6%	66.3%	58.6%	64.8%	0	0	0
CL-PBNR	68.0%	66.1%	63.0%	65.7%	15,860	18,544	23,942
CL-MCD	69.1%	67.1%	63.8%	66.7%	12,549	15,389	20,830
CL-MCD-E	70.3%	69.3%	65.2%	68.3%	9,425	12,218	17,204
CL-MCD-Ensemble	69.1%	67.2%	62.8%	66.4%	15,833	18,757	24,560
Alg. Ens. (Agreement=2)	68.0%	68.0%	64.4%	66.8%	12,148	15,049	20,737
Alg. Ens. (Agreement=3)	70.4%	69.3%	65.0%	68.2%	8,402	10,936	15,562
Tiny-Imagenet							
Original dataset	59.8%	59.8%	60.2%	59.9%	0	0	0
CL-PBNR	56.4%	56.0%	54.0%	55.5%	31,669	35,040	41,883
CL-MCD	57.5%	56.0%	54.5%	56.0%	27,985	31,452	38,545
CL-MCD-E	57.7%	57.2%	55.9%	56.9%	23,174	26,587	33,257
CL-MCD-Ensemble	57.4%	56.4%	53.7%	55.8%	30,818	34,252	41,528
Alg. Ens. (Agreement=2)	57.4%	57.1%	55.6%	56.7%	26,692	30,310	37,598
Alg. Ens. (Agreement=3)	58.8%	57.2%	56.6%	57.5%	18,422	21,647	27,705

Cross-Method Influence of Varying Noise Rates on Label Error Detection. When comparing the performances across different label noise rates in Figure 10, we observe a trend that for lower noise rates (i.e., smaller shares of label errors in the datasets), detecting such mislabeled data becomes more difficult. We observe this trend for CIFAR-10, CIFAR-100, and Tiny-Imagenet. For MNIST, the label error detection performance plateaued and

