
Active Learning for Noisy Data Streams Using Weak and Strong Labelers

Taraneh Younesian
Delft University of Technology
t.younesian@tudelft.nl

Dick Epema
Delft University of Technology
d.h.j.epema@tudelft.nl

Lydia Y. Chen
Delft University of Technology
y.chen-10@tudelft.nl

Abstract

Labeling data correctly is an expensive and challenging task in machine learning, especially for on-line data streams. Deep learning models especially require a large number of clean labeled data that is very difficult to acquire in real-world problems. Choosing useful data samples to label while minimizing the cost of labeling is crucial to maintain efficiency in the training process. When confronted with multiple labelers with different expertise and respective labeling costs, deciding which labeler to choose is nontrivial.

In this paper, we consider a novel weak and strong labeler problem inspired by humans' natural ability for labeling, in the presence of data streams with noisy labels and constrained by a limited budget. We propose an on-line active learning algorithm that consists of four steps: filtering, adding diversity, informative sample selection, and labeler selection. We aim to filter out the suspicious noisy samples and spend the budget on the diverse informative data using strong and weak labelers in a cost-effective manner. We derive a decision function that measures the information gain by combining the informativeness of individual samples and model confidence. We evaluate our proposed algorithm on the well-known image classification datasets CIFAR10 and CIFAR100 with up to 60% noise. Experiments show that by intelligently deciding which labeler to query, our algorithm maintains the same accuracy compared to the case of having only one of the labelers available while spending less of the budget.

1 Introduction

Obtaining a labeled dataset for training machine learning models is a time consuming and expensive task. The common practise of curating data continuously collects both data and the metadata that serves as labels Hendrycks et al. [2018] from the public domain. This immediately reduces the cost of acquiring a high volume of labeled data for deep learning models but introduces the challenge of handling data streams with noisy labels Han et al. [2018], Wang et al. [2018], i.e., data that is annotated with wrong class labels. Human experts Sheng et al. [2008] are sought after to correct labels to enhance the robustness of learning models against label noise. Cognitive science Yan et al. [2014] has shown that humans are better in answering binary questions, such as True/False questions, and are less skilled in multiple-choice questions, e.g., identifying one out of 100 classes in CIFAR-100 benchmark Krizhevsky et al. [2009]. It is more expensive to use strong labelers who are skilled in directly pointing out true class labels than weak labelers who can only (dis)agree with the provided labels. To cost effectively correct the noisy labels by experts Sheng et al. [2008], Chang et al. [2017], it is imperative to assign the correct tasks to the labelers according to their skills and the

difficulty levels of querying tasks. This becomes particularly challenging in streaming data scenarios due to future uncertainty, e.g., how to allocate the correction effort across the learning horizon.

Online learning from streaming data corresponds to today’s data common practise to train machine learning models with a small set of data that is periodically curated from the public domains. Machine learning models thus need to be learned online from the stream data. Moreover, due to privacy or regulation constraints gdp, or storage limits, data turns are available for the limited duration. Moreover, in such scenarios, combating label noise adds to the challenge.

The prior art tries to enhance the robustness of the deep model training against noisy labels in the off-line scenario by (i) filtering out noisy data through model disagreement Han et al. [2018], Yu et al. [2019], (ii) correcting noisy labels through the estimated noisy corruption matrix Patrini et al. [2017], or (iii) modifying the loss functions Ma et al. [2018], Wang et al. [2019]. Among the related studies, high quality labels from human experts are used to correct labels through the estimate noise corruption matrix, e.g., the trust data set in Distillation Li et al. [2017] and GLC Hendrycks et al. [2018]. The trust data sheds light on how a small percentage of high quality labels can prevent deep models from accuracy degradation due to label noise. However, the focus on the off-line datasets renders existing approaches insufficient for handling noisy data streams only a subset of which can be learnt at a time. Indeed, on-line learning from noisy labels can lead to a much more severe degradation of accuracy than the typical off-line case. Moreover, such a subset of trust data is randomly selected and the cost of acquiring additional data label yet to be modeled.

In this paper we address the challenge of training a deep classifier from noisy labeled data streams that are collected over time and that can only be learnt for a limited time. We propose an active learning framework, DuoLab, which aims to learn a robust classifier by cost-effectively assigning the labelcorrection task to either a strong or a weak labeler within a labeling budget. Extensively querying the strong labeler can easily lead to budget exhaustion, whereas the weak labeler might require multiple queries to achieve the cleansing goal. DuoLab consist of four steps: filtering out the suspicious data, choosing diverse samples, informativeness ranking, and selecting labelers. While training the classifier, DuoLab leverages the output of its classifier to filter suspicious data samples. As for label cleansing, we propose a labeler selection function Q that combines the overall model confidence and the informativeness of individual samples.

