

# Analysing the Noise Model Error for Realistic Noisy Label Data

Michael A. Hedderich, Dawei Zhu, Dietrich Klakow

Saarland University, Saarland Informatics Campus, Germany  
 {mhedderich,dzhu,dietrich.klakow}@lsv.uni-saarland.de

## Abstract

Distant and weak supervision allow to obtain large amounts of labeled training data quickly and cheaply, but these automatic annotations tend to contain a high amount of errors. A popular technique to overcome the negative effects of these noisy labels is noise modelling where the underlying noise process is modelled. In this work, we study the quality of these estimated noise models from the theoretical side by deriving the expected error of the noise model. Apart from evaluating the theoretical results on commonly used synthetic noise, we also publish NoisyNER, a new noisy label dataset from the NLP domain that was obtained through a realistic distant supervision technique. It provides seven sets of labels with differing noise patterns to evaluate different noise levels on the same instances. Parallel, clean labels are available making it possible to study scenarios where a small amount of gold-standard data can be leveraged. Our theoretical results and the corresponding experiments give insights into the factors that influence the noise model estimation like the noise distribution and the sampling technique.

## Introduction

One of the factors in the success of deep neural networks is the availability of large, labeled datasets. Where such labeled data is not available, distant and weak supervision have become popular. Related but different to semi-supervised learning, in distant supervision, the unlabeled data is automatically annotated by a separate process using e.g. rules and heuristics created by expert (Ratner et al. 2016) or exploiting the context of images (Mahajan et al. 2018). For tasks like information extraction from text (Mintz et al. 2009), this has become one of the dominant techniques to overcome the lack of labeled data.

While distant supervision allows generating labels in a cheap and fast way, these labels tend to contain more errors than gold standard ones and training with this additional data might even deteriorate the performance of a classifier (see e.g. Fang and Cohn (2016)). Effectively leveraging this noisily-labeled data for machine learning algorithms has, therefore, become a very active field of research. One of the major approaches is the explicit modeling of the noise. This

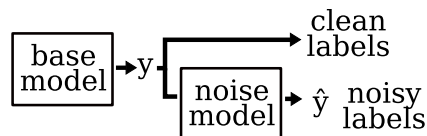


Figure 1: Visualization of the general noise model architecture. The *base model* works directly on the clean data and predicts the clean label  $y$ . For noisily-labeled data, a *noise model* is added after the base model’s predictions.

general concept is task-independent and can be added to existing deep learning architectures. It is visualized in Figure 1. The *base model* is the model that was originally developed for a specific classification task. It is directly used during testing and when dealing with other clean data. When working with noisily-labeled data during training, a *noise model* is added after the base model’s output. The noise model is an estimate of the underlying noise process. The training process of the base model can benefit from it as the noise model can be seen as changing the distribution of the labels from the clean to the noisy. This will be properly defined below.

Many works on noise modeling assume that no manually annotated, clean data is available. The recent trend in few-shot learning and specific works like (Lauscher et al. 2020; Hedderich et al. 2020a) have shown, however, that it is both realistic and beneficial to assume a small amount of manually labeled instances. This motivates us to study in this work scenarios where a small amount of clean, gold-standard data, as well as a large amount of noisily labeled data, are available.

In this paper, we focus on the quality of the noise model. The better such a model is estimated, the better it can model the relationship between clean and noisy labels. We are interested in the factors on which the quality of the noise model depends. We show, both theoretically and empirically, the influence of the clean data and the noise properties. We also propose to adapt the sampling technique to obtain more accurate estimates. Apart from helping to improve the understanding of these noise models in general, we hope that the insights can also be useful guidance for practitioners.

In the noisy labels literature, theoretical insights are often evaluated only on simulated noise processes, e.g. on MNIST with added single-label-flip noise (Reed et al. 2015; Bekker

and Goldberger 2016; Goldberger and Ben-Reuven 2016; Patrini et al. 2017; Han et al. 2018). This synthetic noise has the advantage that it can be controlled completely and allows to rigorously and continuously evaluate aspects like the noise level. However, certain assumptions about the noise have to be taken. And these are usually the same assumptions that are chosen for the noise model itself. It might, therefore, not be too surprising that such a model is suited for such a noise. Recently, efforts have been taken to also evaluate on more realistic scenarios, mostly in the vision domain, e.g. the Clothing1M dataset by Xiao et al. (2015). We want to add to this by making available a noisy label dataset from the natural language processing (NLP) domain based on an existing named entity recognition (NER) corpus. It provides parallel clean and noisy labels for the full data allowing to evaluate different scenarios of resource availability of both clean and noisy data. This new dataset also contains properties that can make learning with noisy labels more challenging such as skewed label distributions and a noise level higher than the true label probability in some settings. In contrast to existing work, we provide seven different sets of noisy labels, each obtained by a realistic noise process via different heuristics in the distant supervision. This makes it possible to experiment with different noise levels for the same instances. The dataset along with the code for the experiments is made publicly available.<sup>1</sup> A version with extended appendix is available online.<sup>2</sup>

### Our key contributions:

- A derivation of the expected error of the noise model estimated from pairs of clean and noisy labels with empirical verification of the derived results on both simulated and realistic noisy labels.
- A set of experiments analyzing how the noise model estimation influences the test performance of the base model.
- NoisyNER, an NLP dataset with noisy labels obtained through non-synthetic, realistic distant supervision that also provides different levels of noise and parallel clean labels.