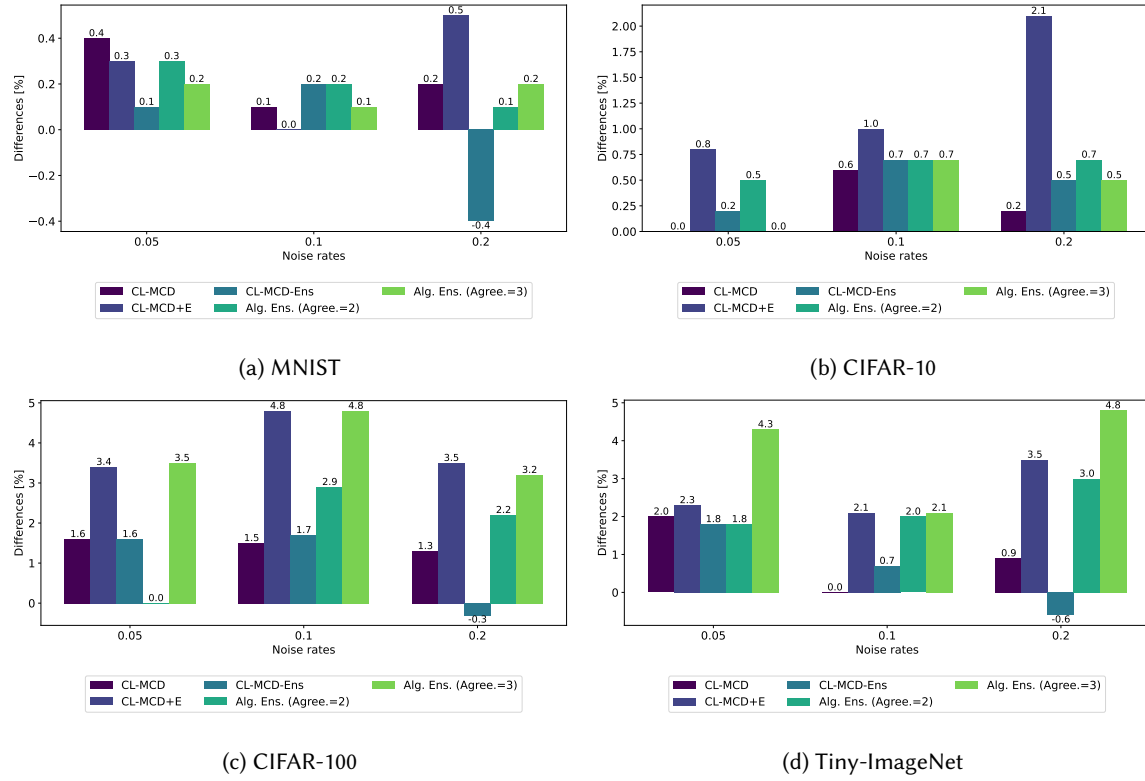


Fig. 9. Comparing performance improvements based on the proposed algorithms compared to CL-PBNR when excluding the label errors identified in Stage 1 in the training datasets. Test datasets remain unchanged for evaluation. Read: At noise rate $\tau_1 = 0.05$, cleaning with the CL-MCD algorithm leads to an accuracy gain of 0.4% compared to cleaning with CL-PBNR.

appears less sensitive to variations in the noise rate. We explain this behavior with a generally high initial model performance on MNIST due to the simplicity of the dataset. In contrast, for datasets where the initial model performance is lower (e.g., CIFAR-100 and Tiny-ImageNet), the trend is particularly pronounced. For the CIFAR-100 dataset, the average F1 score reduces almost by half from 56.7% at $\tau_3 = 0.2$ to 28.9% at $\tau_1 = 0.05$. Similarly, for Tiny-ImageNet, the average F1 score declines from 58.5% at $\tau_3 = 0.2$ to 25.5% at $\tau_1 = 0.05$.

Overall, the findings from this first sensitivity analysis align with previous work (Northcutt, Athalye, et al. 2021) and demonstrate that label error detection becomes more difficult for lower noise rates (i.e. when the relative share of mislabeled data is low). Importantly, our proposed algorithms have demonstrated particularly high performances in such settings (see Evaluation Stage 1) and are tailored to remedy such challenges.

Relationship Between Initial Model Accuracy and Label Error Detection Performance. As a second sensitivity analysis, we investigate the relationship between the initial performance of the ML model and the accuracy of the label error detection. Recall that the initial model accuracy refers to the performance on the original dataset before identifying or removing any potential label errors. In Figure 11, we observe a strong correlation between the initial performance of the ML model and the label error detection performance across the datasets.

Noise Rate τ	0.05	0.1	0.2
MNIST	95.3%	95.1%	94.3%
CIFAR-10	59.4%	63.7%	64.3%
CIFAR-100	28.9%	43.1%	56.7%
Tiny-ImageNet	25.5%	40.7%	58.5%




Fig. 10. Mean F1 score across all algorithms per noise rate and dataset.

	Initial Acc.	F1 Score
MNIST	98.8%	94.9%
CIFAR-10	83.4%	62.5%
CIFAR-100	64.8%	42.9%
Tiny-ImageNet	59.9%	41.6%

Pearson correlation between the two set of scores is $r = 0.98$. Two-sided t-tests with $n = 4$, $df = 2$ result in a test statistic of $t = 6.37$ and a p-value of 0.02.




Fig. 11. Mean initial accuracy of the models and the mean F1 score across all algorithms and noise rates for each dataset together with the correlation and significance measures.

We measure this correlation with the Pearson correlation coefficient r (Schober et al. 2018), which ranges between -1 and 1, where a perfect correlation is represented by $r = 1$. A correlation coefficient of zero indicates the absence of a linear relation between two variables. For the investigated relationship of the initial performance of the ML model and the label error detection performance, we observe a Pearson correlation coefficient of $r = 0.98$, indicating a strong positive correlation. Note that, because r is just an estimate of the population correlation ρ , it is necessary to test for significance (Turney 2022). For this purpose, we formulate the following hypotheses. The null hypothesis (H_0) refers to that the initial model accuracy and the label error detection performance are not significantly linearly correlated ($\rho = 0$). In contrast, in the alternative hypothesis (H_a), both variables are significantly linearly correlated ($\rho \neq 0$). The sample size n corresponds to the number of datasets, i.e., $n = 4$, and the degrees of freedom equals $df = n - 2 = 2$. We leverage the two-tailed t-test to assess the significance of the correlation coefficient at a significance level of $\alpha = 0.05$. We first calculate the test statistic (also called t -value) $t = \frac{r\sqrt{n-2}}{\sqrt{1-r^2}}$ (Cohen 1988, p. 76). We then calculate the p -value, which indicates how likely it is that the correlation between the two variables is random, using a two-sided t -distribution with $df = 2$ (Turney 2022). The test returns a p -value of 0.02, which means we can reject the null hypothesis (H_0) and accept the alternative hypothesis (H_a). This means that the correlation between the initial accuracy of the ML model and the label error detection performance is statistically significant. In other words, a higher initial accuracy of the ML model leads to a higher label error detection performance.