Our contributions are the following. First we design a cost and skill aware active learning framework for noisy data streams. Secondly, by leveraging the diversity of the labelers we are able to greatly enhance the robustness of deep models even in challenging on-line learning scenarios. The proposed Q function can effectively assess the model confidence and informativeness of data samples and assign the suitable labelers accordingly.

2 Related Work

We summarize the related work in the context of robust learning, active learning, and streaming data—the main themes discussed in this paper.

Robust learning against noise labels. Training a robust deep classifier against noisy labels is an active research field Patrini et al. [2017], Ma et al. [2018], Reed et al. [2015], Han et al. [2018]. The main focus is on the off-line scenario where the training data is available at once and can be learned unlimited. To build robust classifier, the prior art either filters out the suspicious noisy data by the intermediate results of deep model Han et al. [2018], Yu et al. [2019] or derives noise resilient loss function Ma et al. [2018], e.g., symmetric cross entropy Wang et al. [2019]. Konstantinov et al. Konstantinov and Lampert [2019] considers multiple entrusted sources and assign weights to the sources according to their quality quantified by the difference between the source and target distributions.

Active learning Active learning Settles [2009] is typically drawn to leverage the expert knowledge to provide the label information for the useful unlabeled data samples that might be limited in the real world data set Sourati et al. [2018], Yan et al. [2018]. Having multiple expert, as the experts may vary in their experience, it is important to match the expertise of labelers to the labeling tasks. While Yan et al. [2016] considers imperfect labelers that may abstain from labeling, the study in Huang et al. [2017] assumes having multiple labelers with different costs and qualities and actively selects both samples and the labelers considering sample usefulness and labeler’s accuracy and cost, assuming

that all the labelers are prone to make mistakes. The study Sheng et al. [2008] focuses on the selection of informative samples in the presence of several non-expert labelers via majority voting of the of the labelers. Considering the same framework, an extension of unbalanced labels is studied in Zhang et al. [2015]. However, they fail to leverage the labelers based on their expertise in labeling. In contrast, Tanno et al. [2019] considers several noisy annotators with unknown expertise and jointly estimates the confusion matrix of the annotators and the true label distribution. Furthermore, the study Zhang and Chaudhuri [2015] designs a difference classifier that predicts where the non-expert labeler differs from the expert in a binary classification task. Moreover, a line of research focuses on multiple uncertain experts' view in clustering Chen et al. [2018], Chang et al. [2017], however, they don't benefit from active learning for informative sample identification.

Streaming Data As an emerging challenge in real-world scenarios, online deep learning is largely overlooked by the prior art. Some studies have attempted to overcome the convergence problem of online deep learning using the idea of sliding window Zhou et al. [2012] and Lee et al. [2016]. the authors in Sahoo et al. [2018] had a different approach to the problem by dynamically changing the depth of the network from shallow to deep to adapt the model capacity. Noisy data streams have been studied by Chu et al. [2004] and Zhu et al. [2006] using an ensemble of classifiers to build a robust model that maximizes the data likelihood. However, the focus of this research was on traditional machine learning models and is not scalable to deep networks.

3 Dealing with Weak and Strong Labelers

In this paper we consider the following on-line learning scenario, which is illustrated in Figure 1. The data periodically, at consecutive time steps $t, t + 1, \dots$, streams into the classifier \mathcal{C} in small batches D for training. Each arriving data batch is used for training the classifier for a limited duration. Here, we assume the most stringent case in which each batch of data can only be used for training the classifier until the arrival of the next batch.

Each data instance $(\mathbf{x}_j, \tilde{y}_j)$ contains feature inputs $\mathbf{x}_j \in \mathcal{X} \subset \mathbb{R}^d$ and a potentially noisy label $\tilde{y} \in \mathcal{Y} := \{1, \dots, N\}$. Due to annotation errors or even malicious attacks Patrini et al. [2017], the collected label (\tilde{y}) can be corrupted from the original label (y). The classifier can either learn directly from the whole batch of dataset or it can first selectively correct labels and then learn from the combined. clean and cleansed data. We will now go through the components of our method DuoLab in Figure 1; filtering, clustering and the Q function are explained in Section 4.

There are two types of labelers, namely a weak and a strong one, available to correct the noisy labels subject to a budget. The weak labeler \mathcal{W} is not an expert in labeling and instead of assigning a label to each unlabeled data instance, it can answer questions about its label with "yes/no". When presented with the true label, \mathcal{W} gives the answer "yes", and "no" otherwise. The strong labeler \mathcal{S} on the other hand, knows the true label of the data; however, it is very costly to use. We assume \mathcal{S} to be c times as expensive as \mathcal{W} to label each data point. Therefore, c times asking \mathcal{W} with "yes/no" questions is equivalent to asking \mathcal{S} for the true label once.