## Background

We are given a dataset  $D$  consisting of instances  $(x, \hat{y})$  where  $\hat{y}$  is a noisy label that can differ from the unknown, clean/true label  $y$ . Both clean and noisy label have one of  $k$  possible label values/classes. We assume that the change from the clean to the noisy label was performed by a probabilistic noise process described as  $p(\hat{y} = j|y = i)$ . This describes the probability of a true label  $y$  being changed from value  $i$  to the noisy label  $\hat{y}$  with label value  $j$ . With probability  $p(\hat{y} = j|y = j)$  the label value is kept unchanged. This is a common approach to describe noisy label settings. Under this process, a uniform noise (Larsen et al. 1998) with noise level  $\epsilon$  is obtained with

$$p_{\text{uni}}(\hat{y} = j|y = i) = \begin{cases} 1 - \epsilon, & \text{for } i = j \\ \frac{\epsilon}{k-1}, & \text{for } i \neq j \end{cases}, \quad (1)$$

<sup>1</sup><https://github.com/uds-lsv/noise-estimation>

<sup>2</sup><https://arxiv.org/abs/2101.09763>

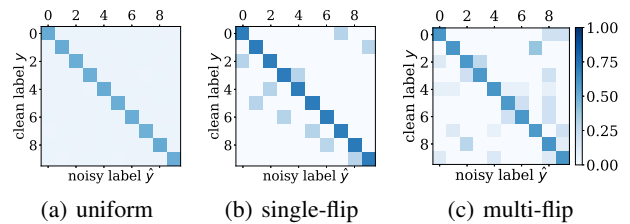


Figure 2: Different noise processes visualized as noise matrices  $M$ : uniform noise ( $\epsilon = 0.5$ ), single-flip noise ( $\epsilon = 0.3$ ) and multi-flip noise ( $\epsilon = 0.4$ ).

and a single label-flip noise (Reed et al. 2015) via

$$p_{\text{nif}}(\hat{y} = j|y = i) = \begin{cases} 1 - \epsilon, & \text{for } i = j \\ \epsilon, & \text{for one } i \neq j \\ 0, & \text{else} \end{cases}. \quad (2)$$

In this work, we will use the more general form that just requires that a noise process can be described with a valid noise transition probability  $p(\hat{y} = j|y = i)$  (Bekker and Goldberger 2016), i.e.

$$\sum_{j=1}^k p(\hat{y} = j|y = i) = 1 \text{ and } p(\hat{y} = j|y = i) \geq 0 \forall i, j. \quad (3)$$

This allows to model more complex noise processes where a true label can be confused with multiple other labels at different noise rates/probabilities. We call this multi-flip noise. This definition generalizes to multi-class classification what in the binary case Natarajan et al. (2013) and Scott, Blanchard, and Handy (2013) describe as class-conditional and asymmetric noise. It is called Markov label corruption by van Rooyen and Williamson (2017).

The noise process can also be described as a matrix  $M \in \mathbb{R}^{k \times k}$  where

$$M_{i,j} = p(\hat{y} = j|y = i). \quad (4)$$

The matrix  $M$  is called confusion or noise transition matrix. See Figure 2 for example matrices with uniform, single-flip and a more complex multi-flip noise.

For a multi-class classification setting, this noise process results in the following relation:

$$p(\hat{y} = j|x) = \sum_{i=1}^k p(\hat{y} = j|y = i)p(y = i|x). \quad (5)$$

This is used to adapt the predictions from the clean label distribution  $p(y = i|x)$  (learned by the base model) to the noisy label distribution  $p(\hat{y} = j|x)$  for the noisily labeled data via the noise process or model  $p(\hat{y} = j|y = i)$ .

Several ways have been proposed to estimate the noise model when only noisy data is available. This includes the use of expert insight (Mnih and Hinton 2012), EM-algorithm (Bekker and Goldberger 2016; Paul et al. 2019; Chen et al. 2019a), backpropagation with regularizers (Sukhbaatar et al.

2015; Luo et al. 2017) and the estimates of a pretrained neural network (Goldberger and Ben-Reuven 2016; Patrini et al. 2017; Dgani, Greenspan, and Goldberger 2018; Wang et al. 2019). For settings where a small portion of clean data can be used, noise model estimation methods have been proposed by Xiao et al. (2015); Fang and Cohn (2016); Hedderich and Klakow (2018); Hendrycks et al. (2018); Lange, Hedderich, and Klakow (2019) and Hedderich et al. (2020a).

Out of the different, similar approaches to estimate a noise model given a small amount of clean labels, we will follow the specific definitions of Hedderich and Klakow (2018) which are based on (Goldberger and Ben-Reuven 2016) and also used in (Lange, Hedderich, and Klakow 2019). The authors assume that a small dataset  $D_C$  with clean labels exists, i.e.  $(x, y) \in D_C$  is known. The size of this clean dataset is much smaller than that of the noisy dataset  $D$ . They then relabel the instances in  $D_C$  with the same mechanism that was used for  $D$  (e.g. distant supervision) to obtain  $\hat{y}$ . This results in a set  $S_{NC}$  of pairs  $(y, \hat{y})$  with  $|S_{NC}| = |D_C|$ . The noise matrix  $M$  is then estimated as  $\widetilde{M}$  using the transitions from  $y$  to  $\hat{y}$  (or equivalently the confusion between  $y$  and  $\hat{y}$ ):

$$\widetilde{M}_{ij} = \frac{m_{ij}}{n_i} = \frac{\sum_{(y, \hat{y}) \in S_{NC}} \mathbb{1}_{\{y=i, \hat{y}=j\}}}{\sum_{(y, \hat{y}) \in S_{NC}} \mathbb{1}_{\{y=i\}}}, \quad (6)$$

where  $m_{ij}$  is the number of times that the label was observed to change from  $i$  to  $j$  due to the noise process and  $n_i$  is the number of instances in  $D_C$  with label  $y = i$ .  $\widetilde{M}$  is the estimated model of the noise process. This noise model is then integrated into the training process using Equation 4 and 5 and as visualized in Figure 1.

### Expected Error of the Noise Model

The noise model obtained in Equation 6 is an approximation of the underlying true noise process estimated on a small number of instances  $D_C$ . In this section, we derive a formula for the expected error of the estimated noise model  $\widetilde{M}$ . This gives us insights into the factors that influence the noise model’s quality as well as their effect.

**Assumptions** In the following proofs, we assume that  $M$  describes a noise process following Equations 3 and 4.  $\widetilde{M}_{ij}$  is estimated using Equation 6.

We study two sampling techniques on how to obtain the set of clean and noisy label pairs  $|S_{NC}|$ . Commonly, a fixed number of unlabeled instances  $n$  is obtained and then manually annotated with gold labels. The value of  $n_i$  then follows the distribution of classes in the data. We call this **Variable Sampling** as the value of  $n_i$  varies.

In contrast to that, for **Fixed Sampling**, for each label value  $i$ , we sample  $n_i$  instances with  $y = i$ . This could be conducted e.g. by asking annotators to provide a specific number of labeled instances per class. In this case,  $n_i$  is fixed. For readability, we write  $\mathbb{E}$  for  $\mathbb{E}_{\sim S_{NC}}$  and analogously for  $\text{Var}$  and  $\text{Cov}$ . We assume that the instances are sampled independently.