Relationship between Label Error Detection Performance and Final Model Accuracy. In the third sensitivity analysis, we focus on the relationship between the performance of the label error detection and the final model accuracy. Recall that the final model accuracy refers to the performance of the ML model, which was trained on the training dataset after removing detected label errors. In Figure 12, we report the average label error detection performance (F1 score) for each algorithm together with the average final accuracy averaged across all datasets. We observe that algorithms that perform better in detecting label errors lead to an improved final model accuracy. We again measure the Pearson correlation and test the statistical significance.

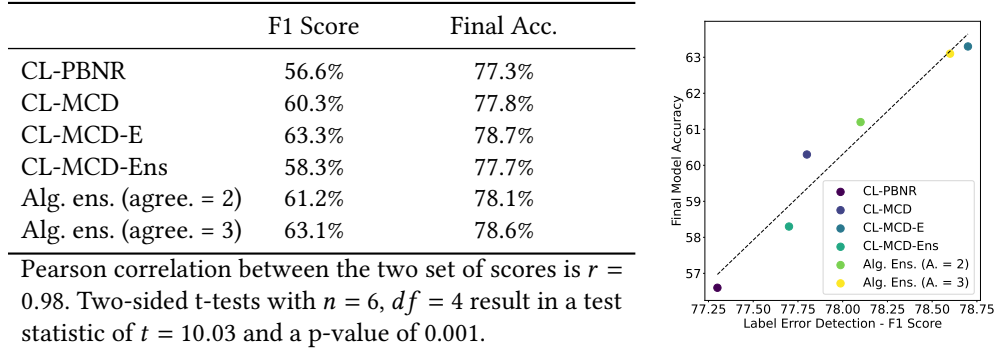


Fig. 12. Correlation between the mean label error detection performance and the mean final accuracy after cleaning all training datasets with the algorithms.

The calculated correlation coefficient for the given sample is $r = 0.98$. This suggests that the label error detection performance and the final model accuracy are strongly correlated—i.e., the better the label error detection, the higher the final model accuracy. To test the significance of the correlation, we again employ a two-sided t-test at a significance level of $\alpha = 0.05$. The sample size n corresponds to the number of algorithms, i.e., $n = 6$, and the degrees of freedom equals $df = n - 2 = 4$. We define the null hypothesis (H_0) as the absence of a significant linear correlation between the label error detection performance and the final model accuracy ($\rho = 0$). The alternative hypothesis (H_a) refers to the presence of a significant linear correlation ($\rho \neq 0$). We obtain a p-value associated with the test statistic of $p = 0.001 < 0.05 = \alpha$, such that the null hypothesis H_0 can be rejected, and the alternative hypothesis H_a can be accepted. Thus, the label error detection performance and the final model accuracy are significantly linearly correlated. In other words, the higher the performance in label error detection, the higher the final model accuracy.

Overall, our sensitivity analyses reveal three important properties of the examined algorithms for label error detection. First, the higher the relative share of label errors in the dataset, the higher the detection accuracy is. Second, the more accurate the models are on the original, uncleaned dataset, the higher the accuracy in the detection of label errors. Third, the higher the accuracy in label error detection is, the higher the model performance on the cleaned datasets. Our analyses underline the importance of removing label errors from datasets to achieve a significant increase in the final model accuracy.