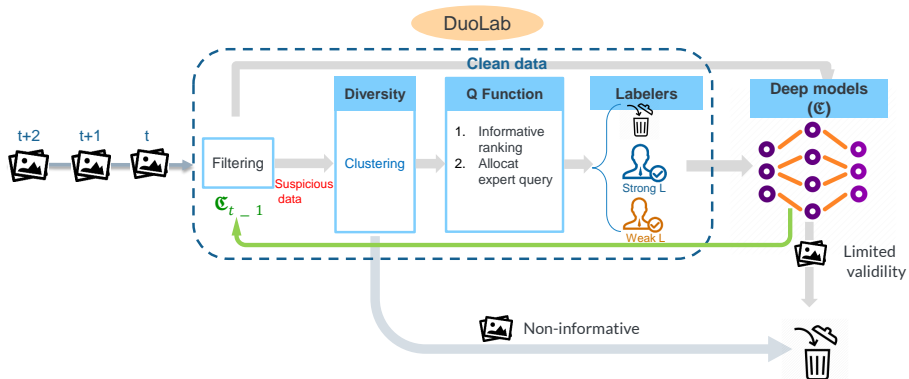


Figure 1: On-line learning scenario with noisy labels: data streams, classifier, and strong/weak labelers

Ultimately, the objective is to train a classifier in the presence of noisy labels considering these two labelers with a limited budget B per batch for labeling to achieve a certain level of accuracy. Detecting noisy samples and deciding which data to ask from which labeler is the main challenge of this problem. Since there exists a limited resource, i.e., the labeling budget, it is crucial to spend it on cleansing the representative and useful noisy samples. Therefore, the identification of the informative samples is one of the goals to solve this problem. Moreover, another challenge is, when relying on the weak labeler, how many "yes/no" questions to ask, and if requiring the strong labeler, how to allocate its budget for the arriving batch of data over time. To summarize, we aim to address the following questions in this paper: (i) how to detect noisy samples, (ii) how to identify informative data samples, and (iii) how to allocate the cleansing budget across the different labelers.

4 DuoLab

We propose an algorithm called DuoLab, which detects the noisy labeled data and decides how to find their true label within the budget and to train a deep multi-class classifier over time. To cost-effectively train an accurate classifier, DuoLab leverages the information provided by the labelers as well as the model’s prediction. DuoLab consists of the four stages of *filtering*, *clustering*, *ranking* (informative sample selection), and *labeler selection*. We explain each of these steps below. Here, we specifically consider convolution neural networks (CNN) as the underlying classifier. The general design of DuoLab can be applied on different types of classifier.

4.1 Filtering: Identifying the Suspicious Data

The first step of DuoLab is to identify the suspicious data samples that might have corrupted labels. We aim to filter out the noisy samples by leveraging the model’s prediction and the confidence. We assume that the CNN classifier is initially trained with a small clean labeled set D^I . After the initial training, data samples $\{(\mathbf{x}_j, \tilde{y}_j)\}$ arriving in batches are provided to the model to predict each sample’s label. To filter out the noisy samples, for each sample \mathbf{x}_j we compare the highest and second-highest predicted label¹ \hat{y}_j^1 and \hat{y}_j^2 with the given label \tilde{y}_j , defining the sequence of n highest predicted labels for \mathbf{x}_j by \mathcal{C} as $\Omega = \{\hat{y}_j^n\}_{n=1}^N$. We consider the samples that have $\hat{y}_j^1 = \tilde{y}_j$ or $\hat{y}_j^2 = \tilde{y}_j$ to be clean, and add them to the clean set $C = \{(\mathbf{x}_j^c, y_j)\}$. The other samples are suspected to be noisy, and are collected in the suspicious set $U = \{(\mathbf{x}_j^u, \tilde{y}_j)\}$. Relaxing the filtering criteria to include lower-ranked predicted labels, e.g., the third and fourth-highest predicted labels, may lead to more false positives, i.e., noisy samples that are identified as clean ones. Contrarily, more stringent filtering criteria, e.g., only using the highest predicted label, results into more false negatives.

4.2 Clustering: Adding Diversity

Since cleaning all the noisy samples is very expensive, our goal is to select samples that properly represent the dataset and are highly informative. It has been shown Wang et al. [2017] that only relying on the informative sample selection may result in selecting similar samples that would cause a waste of the labeling budget. To add diversity to the selected samples, we apply kmeans clustering on the extracted features of the suspicious set U . The features are extracted from the network’s last layer’s output before the softmax. Next, we select the ℓ most informative samples of each cluster using active learning as explained in the next section, and discard the rest.

4.3 Ranking: Informativeness

When dealing with a large set of data that requires (re)labeling with a limited budget, the key is to identify the data samples the model could gain from the most when being trained on them. Active learning aims to find the most informative data and asks their true labels from the labelers. We rank the samples based on their informativeness, select the top ℓ ones from each cluster, and put them in the set $K = \{(\mathbf{x}_j^k, \tilde{y}_j)\}$. The samples in K are then processed in decreasing order of their informativeness index i , where $i = 1$ is the index of the most informative sample in K .