As quality metric for evaluating the noise model, we use squared error which is in this matrix case the square of the Frobenius norm

$$SE = \|M - \widetilde{M}\|_F^2 = \sum_{i=1}^k \sum_{j=1}^k (M_{ij} - \widetilde{M}_{ij})^2. \quad (7)$$

#### Theorem 1

The expected squared error of the noise model is

$$\mathbb{E}[SE] = \sum_{i=1}^k \sum_{j=1}^k \text{Var}[\widetilde{M}_{ij}].$$

*Sketch of the proof:* We show that  $\widetilde{M}$  is an unbiased estimator, i.e.  $\mathbb{E}[\widetilde{M}_{ij}] = M_{ij}$ . From that, Theorem 1 can be followed. For space reasons, the full proofs are given in the Appendix.

**Theorem 2a** Assuming *Variable Sampling*, it holds  $\text{Var}[\widetilde{M}_{ij}] = M_{ij}(1 - M_{ij}) \sum_{n_i=1}^n P(n_i) \frac{1}{n_i}$  where  $P(n_i)$  is the probability of sampling  $n_i$  instances with label  $y = i$  from the data.

**Theorem 2b** Assuming *Fixed Sampling*, it holds  $\text{Var}[\widetilde{M}_{ij}] = \frac{M_{ij}(1 - M_{ij})}{n_i}$ .

*Sketch of the proof:* The proofs for both variants of the theorem work on the main insight that given  $n_i$ , the value of  $m_{ij}$  follows a multinomial distribution defined by  $M_{ij}$ .

Combining Theorem 1 and 2, we obtain a closed-form solution for the expected error of the estimated noise model for both Fixed and Variable Sampling. From this, we can see that

- the error changes with the amount of sampled instances by factor  $\frac{1}{n_i}$ .
- the error depends on the noise distribution as well as the level of noise  $M_{ij}$ . In the single-flip scenario, it reaches its maximum when the noise is as dominant as the true label value.
- Fixed Sampling obtains lower error than Variable Sampling in most cases.

These results are visualized and experimentally verified below.

### Data with Synthetic Noise

Experiments with synthetic or simulated noise allow fine-grained control of the noise level and type of noise. An existing dataset is taken and the labels are assumed to be all correct and clean. Then, to obtain a noisy label dataset, for each instance, the label is flipped according to the noise process. Reed et al. (2015) and Goldberger and Ben-Reuven (2016) use the MNIST dataset (LeCun et al. 1998) and apply single-flip noise (Equation 2) to obtain the noisy labels. We follow their approach and label-flip pattern. Additionally, we also generate noisy labels with uniform noise (Equation 1) and a more complex, multi-flip noise where one label can be changed into one of several incorrect labels (Equation in the Appendix). All three noise types are visualized in Figure 2. We see the multi-flip noise as the most realistic of these synthetic noises, as it resembles most the two realistic datasets presented in the next section.

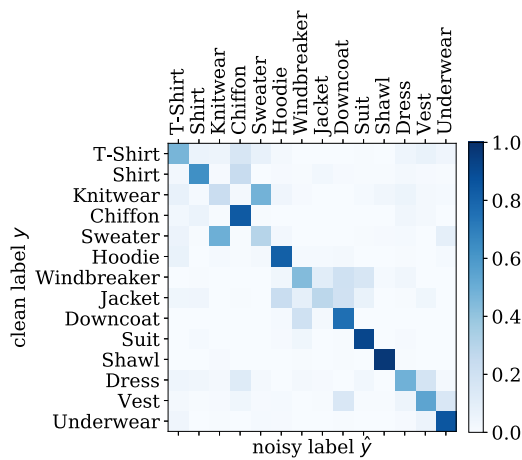


Figure 3: Noise matrix for Clothing1M computed over all pairs of clean and noisy labels.

### Data with Realistic Noise

While evaluating on synthetically generated noise is popular and allows for an easy evaluation in a controlled environment, it is limited by the assumptions on the noise. Certain assumptions are taken when building a model of the noise and the same assumptions are used to generate the noisy labels.

In real-world scenarios, some of these assumptions might not apply. Inspecting realistic noise matrices (Figures 3 and 4), it is already quite obvious that these do not resemble the popular uniform or single-flip noise. We think it is therefore important to also evaluate on more realistic data that does not rely on the noise being simulated. Nevertheless, having parallel clean and noisy labels for the same instances is very useful as it allows e.g. to compute the upper bounds of training on clean data compared to training with noisy labels. In this specific work, it is required to obtain an approximation of the true noise pattern and to flexibly vary the number of clean labels. The Clothing1M dataset by Xiao et al. (2015) and our newly proposed NoisyNER dataset offer these possibilities.

### Clothing1M

The Clothing1M dataset is part of a classification task to label clothing items present in an image. The noisy labels were obtained through a distant supervision process that used the text context of the images appearing on a shopping website. For 37k images, both clean and noisy labels are available. The percentage of correct labels in the noisy data is 38% and a visualization of the noise is given in Figure 3. One can see that the noise distributes neither uniformly nor is there a single label flip. Rather a label tends to be confused with several other related labels, e.g. a "Jacket" with a "Hoodie" and a "Downcoat".

### NoisyNER

In this work, we propose another noisy label dataset. It is from the text classification domain with word-level labels

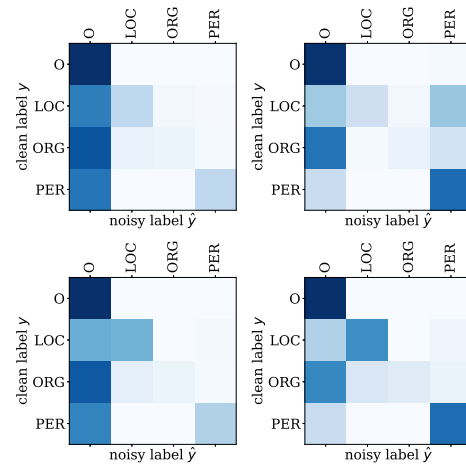


Figure 4: Noise matrices for NoisyNER (label sets 1, 3, 4 and 7) computed over all pairs of clean and noisy labels.