7 Discussion and Conclusion

The goal of this work was to improve state-of-the-art label error detection, i.e., CL, with the help of model uncertainty. For this purpose, different label error detection algorithms based on MCD inference were developed

and experimentally evaluated against a baseline (CL-PBNR) within a two-step evaluation procedure. Using MCD allowed us to incorporate model uncertainty into our algorithms and thus account for it during the label error detection. In addition, ensemble learning techniques were used to further improve the performance and generalizability of our proposed algorithms. After passing the preliminary evaluation on the test sets of four prominent ML benchmark datasets in Stage 1, the developed algorithms were then comprehensively evaluated in Stage 2 on the entire benchmark datasets. Besides the label error detection performance, which we measured by the F1 score, we also assessed the practical implications in terms of (positive) effects on the model accuracy after the removal of the label errors identified. The results of these experiments will be interpreted and discussed in the following. Thereby, we will first examine the label error detection performance before focusing on the implications of our algorithms on the model accuracy. The results obtained in Stage 2 show that our proposed algorithms outperform the baseline (CL-PBNR) in terms of label error detection performance across all datasets and noise rates. The only exceptions, which we will come to later, are the CL-MCD-E algorithm and the Algorithm Ensemble (Agreement=3) at noise rate $\tau_3 = 0.2$ on the MNIST dataset. The reason for the overall better label error detection performance, i.e., the overall better F1 score, can be attributed to the higher precision of all our algorithms across all datasets and noise rates. In the following, we highlight important implications resulting from our work and discuss its limitations, followed by the concluding remarks.

7.1 Replacing Softmax Probabilities with Monte Carlo Dropout Probabilities

The overall higher precision of our algorithms can mainly be justified by the use of MCD. When replacing softmax probabilities (see CL-PBNR) with MCD probabilities (see CL-MCD), the average precision increases by 12.2% across all noise rates and datasets. The reason for the increase in precision is that samples that are difficult to classify for the model (i.e., higher model uncertainty) are filtered out in the first step of CL (*Count*). Thus, these samples cannot be counted as label errors in the confident joint. Consequently, by using MCD probabilities, fewer samples for which the model is uncertain are incorrectly flagged as label errors, which increases the precision of the algorithms compared to CL-PBNR. In contrast to softmax probabilities, MCD probabilities additionally allow the employ of uncertainty measures, such as entropy, on top. This further increases the performance of the label error detection by, on average, 25.8% in precision when comparing the CL-MCD-E algorithm with the baseline. Thus, MCD probabilities not only improve the performance in general, but they additionally facilitate the utilization of additional metrics to refine label error detection.

7.2 Prerequisites for Label Error Detection

An important ingredient for an improvement in the final model accuracy is successful label error detection. Our experiments have demonstrated that the performance in label error detection is correlated with the initial model accuracy. This means that when the initial accuracy of the model is low, as, for example, with the Tiny-ImageNet dataset, the absolute performance in detecting label errors is low as well—eventually preventing an improvement in the final model accuracy. The lower final model performance is reflected in the label error detection performance, where the precision is substantially lower compared to, e.g., the CIFAR-10 dataset. This causes many images to be incorrectly removed, and the final accuracy is negatively affected by removing label errors from the training dataset. Thus, in order to achieve improvements in the final model accuracy, sufficient initial model accuracy is required. Overall, this emphasizes the importance of the complementarity and co-existence of data-centric and model-centric AI, where model-centric AI provides a substantial initial model accuracy, which is required for data-centric detection and removal of label errors. In addition, our experiments revealed that label error detection algorithms are vulnerable to low noise rates (i.e. when only a few label errors exist in the training data). In these cases, we observed the lowest performance of all algorithms. Since in real-world and benchmark datasets, the noise rates typically range between 1% and 10%, it is particularly important to account for this behavior. Thus,

another prerequisite for successful label error detection is algorithms that reliably identify label errors even in very imbalanced settings. We understand UQ-LED algorithms as a first approach in this direction.