To determine the informativeness of a sample, we estimate the prediction uncertainty of the network via the prediction probability $\mathbf{p}(\mathbf{x}_j)$ at the output of the softmax layer. We employ a well-known

¹The ranks of the predicted labels are determined by the value of softmax output of CNN.

active learning method called **Best-versus-second-best (BvSB)** Joshi et al. [2009] to measure the uncertainty of the model while classifying a sample. This method considers samples for which the difference between the probabilities of the two most likely classes is the smallest, meaning that these samples could have easily been classified either way. That is, using p_{best} and $p_{second-best}$ to denote the probabilities of the most likely and second most likely class, the smaller the value of $I(\mathbf{x}_j) = p_{best}(\mathbf{x}_j) - p_{second-best}(\mathbf{x}_j)$, the higher the informativeness of \mathbf{x}_j .

4.4 Labeler Selection: Cost Optimization

The next step is to decide which labeler to choose considering their cost to cleanse the noisy labels. We introduce a cost sensitive labeler selection function Q that solicits the appropriate labeler for each data sample in K . We consider three criteria to be crucial in the process of labeler selection: *i*) the informativeness of the sample, *ii*) the cost of the labelers, and *iii*) the reliability of the model. Our aim is to spend more budget on the informative samples, and therefore we set the value of the selection function to be higher if a sample is more informative, and to guarantee cleansing those samples we aim to assign the most informative samples to the strong labeler. As the second criterion is the cost of the labelers, the challenge is to decide how to allocate the limited budget between the strong and weak labelers. We define the selection function to be lower if the cost is high, resulting in higher chance of the selection of the weak labeler in costly situations. Moreover, the higher the value of the function, the higher the chance to ask the strong labeler for labeling. Furthermore, if the model is not performing well enough, it means we need more cleansing so the value of the function should increase.

Cost Sensitive Labeler Selection Function. We will now define the cost sensitive labeler selection function Q and show how it takes into account the three labeler-selection criteria mentioned above. Before we do so, we define the following measure $L_V(t)$ for the overall model reliability up to the current batch t using the cross-entropy loss on the clean validation set D^V :

$$L_V(t) = \sum_{v=1}^V p(\mathbf{x}_v, y) \log(p(\mathbf{x}_v, \hat{y}^1)). \quad (1)$$

A lower value of $L_V(t)$ indicates a higher reliability. In addition, we denote by $E_{\mathcal{W}}(t)$ and $E_{\mathcal{S}}(t)$ the numbers of queries from the current batch t that have been asked sofar from the two labelers. The function Q is only applied when the current total cost does not exceed the budget, that is, when $E_{\mathcal{W}}(t) + cE_{\mathcal{S}}(t) \leq B$. The value of the labeler selection function Q on a data sample \mathbf{x}_j^u is now defined as

$$Q(\mathbf{x}_j^u) = \frac{L_V(t)}{I(\mathbf{x}_j^u)cE_{\mathcal{S}}(t)}, \quad (2)$$

which has a higher value when the model is less reliable, when the sample is more informative, and when the strong labeler \mathcal{S} has not yet been queried very often—the more \mathcal{S} is queried during a batch, the more difficult it becomes to query it for later samples.

Algorithm 1 shows the steps of our proposed method. As indicated in the algorithm, for each data arrival, after detecting the noisy samples we calculate their values of the cost sensitive selection function Q . If $Q(\mathbf{x}_j^u) > \bar{Q}$ for some threshold \bar{Q} , \mathbf{x}_j^u is presented to \mathcal{S} to be labeled and added to the clean set C , otherwise we query \mathcal{W} . If the answer is "No", we repeat the steps until we get the clean label or we exceed the budget. In the latter case, we discard the sample.

A key challenge about the weak labeler is how to proceed upon receiving "No" for a sample. A naive approach is randomly choosing from the remaining labels and continuing until receiving "Yes". Here, we introduce a more intelligent way by asking \mathcal{W} for the label of each sample in the order of Ω , which is the sorted sequence of the labels based on the classifier's prediction probabilities (see Section 4.1); this presumably results in fewer \mathcal{W} queries. Furthermore, we introduce a parameter \bar{w} indicating the maximum number of queries to the weak labeler per sample, which prevents wasting too much of the budget on a single sample and allows for more exploration.