for named entity recognition (NER). The labels are persons, locations and organizations. The language is Estonian, a typical low-resource language with a demand for natural language processing techniques. The text and the clean labels were collected by Laur (2013) through expert annotations (Tkachenko, Petmanson, and Laur 2013). The noisy labels are obtained through a distant supervision/automatic annotation approach. Using a knowledge base for the distant supervision was first introduced by Mintz et al. (2009) and is still commonly used, e.g. by Lison et al. (2020). In our case, lists of named entities were extracted from Wikidata and matched against the words in the text via the ANEA tool (Hedderich, Lange, and Klakow 2021). If a word/string appears in a list of entities, it is labeled as the corresponding entity. This allows to quickly annotate large text corpora without manually labeling each word. However, these labels are not error-free. Reasons include non-complete lists of entities, grammatical complexities of Estonian that do not allow simple string matching or entity lists in conflict with each other (e.g. "Tartu" is both the name of an Estonian city and a person name). Heuristic functions allow to leverage insights from experts efficiently (Ratner et al. 2020) and they can also be applied to correct some of these error sources, e.g. by normalizing (lemmatizing) the grammatical form of a word or by excluding certain high false-positive words. Details are given in the Appendix.

In contrast to Clothing1M, we provide seven sets of labels that differ in the noise process and therefore in the amount of noise and the noise pattern. Four of these are visualized in Figure 4 with the rest in the Appendix. Table 1 lists an overview of the percentage of correct labels. The different labels are obtained by using varying amounts of heuristics during the automatic annotation process. This reflects different levels of manual effort that one might be able to spend on the creation of distantly supervised data. Having these multiple sets of labels for the same instances allows to directly evaluate for different noise levels while excluding side effects that differing features might have.

LABEL SET	1	2	3	4	5	6	7
PRECISION	67	73	37	75	48	53	59
RECALL	18	27	31	27	41	41	49
F1	28	39	34	40	44	46	54

Table 1: Percentages of correct labels in the NoisyNER dataset for the seven different label sets. For each subsequent label set, more manual effort in the labeling heuristics was invested. As the label distribution is highly skewed, precision, recall and F1 score are reported. Following the standard approach for named entity recognition (Tjong Kim Sang and De Meulder 2003), the micro-average is computed excluding the non-entity label.

We want to highlight several properties of the dataset:

- Clean labels are available for all instances. This allows studying different splits of clean and noisy labels as well as computing upper bounds of performance on only clean data.
- The distribution of the labels is skewed. Out of the ca. 217k instances, ca. 8k are persons, 6k are locations and 6k are organizations. All other instances are labeled as non-entity O. Such a skewed label distribution is typical for a named entity recognition dataset.
- In past works, experiments are often only performed until the probability of a noisy class reaches the probability of the true class, i.e. it is assumed that  $p(\hat{y} = j|y = i) < p(\hat{y} = i|y = i) \forall j \neq i$ . This assumption does not hold for several of our label sets which can make learning on the data more challenging.
- While not studied in this work, the labels in the dataset also contain sequential dependencies. A clean or noisy label can span over several words/instances, e.g. for the mention of a person with a first and last name. These sequential dependencies could be leveraged in future work.

### Analysis of the Noise Model Error

In this section, the theoretically expected squared error between the noise model estimate and the true noise matrix is compared to the empirically measured one. We vary the two parameters found in Theorem 2: the amount of sampled data  $n_i/n$  and the amount of noise  $M_{ij}$ . For Clothing1M only the data size can be varied while for NoisyNER the variation of the sample size can be compared across different noise levels and noise distributions.

### Experimental Setup

From the full dataset, a small subset  $D_C$  and corresponding  $S_{NC}$  is sampled uniformly at random either using the Fixed or Variable Sampling approach. The noise model  $\hat{M}$  is estimated on this sample and compared to the true noise process  $M$ . The process is repeated 500 times and the average empirical squared error is reported as well as its standard deviation on the error bars.

For the synthetic noisy labels, the true noise process is known by construction. For the realistic datasets, the true noise process is unknown. Instead, the noise matrix  $M$  is computed over the whole data as an approximation. For the distribution  $P(y)$ , that is part of the Variable Sampling formula, we assume a uniform distribution for MNIST and a multinomial distribution for Clothing1M and NoisyNER with the parameters of the distribution estimated over the whole dataset. In praxis, one might also rely on expert knowledge for this distribution (e.g. from Augenstein, Derczynski, and Bontcheva (2017) for NER).

### Results & Analysis

On the synthetic data (Figure 5), the theoretically expected error of the noise model follows closely the empirical measurements. This holds across different noise types, sample sizes and noise levels. Only for a very small set of clean labels in combination with a high noise level, there is a slight deviation.

As stated above, we can see the influence of the sample size and noise process. The error changes with the amount of sampled instances by factor  $\frac{1}{n_i}$  and it depends on the noise distribution as well as the level of noise  $M_{ij}$ . For the evaluated scenarios, due to the additional dependency on the clean label distribution  $P(y)$ , the Variable Sampling technique has a higher expected error than the Fixed Sampling approach, especially for settings with large noise. From the empirical experiments, one can see that the variance of the noise model error mostly depends on the sample size.

The theoretical and experimental results also match on the realistic noisily labeled datasets Clothing1M and NoisyNER (Figure 6). Again, only for the very low sample size, one can observe a deviation. It is interesting to note that for NoisyNER the estimation error of the noise model is higher for the data with overall lower noise level (measured in F1 score in Table 1). This is due to how the noise distribution changes in the realistic setting. The difference between Fixed Sampling and Variable Sampling is most noticeable for NoisyNER increasing the error by a factor of around 3 (cf. plots in the Appendix). This suggests that in practice, especially for such skewed distributions, a sampling technique is beneficial which focuses on each label separately.

### Analysing the Base Model Performance

In the previous sections, we studied the factors on which the quality of the noise model’s estimate depends. The noise model is part of a larger classifier and it is combined with the actual base model that performs the task-specific classification (cf. Figure 1). In this section, we evaluate in different experiments the base model on clean test data and analyze the effects of the noise modelling and the noise model estimation on the base model and its task performance.

As in the experimental setup detailed in the previous section, a clean subset  $D_C$  of the data is sampled and the noise model estimated on it. The base model is then trained directly on the clean data  $D_C$ . The noise model is added to the base model for training on the noisy dataset  $D$  (cf. Figure 1). For evaluation, a test set is held-out from the data.