7.3 Efficiency Considerations

Our findings show that removing label errors from training datasets can significantly increase the performance of deep learning models. However, researchers and practitioners particularly need to monitor the number of data samples that are removed by the algorithms. Our proposed algorithms generally exhibit a higher precision, which means that fewer samples are removed from the dataset by mistake. If, however, the precision of a label error detection algorithm is low at a very high recall, too many images are incorrectly removed during cleaning. This can have an overall negative effect on model training since there might be too little training data per class available to learn the class representations correctly. Consider the following example: We synthetically add 5,000 label errors to a dataset with 50,000 samples in total (i.e., noise rate of 10%). A label error detection algorithm with a recall of 100% and a precision of 25% will incorrectly remove 25,000 images, such that in the best case, 20,000 correctly labeled images are removed. Then, instead of 50,000 images, the model has only 30,000 correctly labeled images available for training, which might be insufficient to learn the class representations. However, if the algorithm achieves a substantially lower recall of 70% and a slightly higher precision of 35%, the algorithm removes only 10,000 images. This means that in the best case, only 5,000 correctly labeled images are removed, and a resulting 45,000 correctly labeled images are available for model training. These considerations are not only important from a performance perspective but also from an efficiency perspective. Labeling large bodies of data for supervised machine learning is costly, especially in high-stakes domains where human experts have to annotate the data. Removing correctly labeled images from the datasets is associated with significant costs in terms of unnecessary labeling effort. This makes a higher precision in label error detection particularly relevant.

7.4 Applicability Beyond Computer Vision

While this paper demonstrates the effectiveness of UQ-LED algorithms in image classification tasks, the applicability of UQ-LED extends to other domains within machine learning. For instance, natural language processing (NLP) tasks, such as sentiment analysis, frequently encounter label noise due to subjective human annotation processes. Specifically, mislabeled reviews can severely impact the performance and interpretability of sentiment models. UQ-LED can potentially be adapted for sentiment analysis by leveraging uncertainty measures to detect anomalous labels indicative of potential annotation errors. By quantifying the predictive model uncertainty on text data, one could effectively identify reviews for which the sentiment label does not align with the textual cues recognized by the model. Similarly, time series analysis tasks, which commonly experience noisy labels due to temporal data inconsistencies or annotation drift, could benefit from UQ-LED algorithms. Future research could investigate domain-specific adaptations of UQ-LED algorithms, tailoring uncertainty quantification strategies and evaluation methods to effectively handle label noise within these distinct machine learning scenarios.

7.5 Evaluation Strategies in Real-World Scenarios

The evaluation of label error detection methods, including UQ-LED, often apply synthetic noise generation to leverage the full availability of ground truth labels. Real-world settings pose several practical challenges, notably the high cost and extensive time required to obtain verified, ground-truth labels. To address these constraints, alternative evaluation strategies can be helpful. One potential approach is the use of proxy measures, like observing improvements in performance on downstream tasks across several different models following the removal of identified label errors. Additionally, conducting user studies on subsets of the data to qualitatively assess the validity and utility of detected label errors offers another promising approach. These alternative

evaluation strategies could provide valuable insights into the effectiveness of UQ-LED in practical scenarios, complementing the insights obtained from synthetic noise evaluations.

7.6 Limitations

Like any work, this study is not free of limitations. The original CL method of (Northcutt, Jiang, et al. 2021), which we used as the foundation for our own algorithms, is fully model agnostic and can be used not only for neural networks but also, for example, for regression problems, support vector machines, or decision trees. Our proposed algorithms build on Monte Carlo Dropout and, thus, are specific to DNNs. However, we want to emphasize that the algorithms are agnostic to different architectures of deep neural networks since dropout layers can be added to any DNN. In addition, our results may be affected by actual label errors that have been inherently present in the original datasets before adding synthetically generated label noise. This may result in the following two effects on the performance measurements of label error detection performance. First, if label errors have been present before adding synthetic errors and the algorithms successfully detect them, their *precision* would be negatively affected because, in contrast to the synthetically added label errors, the already existing label errors were not marked as such by us before the evaluation. As a result, the absolute precision of all algorithms may be higher than we measured. Notably, this applies to all algorithms, including the CL-PBNR baseline, thus not affecting the comparison of algorithms and our findings. Second, pre-existing label errors could automatically be corrected when generating synthetic label errors. Since human annotators often mislabel classes that are similar, our label noise generator considers class similarities during noise generation. As a result, the label errors unintentionally corrected during the noise generation cannot be identified as such by the different algorithms. Consequently, in case of such an event, the *recall* of the algorithms would be reduced since we have previously marked the now corrected samples as label errors. Again, this limitation applies equally to all algorithms. Thus the overall core message of this work, that MCD-based algorithms with high precision are preferable for label error detection, remains. If future work would like to address this limitation, it would be a good idea to clean the dataset using one of the algorithms before adding label noise. However, we decided against this to avoid influencing our results with the choice of the cleaning algorithm.