5 Experiments

We evaluate DuoLab on CIFAR-10 and CIFAR-100 whose data instances are corrupted with label noise and are on-line streamed to the classifier. We compare DuoLab with the state of the art in

Algorithm 1: Algorithm of DuoLab: filtering, clustering, labeler selection, and training

Input : Initial dataset D^I , Data batches D , weak labeler \mathcal{W} , strong labeler \mathcal{S} , cost of strong labeler c , maximum number of weak queries per sample \bar{w} , budget B , clustering parameter k

Output : Training set C for the classifier \mathcal{C}

```
1 Train  $\mathcal{C}$  with  $D^I$ 
2 foreach arriving  $D = \{(\mathbf{x}_j, \tilde{y}_j)\}$  and each  $(\mathbf{x}_j, \tilde{y}_j)$  do
3   if  $(\tilde{y}_j^1 = \tilde{y}_j)$  or  $(\tilde{y}_j^2 = \tilde{y}_j)$  then
4      $C = \{(\mathbf{x}_j^c, y_j), \mathbf{x}_j^c := \mathbf{x}_j\}$  #clean set
5   else
6      $U = \{(\mathbf{x}_j^u, \tilde{y}_j), \mathbf{x}_j^u := \mathbf{x}_j\}$  #suspicious set
7   Apply Kmeans clustering on  $U$ .
8   From each cluster select the most informative  $k$  samples and add to  $K = \{(\mathbf{x}_j^k, \tilde{y}_j)\}$ . #Discard the rest
9   for  $\mathbf{x}_i$  in  $K$  where  $i$  is the informativeness index do
10     $w = 0$ 
11    if  $(Q > \bar{Q})$  and  $(E_W(t) + cE_S(t) + c \leq B)$  then
12      Query  $\mathcal{S}$ , update  $E_S$  and  $Q$ 
13      Add  $\mathbf{x}_i$  to  $C$ 
14    else if  $(E_W(t) + cE_S(t) + 1 \leq B)$  and  $(w < \bar{w})$  then
15      Query  $\mathcal{W}$  based on  $\Omega$ , update  $E_W$  and  $Q$ ,  $w = w + 1$ 
16      if (The answer is "Yes") then
17        Add  $\mathbf{x}_i$  to  $C$ 
18      else
19        Go to step 14
20    else
21      Discard  $\mathbf{x}_i$ 
```

different, varying learning scenarios, i.e., availability of labelers, cost, and the size of the on-line batch. We demonstrate how DuoLab improves the accuracy of the classifier using filtering and cleansing by the weak and strong labelers within the limited budget.

5.1 Experimental Setup

Training Network’s Characteristics. As DuoLab classifier for CIFAR-10 and CIFAR-100, we use two Convolutional Neural Network (CNN) architectures defined in Ma et al. [2018] with ReLU activation functions, softmax activation as image classifier and cross-entropy as loss function. We train the models by using stochastic gradient descent with momentum 0.9, learning rate 0.1, and weight decay 2×10^{-3} . DuoLab and all baselines are implemented using Keras v2.2.4 and Tensorflow v1.12, except co-teaching which uses PyTorch v1.1.0.

Label Noise Injection To inject noise into the labels, we select samples with the probability equal to the noise rate and replace the true label with a random different label with uniform probability. We use noise rates of 30% and 60% in the training set. The test and validation sets are clean.

Noise Resilient Baselines We put DuoLab in the context of other noise-resistant techniques drawn from the related work and adapted to the online scenario: **D2L** Wang et al. [2018] estimates the dimensionality of subspaces during training to adapt the loss function, **Forward** Patrini et al. [2017] corrects the loss function based on the noise transition matrix, **Bootstrap** Reed et al. [2015] (hard and soft) uses convex combination of the given and predicted labels for training, and **Co-teaching** Han et al. [2018] exchanges mini-batches between two networks trained in parallel.

Training Parameters for CIFAR-10 and CIFAR-100. For CIFAR-10, the initial training set consists of 4K samples, and 1K samples are considered as the validation set to measure model reliability. The remaining 55K samples are divided into 45K training and 10K testing set. The training data arrive in 45 batches of 1K samples with 30% and 60% noise in an online setting. The initial set is trained for 60 epochs and the rest are trained 20 epochs per batch. The epochs are selected in a way to let the network train enough but also avoid overfitting due to small batch size. We set the cost sensitive function threshold Q to 10 and the cost of the strong labeler to $c = \{2, 10\}$.

Table 1: The accuracy, the average numbers of queries per labeler per batch, and the TP and FP (true and false positive) rates for DuoLab and different labeler selection baselines with budget $B = 250$ and $\bar{w} = 2$ for CIFAR-10 ("-" means not applicable).