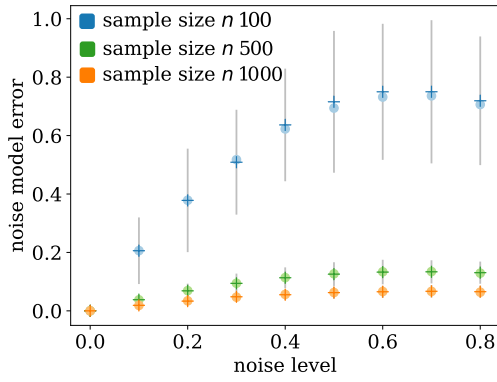
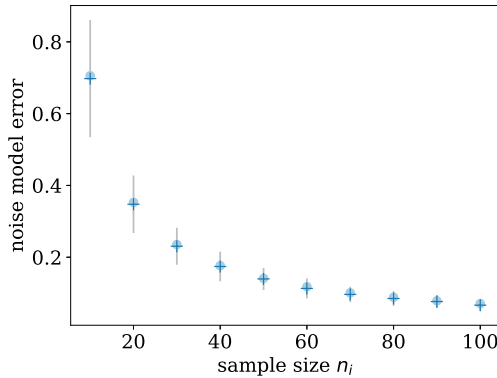
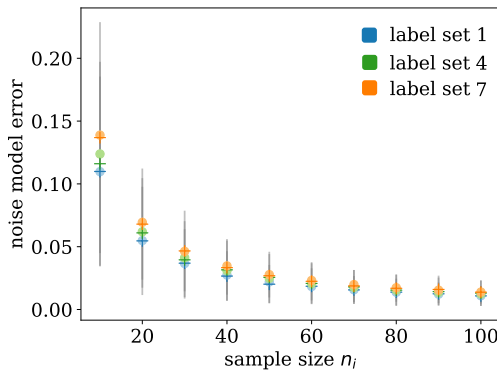


Figure 5: Comparison between the theoretically expected mean error (circle marker) and the empirically measured mean error (cross) of the noise model on the MNIST dataset with *multi-flip noise* and with *Variable Sampling*. The x-axis varies the *noise level*. Plots with all noise types, for Fixed Sampling and for sample size on the x-axis are given in the Appendix.



(a) Clothing1M



(b) NoisyNER

Figure 6: Comparison between the theoretically expected mean error (circle marker) and the empirically measured mean error (cross) of the noise model with *Fixed Sampling*. The x-axis varies the *sample size*. Variable Sampling and the remaining label sets can be found in the Appendix.

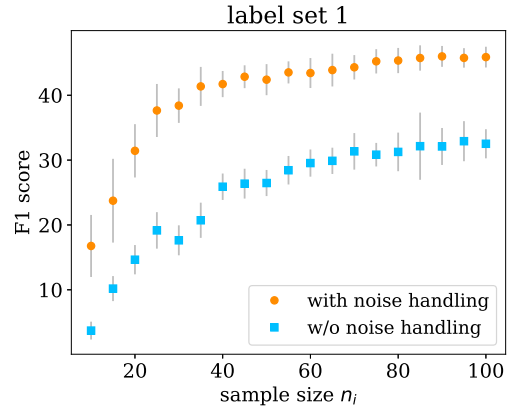


Figure 7: Mean test performance (F1 score) of the base model on NoisyNER label set 1 on clean and noisy data with and without noise handling. Using Fixed Sampling. Further plots in the Appendix.

We follow the training procedure of Hedderich and Klakow (2018) and the training details are given in the Appendix. Due to the longer runtimes, the experiments are repeated 50 times for Clothing1M and NoisyNER. The error bars show the empirical standard deviation.

### Effect of Noise Handling

Figure 7 shows the test performance of the base model for increasing size of  $D_C$ . It compares training the base model directly on the clean and the noisy label data to handling the noisy labels via a noise model. The latter improves the results in most cases. This confirms past findings that noise handling is an important technique to leverage distantly supervised training data. Comparing the different noise levels on NoisyNER, one can see that larger noise levels also result in larger improvements via noise handling. Only in a few cases with a very small amount of clean training samples does noise handling deteriorate the results, possibly due to a bad estimation of the noise model (see Appendix).

### Relationship between Noise Estimation and Base Model Performance

There are several factors on which the base model's performance depends. These could include the amount of clean and noisy training data, the noise distribution and the quality of the estimated noise model. Here, we show experimentally on Clothing1M and NoisyNER that the noise model estimation error directly influences the performance of the base model.

In Figure 8, the expected noise model error is plotted against the test performance of the base model for Fixed Sampling. They show a clear negative correlation. The influence of the noise distribution is visible in the different slopes in the plots for the different noisy label sets of NoisyNER. For all settings, the Pearson Correlation between the mean test performance and the expected noise model error is at least  $-0.96$ .



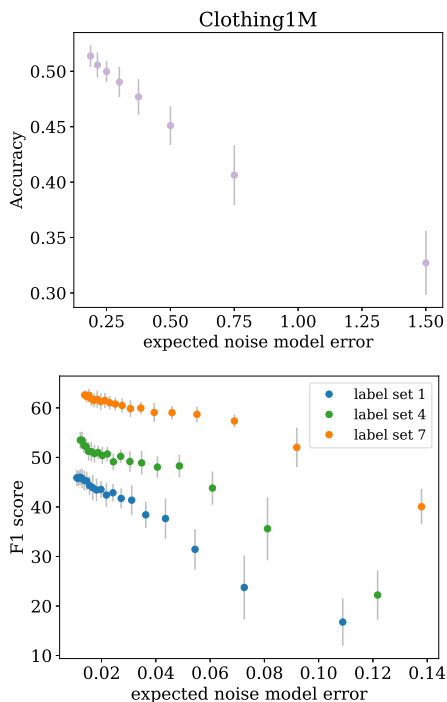


Figure 8: Relationship between the theoretically expected noise model error and the test performance (Accuracy/F1 score) of the base model for Clothing1M and NoisyNER (label set 1, 4 and 7) with Fixed Sampling. Each point corresponds to one sample size  $n_i$  (cf. Figure 7). Additional plots in the Appendix.