7.7 Concluding Remarks

Supervised machine learning is heavily influenced by the quality of the labeled data. Correct labels are an essential prerequisite for accurate and reliable model output. Removing erroneous labels from the training dataset is a vital data-centric approach to improve the model performance, which we enhance by building on methods for uncertainty quantification based on approximations of Bayesian neural networks. Our findings suggest that uncertainty quantification-based label error detection can significantly increase the performance of machine learning models and, thus, underlines the importance of well-curated datasets for supervised machine learning.

References

- I. Alarab, S. Pragoonwit, and M. I. Nacer. Apr. 2021. "Illustrative Discussion of MC-Dropout in General Dataset: Uncertainty Estimation in Bitcoin." *Neural Processing Letters*, 53, 2, (Apr. 2021), 1001–1011. doi:[10.1007/s11063-021-10424-x](https://doi.org/10.1007/s11063-021-10424-x).
- G. Algan and I. Ulusoy. 2020. "Label Noise Types and Their Effects on Deep Learning." *arXiv preprint arXiv:2003.10471*, 0, 0. doi:[10.48550/ARXIV.2003.10471](https://doi.org/10.48550/ARXIV.2003.10471).
- D. Angluin and P. Laird. Apr. 1988. "Learning From Noisy Examples." *Machine Learning*, 2, 4, (Apr. 1988), 343–370. doi:[10.1023/A:1022873112823](https://doi.org/10.1023/A:1022873112823).
- L. Aroyo, M. Lease, P. Paritosh, and M. Schaeckermann. 2022. "Data Excellence for AI: Why Should You Care?" *Interactions*, 29, 2, 66–69.
- C. M. Bishop. 2016. *Pattern Recognition and Machine Learning*. eng. (Softcover reprint of the original 1st edition 2006 (corrected at 8th printing 2009) ed.). Information Science and Statistics. Springer New York, New York, NY. ISBN: 978-1-4939-3843-8.
- H. Cheng, Z. Zhu, X. Li, Y. Gong, X. Sun, and Y. Liu. 2020. "Learning with Instance-Dependent Label Noise: A Sample Sieve Approach." *arXiv preprint arXiv:2010.02347*, 0, 0.