Noise	30%						60%				
Method	c	Acc(%)	no. S	no. W	TP(%)	FP(%)	Acc(%)	no. S	no. W	TP(%)	FP(%)
DuoLab	2	76.13	22.0	57.0	61.43	3.81	69.44	38.4	63.0	33.57	7.84
DuoLab	10	75.45	4.7	52.5	60.90	3.80	67.42	9.8	51.8	32.95	7.71
DuoLab + Kmeans	2	75.99	15.3	67.1	61.26	3.66	68.61	21.5	78.7	33.38	7.84
DuoLab + Kmeans	10	75.33	4.1	56.1	60.70	3.78	67.53	21.5	78.7	33.38	7.84
Only \mathcal{S}	10	74.96	25.0	-	60.80	3.78	66.84	25.0	-	32.84	7.81
Only \mathcal{W}	-	75.11	-	67.6	60.87	3.84	68.93	-	89.5	33.53	7.75
Clean All Suspicious	-	77.26	343.2	-	61.72	3.91	75.71	-	575.1	35.03	7.46
No AL(only Filter)	-	73.60	-	-	60.77	3.78	63.34	-	-	32.26	8.11
Opt Filter	-	77.78	-	-	70.00	-	72.56	-	-	40.00	-
No Filter	10	62.21	7.3	96.7	-	-	36.91	17.7	24.2	-	-

For CIFAR-100 we increase the arriving data batch size to 9K and 60 epochs per batch to cope with the higher complexity. The cost sensitive function threshold \bar{Q} was set to 150 and the cost of the strong labeler to $c = \{5, 25\}$. In both datasets, the threshold is determined experimentally. In general, a high threshold, e.g., $\bar{Q} = 150$, leads to using less the strong labeler and more the weak labeler.

For both CIFAR-10 and CIFAR-100, we set the labeler budget of DuoLab to be 25% of the batch size, which amounts to 250 and 2250 unit cost per batch for each dataset, respectively. The higher budget leads to more cleansing and therefore, higher accuracy. Whenever the accuracy of the network over the validation sets drops for $r\%$ from the start to the end of a batch, we rollback the model weights to the previous batch's parameters before processing the following one. Rollback uses $r = 20\%$. We report the highest accuracy among batches for each experiment. All the experiments are repeated three times and the average is reported in the tables. We show the highest accuracy among the first six rows of the tables 1 and 2 in bold.

5.2 CIFAR-10 Results

The effect of filtering. Here we compare our results with three baselines: *Opt Filter*, where there is an optimal filter that can perfectly select all the clean samples and discard all the noisy ones, *No Filter*, where there is no filtering, and *No AL*, where we filter out samples using our filtering method but no further cleansing is done on the suspicious samples by active learning. As the results show in Table 1, *Opt Filter* has a high accuracy due to omitting all the noisy samples and training over a relatively large clean set (700 and 400 clean samples for 30% and 60% noise). However, despite having a smaller clean data and some noisy samples due to false positive (FP) detection, our method is able to get close in terms of accuracy. Comparing to the *No AL* results, one can observe the effectiveness of our filtering together with active cleansing.

The effect of clustering. As Table 1 shows, clustering results in cleaning more suspicious samples compared to DuoLab without clustering; however, since it fails to cleanse fewer highly informative samples by \mathcal{S} , it results in lower accuracy.

The effect of the labeler selection. We compare DuoLab with two extreme labeler selection baselines: i) *Only \mathcal{S}* , where only the strong labeler is available to be queried, and ii) *Only \mathcal{W}* , where only the weak labeler is available to be queried. Moreover, to show the upper and the lower bound of the accuracy, we present the accuracy of the case where all the suspicious samples are cleansed and also when no cleansing is done, namely *No AL*. As the results show in Table 1, with lower cost for the strong labeler, i.e., with $c = 2$, it is more beneficial to use only the strong labeler since it leads to cleaning more suspicious samples. The reason is that due to \bar{w} , the weak labeler fails to give the true label if the true label is not among the top predictions of the classifier. However, with higher costs, using only \mathcal{S} results in cleaning a very small number of samples. Contrarily, although using only the weak labeler might result in cleaning more data, missing some highly informative samples due to the limitation caused by \bar{w} results in a low performance. Our proposed method outperforms these two cases by benefiting from both cleansing the highly informative samples with the strong labeler, and minimizing the cost with the weak labeler.

Table 2: The accuracy, the average numbers of queries per labeler per batch, and the TP and FP (true and false positive) rates for DuoLab and different labeler selection baselines with budget $B = 2250$ and $\bar{w} = 2$ for CIFAR-100 ("- means not applicable).