Increasing the number of clean labels is not only beneficial to the estimation of the noise model but also to the base model itself. To remove this effect, the experiment is repeated with a fixed amount of training instances for the base model ( $n_i = 50$  for NoisyNER and  $n_i = 25$  for Clothing1M) and a varying amount of clean labels for the noise model estimation (plots in the Appendix). The same linear relationship can be seen. The Pearson Correlation is again at least  $-0.96$  for all settings.

### Effect of Sampling during Estimation

Above, we saw both from the theoretical analysis as well as in the experiments that Variable versus Fixed Sampling has a strong influence on the noise model estimation error. Figure 9 shows that this effect also transfers to the performance of the base model on the test set. Fixed Sampling consistently outperforms Variable Sampling on the average performance across different noise levels. It also reduces the variance of the reached test performance, another important issue models trained on noisy labels suffer from. The same holds again when a fixed amount of training instances for the base model is used to remove the effect of just having more clean data.

As far as we are aware, Variable Sampling is more common in the literature, suggesting that Fixed Sampling could be a useful and model-independent way to improve noise handling.

### Other Related Work

Apart from the publications mentioned in the previous sections, we want to highlight the following related work. The survey by Frenay and Verleysen (2014) gives a comprehensive overview of the literature on learning with noisily labeled data. More recently, Algan and Ulusoy (2019) and Hedderich et al. (2020b) surveyed deep learning techniques for noisy labels in the image classification domain and the NLP domain respectively.

Other aspects of learning with noisy labels have been studied from a theoretical viewpoint. For the binary classification setting, Liu and Tao (2016a) prove an upper bound for the noise rate. Natarajan et al. (2013) and Patrini et al. (2016) derive upper bounds on the risk when learning in a similar noise scenario. Assuming a known and correct noise model, van Rooyen and Williamson (2017) derive upper and lower bounds on learning with noisy labeled data. The existence of anchor points is an important assumption for several recent works on estimating a noise matrix, i.a. (Liu and Tao 2016b; Patrini et al. 2017; Xia et al. 2019). Yao et al. (2020) propose to learn an intermediate class to reduce the estimation error of the noise matrix. Chen et al. (2019b) propose a formula to calculate the test accuracy assuming, however, that the test set is noisy as well.

On the empirical side, Veit et al. (2017) model the relationship  $p(y|\hat{y})$ , i.e. the opposite of the noise models in this work. They also use pairs of clean and corresponding noisy labels. Rahimi, Li, and Cohn (2019) include the concept of noise transition matrices in a Bayesian framework to model several sources of noisy supervision. In their few-shot evaluation, they also leverage pairs of clean and noisy labels. Ortego et al. (2019) study the behaviour of the loss for different noise types. Some recent works have taken instance-dependent (instead of only label-dependent) noise into consideration both from a theoretic and an empiric viewpoint (Luo et al. 2017; Menon, van Rooyen, and Natarajan 2018; Cannings, Fan, and Samworth 2018; Cheng et al. 2020; Xia et al. 2020). Luo et al. (2017) successfully modeled instance-dependent noise for an information extraction task, making instance-dependence an interesting future work for the NoisyNER dataset.

In the image classification domain, further noisy label datasets exist. The distantly supervised WebVision (Li et al. 2017) was obtained by searching for images on Google and Flickr where the labels (and related words) are used as search keywords. It contains 2.4 million images but it lacks clean labels for the training data. The Food101N dataset (Lee et al. 2018) was created similarly, focusing, like Clothing1M, on a specific image domain. The authors provide ca. 300k images, ca. 50k of which have additional, human-annotated labels. The noise rate is lower with around 20%. The older Tiny Images dataset (Torralba, Fergus, and Freeman 2008) consists of 80 million images at a very low resolution with 3526 training images having clean labels. Recently, Jiang et al. (2020) proposed artificially adding different amounts of incorrectly labeled images to existing datasets for evaluation. For the task of text sentiment analysis, Wang et al. (2019) proposed a dataset that obtained sentence level labels by exploiting document level ratings.

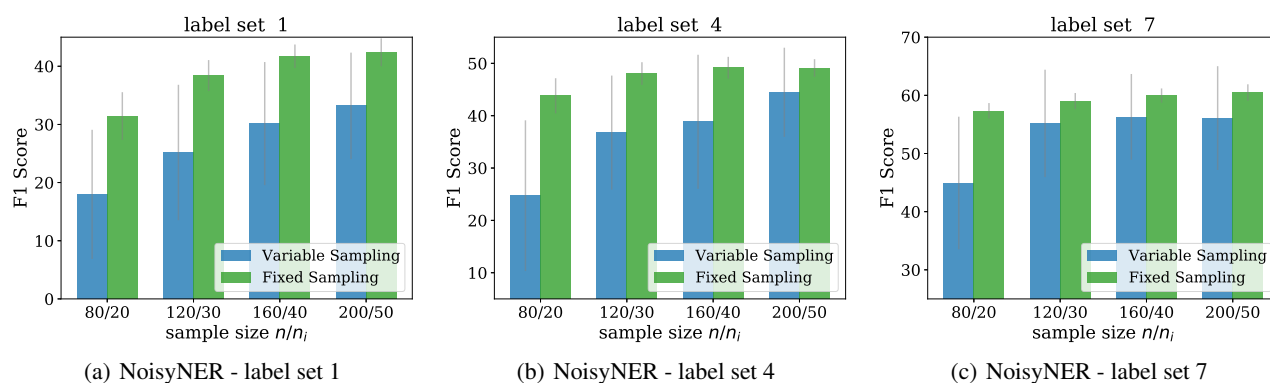


Figure 9: Comparing Variable and Fixed Sampling by mean test performance (F1 score) of the base model on NoisyNER label set 1, 4 and 7. The two approaches were given the same amount of clean data with  $n_i = n/k$  for Fixed Sampling. Further plots in the Appendix.

Clean sentence level labels exist for two of their studied text domains.