- J. Cheng, T. Liu, K. Ramamohanarao, and D. Tao. 2020. "Learning with Bounded Instance- and Label-dependent Label Noise." In: *International Conference on Machine Learning*. Vol. 1. PMLR, 1789–1799.
- J. Cohen. 1988. *Statistical power analysis for the behavioral sciences*. (2nd ed ed.). L. Erlbaum Associates, Hillsdale, N.J. ISBN: 978-0-8058-0283-2.
- F. R. Cordeiro and G. Carneiro. 2020. "A Survey on Deep Learning with Noisy Labels: How to train your model when you cannot trust on the annotations?" In: *2020 33rd SIBGRAPI conference on graphics, patterns and images (SIBGRAPI)*. Vol. 1. IEEE, 9–16.
- F. Daniel, P. Kucherbaev, C. Cappiello, B. Benatallah, and M. Allahbakhsh. 2018. "Quality Control in Crowdsourcing: A Survey of Quality Attributes, Assessment Techniques and Assurance Actions." *ACM Computing Surveys (CSUR)*, 51, 1, 1–40.
- W. Falcon. 2019. *PyTorch Lightning*. (2019). doi:10.5281/ZENODO.3828935.
- B. Frénay and M. Verleysen. May 2014. "Classification in the Presence of Label Noise: A Survey." *IEEE Transactions on Neural Networks and Learning Systems*, 25, 5, (May 2014), 845–869. doi:10.1109/TNNLS.2013.2292894.
- Y. Gal and Z. Ghahramani. 2016. "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning." In: *International Conference on Machine Learning*, 1050–1059.
- G. R. Gare and J. M. Galeotti. 2021. "Exploiting Class Similarity for Machine Learning with Confidence Labels and Projective Loss Functions." *arXiv preprint arXiv:2103.13607*, 0, 0.
- A. Ghosh, H. Kumar, and P. S. Sastry. 2017. "Robust Loss Functions under Label Noise for Deep Neural Networks." In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 31.
- J. Goldberger and E. Ben-Reuven. 2017. "Training deep neural-networks using a noise adaptation layer." In: *International Conference on Learning Representations*. <https://openreview.net/forum?id=H12GRgcxg>.
- M. Haakman, L. Cruz, H. Huijgens, and A. van Deursen. Sept. 2021. "AI lifecycle models need to be revised: An exploratory study in Fintech." *Empirical Software Engineering*, 26, 5, (Sept. 2021), 95. doi:10.1007/s10664-021-09993-1.
- O. H. Hamid. May 2022. "From Model-Centric to Data-Centric AI: A Paradigm Shift or Rather a Complementary Approach?" In: *2022 8th International Conference on Information Technology Trends (ITT)*. IEEE, Dubai, United Arab Emirates, (May 2022), 196–199. ISBN: 978-1-66548-742-9. doi:10.1109/ITT56123.2022.9863935.
- B. Han, Q. Yao, T. Liu, G. Niu, I. W. Tsang, J. T. Kwok, and M. Sugiyama. 2020. "A Survey of Label-noise Representation Learning: Past, Present and Future." *arXiv preprint arXiv:2011.04406*, 0, 0.
- B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, and M. Sugiyama. 2018. "Co-teaching: Robust Training of Deep Neural Networks with Extremely Noisy Labels." *Advances in neural information processing systems*, 31, 1.
- K. He, X. Zhang, S. Ren, and J. Sun. 2016. "Deep Residual Learning for Image Recognition." In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- D. Hendrycks and K. Gimpel. 2016. "A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks." In: *Proceedings of the International Conference on Learning Representations*.
- D. Hendrycks, K. Lee, and M. Mazeika. 2019. "Using Pre-Training Can Improve Model Robustness and Uncertainty." In: *International conference on machine learning*. PMLR, 2712–2721.
- L. V. Jospin, W. Buntine, F. Boussaid, H. Laga, and M. Bennamoun. May 2022. "Hands-on Bayesian Neural Networks – a Tutorial for Deep Learning Users." *IEEE Computational Intelligence Magazine*, 17, 2, (May 2022), 29–48. Publisher: arXiv. doi:10.1109/MCI.2022.3155327.
- A. Kendall, V. Badrinarayanan, and R. Cipolla. 2015. "Bayesian SegNet: Model Uncertainty in Deep Convolutional Encoder-Decoder Architectures for Scene Understanding." *arXiv preprint arXiv:1511.02680*, 0, 0.
- A. Krizhevsky. 2009. *Learning Multiple Layers of Features From Tiny Images*. Tech. rep. University of Toronto, Toronto.
- Y. Le and X. Yang. 2015. "Tiny Imagenet Visual Recognition Challenge." *Stanford University CS 231N*, 7, 7, 3.
- Y. LeCun, C. Cortes, and C. Burges. 2010. "MNIST handwritten digit database." *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2.
- M. Lukasik, S. Bhojanapalli, A. Menon, and S. Kumar. 2020. "Does Label Smoothing Mitigate Label Noise?" In: *International Conference on Machine Learning*. PMLR, 6448–6458.
- M. Mazumder et al.. July 2022. *DataPerf: Benchmarks for Data-Centric AI Development*. Publisher: arXiv. (July 2022). Retrieved Sept. 28, 2022 from <http://arxiv.org/abs/2207.10062>.
- N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari. 2013. "Learning with Noisy Labels." In: *Advances in Neural Information Processing Systems*. Ed. by C. J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger. Vol. 26. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2013/file/3871bd64012152bfb53fdf04b401193f-Paper.pdf>.
- R. M. Neal. 1996. *Bayesian Learning for Neural Networks*. Lecture Notes in Statistics. Vol. 118. Springer New York, New York, NY. ISBN: 978-0-387-94724-2 978-1-4612-0745-0. doi:10.1007/978-1-4612-0745-0.
- C. G. Northcutt, A. Athalye, and J. Mueller. 2021. "Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks." In: *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021) Track on Datasets and Benchmarks*.
- C. G. Northcutt, L. Jiang, and I. L. Chuang. 2021. "Confident Learning: Estimating Uncertainty in Dataset Labels." *Journal of Artificial Intelligence Research*, 70, 1373–1411.