Noise	30%						60%				
Method	c	Acc(%)	no. S	no. W	TP(%)	FP(%)	Acc(%)	no. S	no. W	TP(%)	FP(%)
DuoLab	5	39.07	33.3	102.1	32.18	0.44	34.45	44.0	171.9	17.21	0.97
DuoLab	25	39.27	7.3	102.7	32.57	0.45	34.14	8.9	174.2	17.65	0.90
DuoLab + Kmeans	5	38.45	16.9	105.1	33.14	0.49	34.35	17.3	190.8	17.65	0.97
DuoLab + Kmeans	25	38.98	4.0	100.5	32.30	0.49	32.92	4.0	181.4	17.11	0.97
Only \mathcal{S}	25	39.57	90.0	-	33.14	0.47	33.13	90.0	-	17.12	0.95
Only \mathcal{W}	-	38.25	-	105.9	32.83	0.48	32.94	-	192.2	17.12	0.99
Clean All Suspicious	-	49.05	5085.4	-	40.91	0.45	49.53	6829.6	-	23.38	0.74
No AL(only Filter)	-	38.08	-	-	32.63	0.44	32.32	-	-	16.18	0.98
Opt Filter	-	42.28	-	-	70.00	-	37.33	-	-	40.00	-
No Filter	25	30.86	112.0	-	-	-	13.19	112.0	-	-	-

5.3 CIFAR-100 Results

The effect of filtering. As the results in Table 2 show, considering that CIFAR-100 has more classes than CIFAR-10 and less data per class, filtering is more difficult. However, we succeed in having a very low FP rate by our filtering method, and therefore we achieve an accuracy that is very close to the upper-bound of filtering *Opt Filter*. Comparing the accuracy of *No AL* and DuoLab with *Opt Filter*, one could observe that filtering has a high impact on the performance, especially when the labelers' cleansed sample set is small (less than 2% of the batch size), due to a low budget and more complicated dataset. Moreover, since filtering relies on the model's prediction, the higher accuracy of the model leads to a better filtering, i.e., a higher TP and a lower FP rate, which again impacts the accuracy itself.

The effect of clustering. Similar to CIFAR-10, clustering results in a slightly lower accuracy, since it affects the number of cleansed samples by the labeler. As shown in Table 2, applying clustering on the suspicious data causes less cleansing by the strong labeler, which is equivalent to missing more informative samples.

The effect of the labeler selection. As Table 2 shows, DuoLab performs better than *Only S* and *Only W* in the case of high strong query cost. In this case, DuoLab succeeds in cleansing more samples than *Only S* and therefore achieves a similar or higher accuracy. This difference is more apparent with 60% noise. Moreover, DuoLab achieves a similar cleansed number compared to *Only W* and due to benefiting from \mathcal{S} beside \mathcal{W} , DuoLab guarantees the cleansing of the highly informative samples and thus results in a more accurate model. The higher the cost is, the more effective our method performs.

5.4 Comparison with the Noise Resilient Models

We compare our proposed DuoLab with the baselines described in Section 5.1 under 30% label noise. Table 3 summarizes the results for different noise-resistant model baselines and DuoLab. Analyzing the results in the table shows that although these methods are well-known to combat the label noise, they achieve this by benefiting from a large dataset available for training, and therefore they fall short to adapt to the streaming data setting. On the contrary, DuoLab, benefiting from a novel filtering method and the knowledge of the experts, excellently adapts to the streaming data scenario with small batch size and for the limited learning duration.

Table 3: The accuracy of DuoLab and the noise resilient baselines for CIFAR-10 and CIFAR-100 with 30% noise.

Method	Accuracy (%)	
	CIFAR-10	CIFAR-100
Baselines		
D2L	52.77	11.55
Forward	59.94	25.56
Co-teaching	61.52	29.41
Bootstrap soft	48.95	23.35
Bootstrap hard	49.61	24.02
Our		
DuoLab (c = 2, 5)	76.13	39.07
DuoLab (c = 10, 25)	75.45	39.27

6 Conclusion

In this paper we introduce DuoLab a method to overcome the challenge of training deep neural networks with streaming noisy labeled data by weak and strong labelers. As the strong labeler comes with its cost of labeling and is constrained by a budget, our method decides when to use each labeler to achieve a high level of performance beside cleaning as more samples as possible. As a result of DuoLab benefiting from the merits of both labelers in cleaning the highly informative samples and efficiently spending the budget, we are able to perform better than using only one of the labelers. We achieve this by only cleansing up to 10% of the batch size in CIFAR-10 and 2% in CIFAR-100.

Broader Impact

This core idea of fusing human and artificial intelligence in DuoLab can be broadly applicable to data management and knowledge extraction. The cost- and privacy-awareness in DuoLab addresses the gap between state of the practise, i.e., streaming learning on noisy data, and the state of the art on off-line learning on the clean data. For instance, the rich information embedded in social media or recommendation systems is unfortunately noisy or even fake and can not be easily extracted through the off-line resilient solutions. The proposed DuoLab has a great potential to distill the data and improve the knowledge extraction for the society. On the other hand, if DuoLab fails to defend the training process against noise, the classifier can end up being inaccurate and may even harm the society in case of misuse. The other downside of utilizing DuoLab is that we might indirectly foster the noisy data generation as the cost-efficiency of cleansing data via the labelers improves.