## Conclusion

In this work, we analyzed the factors that influence the estimation of the noise model in noisy label modelling for distant supervision. We derived the expected error of a noise model estimated from pairs of clean and noisy labels highlighting factors like noise distribution and sampling technique. Extensive experiments on synthetic and realistic noise showed that it matches the empirical error. We also analyzed how the noise model estimation affects the test performance of the base model. Our experiments show the strong influence of the noise model estimation and how theoretical insights on e.g. the sampling technique can be used to improve the task performance. A major contribution of this work is our newly created NoisyNER dataset, a named entity recognition dataset consisting of seven sets of labels with differing noise levels and patterns. It allows a simple and extensive evaluation of noisy label approaches in a realistic setting and for different levels of resource availability. With its interesting properties, we hope that it is useful for the community to develop noise-handling methods that are applicable to real scenarios.

## Acknowledgments

We would like to thank the reviewers as well as Gabriele Bettgenhuser and Alexander Blatt for their extensive and helpful feedback. This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Project-ID 232722074 SFB 1102 and the EU Horizon 2020 project ROXANNE under grant number 833635.

## Ethics Statement

Distant supervision and handling of noisy labels reduce the need for costly, labeled supervision. While this could lower the entrance bar for nefarious uses of machine learning, it also makes it possible to train and use machine learning approaches in low-resource scenarios such as under-resourced

languages or applications developed by individuals or small organizations. It can, therefore, be a step towards the democratization of AI.

## References

- Algan, G.; and Ulusoy, I. 2019. Image Classification with Deep Learning in the Presence of Noisy Labels: A Survey. *CoRR* abs/1912.05170.
- Augenstein, I.; Derczynski, L.; and Bontcheva, K. 2017. Generalisation in named entity recognition: A quantitative analysis. *Computer Speech & Language* 44: 645–655.
- Bekker, A. J.; and Goldberger, J. 2016. Training deep neural networks based on unreliable labels. In *Proceedings ICASSP 2016*, 2682–2686.
- Cannings, T. I.; Fan, Y.; and Samworth, R. J. 2018. Classification with imperfect training labels. *CoRR* abs/1805.11505.
- Chen, J.; Zhang, R.; Mao, Y.; Guo, H.; and Xu, J. 2019a. Uncover the Ground-Truth Relations in Distant Supervision: A Neural Expectation-Maximization Framework. In *Proceedings of EMNLP-IJCNLP 2019*, 326–336.
- Chen, P.; Liao, B.; Chen, G.; and Zhang, S. 2019b. Understanding and Utilizing Deep Neural Networks Trained with Noisy Labels. In *Proceedings of ICML 2019*, 1062–1070.
- Cheng, J.; Liu, T.; Ramamohanarao, K.; and Tao, D. 2020. Learning with Bounded Instance and Label-dependent Label Noise. In *Proceedings of ICML 2020*, 1789–1799.
- Dgani, Y.; Greenspan, H.; and Goldberger, J. 2018. Training a neural network based on unreliable human annotation of medical images. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, 39–42.
- Fang, M.; and Cohn, T. 2016. Learning when to trust distant supervision: An application to low-resource POS tagging using cross-lingual projection. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, 178–186.



- Frenay, B.; and Verleysen, M. 2014. Classification in the Presence of Label Noise: A Survey. *IEEE Transactions on Neural Networks and Learning Systems* 25(5): 845–869.
- Goldberger, J.; and Ben-Reuven, E. 2016. Training deep neural-networks using a noise adaptation layer. In *Proceedings of ICLR 2016*.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I. W.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Proceedings of NeurIPS 2018*, 8536–8546.
- Hedderich, M. A.; Adelani, D.; Zhu, D.; Alabi, J.; Markus, U.; and Klakow, D. 2020a. Transfer Learning and Distant Supervision for Multilingual Transformer Models: A Study on African Languages. In *Proceedings of EMNLP 2020*, 2580–2591.
- Hedderich, M. A.; and Klakow, D. 2018. Training a Neural Network in a Low-Resource Setting on Automatically Annotated Noisy Data. In *Proceedings of the Workshop on Deep Learning Approaches for Low-Resource NLP at ACL 2018*, 12–18.
- Hedderich, M. A.; Lange, L.; Adel, H.; Strtgen, J.; and Klakow, D. 2020b. A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios. *CoRR* abs/2010.12309.
- Hedderich, M. A.; Lange, L.; and Klakow, D. 2021. ANEA: Distant Supervision for Low-Resource Named Entity Recognition. *CoRR* abs/2102.13129.
- Hendrycks, D.; Mazeika, M.; Wilson, D.; and Gimpel, K. 2018. Using Trusted Data to Train Deep Networks on Labels Corrupted by Severe Noise. In *Proceedings of NeurIPS 2018*, 10477–10486.
- Jiang, L.; Huang, D.; Liu, M.; and Yang, W. 2020. Beyond synthetic noise: Deep learning on controlled noisy labels. In *Proceedings of ICML 2020*, 4804–4815.
- Lange, L.; Hedderich, M. A.; and Klakow, D. 2019. Feature-Dependent Confusion Matrices for Low-Resource NER Labeling with Noisy Labels. In *Proceedings of EMNLP-IJCNLP 2019*, 3552–3557.
- Larsen, J.; Nonboe, L.; Hintz-Madsen, M.; and Hansen, L. K. 1998. Design of robust neural network classifiers. In *Proceedings of ICASSP 1998*, 1205–1208.
- Laur, S. 2013. Nimeksuste korpus. Center of Estonian Language Resources. doi:10.15155/1-00-0000-0000-0000-00073L.
- Lauscher, A.; Ravishankar, V.; Vulic, I.; and Glavas, G. 2020. From Zero to Hero: On the Limitations of Zero-Shot Language Transfer with Multilingual Transformers. In *Proceedings of EMNLP 2020*, 4483–4499.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11): 2278–2324.
- Lee, K.; He, X.; Zhang, L.; and Yang, L. 2018. CleanNet: Transfer Learning for Scalable Image Classifier Training With Label Noise. In *Proceedings of CVPR 2018*, 5447–5456.
- Li, W.; Wang, L.; Li, W.; Agustsson, E.; and Gool, L. V. 2017. WebVision Database: Visual Learning and Understanding from Web Data. *CoRR* abs/1708.02862.
- Lison, P.; Barnes, J.; Hubin, A.; and Touileb, S. 2020. Named Entity Recognition without Labelled Data: A Weak Supervision Approach. In *Proceedings of ACL 2020*, 1518–1533.
- Liu, T.; and Tao, D. 2016a. Classification with Noisy Labels by Importance Reweighting. *IEEE Trans. Pattern Anal. Mach. Intell.* 38(3): 447–461.
- Liu, T.; and Tao, D. 2016b. Classification with Noisy Labels by Importance Reweighting. *IEEE Trans. Pattern Anal. Mach. Intell.* 38(3): 447–461.
- Luo, B.; Feng, Y.; Wang, Z.; Zhu, Z.; Huang, S.; Yan, R.; and Zhao, D. 2017. Learning with Noise: Enhance Distantly Supervised Relation Extraction with Dynamic Transition Matrix. In *Proceedings of ACL 2017*, 430–439.
- Mahajan, D.; Girshick, R.; Ramanathan, V.; He, K.; Paluri, M.; Li, Y.; Bharambe, A.; and van der Maaten, L. 2018. Exploring the Limits of Weakly Supervised Pretraining. In Ferrari, V.; Hebert, M.; Sminchisescu, C.; and Weiss, Y., eds., *Computer Vision – ECCV 2018*, 185–201. Springer International Publishing.
- Menon, A. K.; van Rooyen, B.; and Natarajan, N. 2018. Learning from binary labels with instance-dependent noise. *Machine Learning* 107(8-10): 1561–1595.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of ACL 2009*, 1003–1011.
- Mnih, V.; and Hinton, G. 2012. Learning to Label Aerial Images from Noisy Data. In *Proceedings of ICML 2012*.
- Natarajan, N.; Dhillon, I. S.; Ravikumar, P.; and Tewari, A. 2013. Learning with Noisy Labels. In *Proceedings of NeurIPS 2013*, 1196–1204.
- Ortego, D.; Arazo, E.; Albert, P.; O’Connor, N. E.; and McGuinness, K. 2019. Towards Robust Learning with Different Label Noise Distributions. *CoRR* abs/1912.08741. URL <http://arxiv.org/abs/1912.08741>.
- Patrini, G.; Nielsen, F.; Nock, R.; and Carioni, M. 2016. Loss factorization, weakly supervised learning and label noise robustness. In *Proceedings of ICML 2016*, 708–717.
- Patrini, G.; Rozza, A.; Menon, A. K.; Nock, R.; and Qu, L. 2017. Making Deep Neural Networks Robust to Label Noise: A Loss Correction Approach. In *Proceedings of CVPR 2017*, 2233–2241.
- Paul, D.; Singh, M.; Hedderich, M. A.; and Klakow, D. 2019. Handling Noisy Labels for Robustly Learning from Self-Training Data for Low-Resource Sequence Labeling. In *Proceedings of NAACL 2019: Student Research Workshop*, 29–34.
- Rahimi, A.; Li, Y.; and Cohn, T. 2019. Massively Multilingual Transfer for NER. In *Proceedings of ACL 2019*, 151–164.