- A. Paszke et al.. 2019. “Pytorch: An Imperative Style, High-Performance Deep Learning Library.” *Advances in Neural Information Processing Systems*, 32, 8024–8035.
- G. Patrini, A. Rozza, A. Krishna Menon, R. Nock, and L. Qu. July 2017. “Making Deep Neural Networks Robust to Label Noise: A Loss Correction Approach.” In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (July 2017). doi:[10.48550/ARXIV.1609.03683](https://doi.org/10.48550/ARXIV.1609.03683).
- N. Sambasivan, S. Kapania, H. Highfill, D. Akrong, P. Paritosh, and L. M. Aroyo. 2021. ““Everyone Wants to Do the Model Work, Not the Data Work”: Data Cascades in High-Stakes AI.” In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–15.
- P. Schober, C. Boer, and L. A. Schwarte. May 2018. “Correlation Coefficients: Appropriate Use and Interpretation.” en. *Anesthesia & Analgesia*, 126, 5, (May 2018), 1763–1768. doi:[10.1213/ANE.0000000000002864](https://doi.org/10.1213/ANE.0000000000002864).
- J. Shu, Q. Xie, L. Yi, Q. Zhao, S. Zhou, Z. Xu, and D. Meng. 2019. “Meta-Weight-Net: Learning an Explicit Mapping For Sample Weighting.” *Advances in neural information processing systems*, 32.
- K. Simonyan and A. Zisserman. 2014. “Very Deep Convolutional Networks for Large-Scale Image Recognition.” In: *Proceedings of the International Conference on Learning Representations*.
- H. Song, M. Kim, D. Park, Y. Shin, and J.-G. Lee. Mar. 2022. *Learning from Noisy Labels with Deep Neural Networks: A Survey*. Publisher: arXiv. (Mar. 2022). Retrieved Sept. 23, 2022 from <http://arxiv.org/abs/2007.08199>.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. 2014. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting.” *Journal of Machine Learning Research*, 15, 56, 1929–1958. <http://jmlr.org/papers/v15/srivastava14a.html>.
- A. Treiss, J. Walk, and N. Kühl. 2021. “An Uncertainty-Based Human-in-the-Loop System for Industrial Tool Wear Analysis.” en. In: *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track*. Vol. 12461. Series Title: Lecture Notes in Computer Science. Springer International Publishing, Cham, 85–100. ISBN: 978-3-030-67669-8 978-3-030-67670-4. doi:[10.1007/978-3-030-67670-4_6](https://doi.org/10.1007/978-3-030-67670-4_6).
- S. Turney. Sept. 2022. *Pearson Correlation Coefficient (r)*. (Sept. 2022). Retrieved Nov. 21, 2022 from <https://www.scribbr.com/statistics/pearson-correlation-coefficient/>.
- D. Wang, S. Prabhat, and N. Sambasivan. 2022. “Whose AI Dream? In search of the aspiration in data annotation.” In: *CHI Conference on Human Factors in Computing Systems*. ACM, 1–16. doi:[10.1145/3491102.3502121](https://doi.org/10.1145/3491102.3502121).
- W. Wang and Z.-H. Zhou. Nov. 2015. “Crowdsourcing label quality: a theoretical analysis.” en. *Science China Information Sciences*, 58, 11, (Nov. 2015), 1–12. doi:[10.1007/s11432-015-5391-x](https://doi.org/10.1007/s11432-015-5391-x).
- X. Xia, T. Liu, N. Wang, B. Han, C. Gong, G. Niu, and M. Sugiyama. 2019. “Are Anchor Points Really Indispensable in Label-Noise Learning?” *Advances in Neural Information Processing Systems*, 32.
- T. Xiao, T. Xia, Y. Yang, C. Huang, and X. Wang. 2015. “Learning from massive noisy labeled data for image classification.” In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Boston, MA, USA, 2691–2699. ISBN: 978-1-4673-6964-0. doi:[10.1109/CVPR.2015.7298885](https://doi.org/10.1109/CVPR.2015.7298885).
- Z. Zhang and M. Sabuncu. 2018. “Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels.” *Advances in Neural Information Processing Systems*, 31.

Received March 11, 2024; accepted May 30, 2025