References

- Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. In *NIPS*, pages 10456–10465, 2018.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor W. Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, pages 8536–8546, 2018.
- Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *IEEE CVPR*, pages 8688–8696, 2018.
- Victor S. Sheng, Foster J. Provost, and Panagiotis G. Ipeirotis. Get another label? improving data quality and data mining using multiple, noisy labelers. In *ACM SIGKDD*, pages 614–622, 2008.
- Yan Yan, Rómer Rosales, Glenn Fung, Subramanian Ramanathan, and Jennifer G. Dy. Learning from multiple annotators with varying expertise. *Mach. Learn.*, 95(3):291–327, 2014.
- Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). 2009. URL <http://www.cs.toronto.edu/~kriz/cifar.html>.
- Yale Chang, Junxiang Chen, Michael H. Cho, Peter J. Castaldi, Edwin K. Silverman, and Jennifer G. Dy. Multiple clustering views from multiple uncertain experts. In *ICML*, pages 674–683, 2017.
- General data protection regulation. URL <https://gdpr.eu/>.
- Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, pages 7164–7173, 2019.
- Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *IEEE CVPR*, pages 1944–1952, 2017.
- Xingjun Ma, Yisen Wang, Michael E. Houle, Shuo Zhou, Sarah M. Erfani, Shu-Tao Xia, Sudanthi N. R. Wijewickrema, and James Bailey. Dimensionality-driven learning with noisy labels. In *ICML*, volume 80, pages 3361–3370, 2018.

- Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *IEEE ICCV*, pages 322–330, 2019.
- Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. In *IEEE International Conference on Computer Vision ICCV*, pages 1928–1936, 2017.
- Scott E. Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. In *ICLR*, 2015.
- Nikola Konstantinov and Christoph Lampert. Robust learning from untrusted sources. In *ICML*, volume 97, pages 3488–3498, 2019.
- Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- Jamshid Sourati, Murat Akçakaya, Deniz Erdogmus, Todd K. Leen, and Jennifer G. Dy. A probabilistic active learning algorithm based on fisher information ratio. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(8):2023–2029, 2018.
- Songbai Yan, Kamalika Chaudhuri, and Tara Javidi. Active learning with logged data. In *ICML*, pages 5517–5526, 2018.
- Songbai Yan, Kamalika Chaudhuri, and Tara Javidi. Active learning from imperfect labelers. In *NeurIPs*, pages 2128–2136, 2016.
- Sheng-Jun Huang, Jia-Lve Chen, Xin Mu, and Zhi-Hua Zhou. Cost-effective active learning from diverse labelers. In Carles Sierra, editor, *IJCAI*, pages 1879–1885. ijcai.org, 2017.
- Jing Zhang, Xindong Wu, and Victor S. Sheng. Active learning with imbalanced multiple noisy labeling. *IEEE Trans. Cybernetics*, 45(5):1081–1093, 2015.
- Ryutaro Tanno, Ardavan Saeedi, Swami Sankaranarayanan, Daniel C. Alexander, and Nathan Silberman. Learning from noisy labels by regularized estimation of annotator confusion. In *CVPR*, pages 11244–11253, 2019.
- Chicheng Zhang and Kamalika Chaudhuri. Active learning from weak and strong labelers. In *NeurIPs*, pages 703–711, 2015.
- Junxiang Chen, Yale Chang, Peter J. Castaldi, Michael H. Cho, Brian Hobbs, and Jennifer G. Dy. Crowdclustering with partition labels. In *AISTATS*, pages 1127–1136, 2018.
- Guanyu Zhou, Kihyuk Sohn, and Honglak Lee. Online incremental feature learning with denoising autoencoders. In *AISTATS*, pages 1453–1461, 2012.
- Sang-Woo Lee, Chung-yeon Lee, Dong-Hyun Kwak, Jiwon Kim, Jeonghee Kim, and Byoung-Tak Zhang. Dual-memory deep learning architectures for lifelong learning of everyday human behaviors. In *IJCAI*, pages 1669–1675, 2016.
- Doyen Sahoo, Quang Pham, Jing Lu, and Steven C. H. Hoi. Online deep learning: Learning deep neural networks on the fly. In *IJCAI*, pages 2660–2666, 2018.
- Fang Chu, Yizhou Wang, and Carlo Zaniolo. An adaptive learning approach for noisy data streams. In *IEEE ICDM*, pages 351–354, 2004.
- Xingquan Zhu, Xindong Wu, and Ying Yang. Effective classification of noisy data streams with attribute-oriented dynamic classifier selection. *Knowledge and Information Systems*, 9(3):339–363, 2006.
- Ran Wang, Xi-Zhao Wang, Sam Kwong, and Chen Xu. Incorporating diversity and informativeness in multiple-instance active learning. *IEEE Trans. Fuzzy Syst.*, 25(6):1460–1475, 2017.
- Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *IEEE CVPR*, pages 2372–2379, 2009.