- Ratner, A.; Bach, S. H.; Ehrenberg, H. R.; Fries, J. A.; Wu, S.; and Ré, C. 2020. Snorkel: rapid training data creation with weak supervision. *VLDB J.* 29(2-3): 709–730.
- Ratner, A. J.; De Sa, C. M.; Wu, S.; Selsam, D.; and Ré, C. 2016. Data Programming: Creating Large Training Sets, Quickly. In Lee, D. D.; Sugiyama, M.; Luxburg, U. V.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 29*, 3567–3575. Curran Associates, Inc.
- Reed, S. E.; Lee, H.; Anguelov, D.; Szegedy, C.; Erhan, D.; and Rabinovich, A. 2015. Training Deep Neural Networks on Noisy Labels with Bootstrapping. In *ICLR 2015, Workshop Track Proceedings*.
- Scott, C.; Blanchard, G.; and Handy, G. 2013. Classification with Asymmetric Label Noise: Consistency and Maximal Denoising. In *Proceedings of COLT 2013*, 489–511.
- Sukhbaatar, S.; Bruna, J.; Paluri, M.; Bourdev, L.; and Fergus, R. 2015. Training Convolutional Networks with Noisy Labels. In *Workshop Track of the International Conference on Learning Representations (ICLR)*.
- Tjong Kim Sang, E. F.; and De Meulder, F. 2003. Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition. In *Proceedings of HLT-NAACL 2003*, 142–147.
- Tkachenko, A.; Petmanson, T.; and Laur, S. 2013. Named Entity Recognition in Estonian. In *Proceedings of the 4th Biennial International Workshop on Balto-Slavic Natural Language Processing, BSNLP@ACL 2013, Sofia, Bulgaria, August 8-9, 2013*, 78–83.
- Torralba, A.; Fergus, R.; and Freeman, W. T. 2008. 80 Million Tiny Images: A Large Data Set for Nonparametric Object and Scene Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 30(11): 1958–1970.
- van Rooyen, B.; and Williamson, R. C. 2017. A Theory of Learning with Corrupted Labels. *J. Mach. Learn. Res.* 18: 228:1–228:50.
- Veit, A.; Alldrin, N.; Chechik, G.; Krasin, I.; Gupta, A.; and Belongie, S. J. 2017. Learning From Noisy Large-Scale Datasets With Minimal Supervision. In *Proceedings of CVPR 2017*, 6575–6583.
- Wang, H.; Liu, B.; Li, C.; Yang, Y.; and Li, T. 2019. Learning with Noisy Labels for Sentence-level Sentiment Classification. In *Proceedings of EMNLP-IJCNLP 2019*, 6285–6291.
- Xia, X.; Liu, T.; Han, B.; Wang, N.; Gong, M.; Liu, H.; Niu, G.; Tao, D.; and Sugiyama, M. 2020. Part-dependent Label Noise: Towards Instance-dependent Label Noise. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Proceedings of NeurIPS 2020*.
- Xia, X.; Liu, T.; Wang, N.; Han, B.; Gong, C.; Niu, G.; and Sugiyama, M. 2019. Are Anchor Points Really Indispensable in Label-Noise Learning? In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Proceedings of NeurIPS 2019*, 6835–6846.
- Xiao, T.; Xia, T.; Yang, Y.; Huang, C.; and Wang, X. 2015. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2691–2699.
- Yao, Y.; Liu, T.; Han, B.; Gong, M.; Deng, J.; Niu, G.; and Sugiyama, M. 2020. Dual T: Reducing Estimation Error for Transition Matrix in Label-noise Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Proceedings of NeurIPS 2020*.