



Elucidating robust learning with uncertainty-aware corruption pattern estimation

Jeongeun Park^a, Seungyoun Shin^b, Sangheum Hwang^c, Sungjoon Choi^{a,*}

^a Department of Artificial Intelligence, Korea University, Seoul 02841, Korea

^b Department Computer Engineering, Dongguk University, Seoul 04620, Korea

^c Department of Data Science, Seoul National University of Science and Technology, Seoul 01811, Korea

ARTICLE INFO

Article history:

Received 24 October 2021

Revised 29 August 2022

Accepted 1 February 2023

Available online 3 February 2023

Keywords:

Robust learning

Training with noisy labels

Uncertainty estimation

Corruption pattern estimation

ABSTRACT

Robust learning methods aim to learn a clean target distribution from noisy and corrupted training data where a specific corruption pattern is often assumed a priori. Our proposed method can not only successfully learn the clean target distribution from a dirty dataset but also can estimate the underlying noise pattern. To this end, we leverage a mixture-of-experts model that can distinguish two different types of predictive uncertainty, aleatoric and epistemic uncertainty. We show that the ability to estimate the uncertainty plays a significant role in elucidating the corruption patterns as these two objectives are tightly intertwined. We also present a novel validation scheme for evaluating the performance of the corruption pattern estimation. Our proposed method is extensively assessed in terms of both robustness and corruption pattern estimation in the computer vision domain. Code has been made publicly available at <https://github.com/jeongeun980906/Uncertainty-Aware-Robust-Learning>.

© 2023 The Authors. Published by Elsevier Ltd.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

In this paper, we focus on the problem of robust learning [5–9,37–39] with its emphasis on elucidating the corruption patterns on the noisy training dataset. Most existing robust learning studies [6,9] assume that the label corruption pattern is solely a function of class information, also known as Class-Conditional Noise (CCN). While this CCN assumption is simple to formulate, it may not be useful in practice in that it is more natural to assume that the noise pattern is a function of input instances which is often referred to as an Instance-Dependent Noise (IDN) learning problem [10].

However, the original IDN learning problem is likely to be infeasible in that it has to estimate a $C \times C$ class transition matrix per input instance. Due to this intractability, recent work on the IDN setting focuses on a simple binary classification problem [11,12] or requires a small clean dataset [10]. To mitigate this issue, we cast the IDN problem into a two-stage problem of first partitioning the input space using uncertainty measures and then estimating the label transition matrix per each group, which will be referred to as a Set-Dependent Noise (SDN) learning problem. In particular,

a specific type of predictive uncertainty, named aleatoric uncertainty, is used to partition the input space. The clusters with high aleatoric uncertainty can be viewed as collective outliers [13], a subset of inputs with a specific noise pattern.

We would like to stress that robust learning and uncertainty estimation problems are intimately related to each other as robust learning deals with the noisy training data, which naturally gives rise to predictive uncertainty. The predictive uncertainty can be decomposed into epistemic and aleatoric uncertainty. The former focuses on the reducible part of the uncertainty (model uncertainty), which may come from the lack of training data. In contrast, the latter comes from the irreducible part (data uncertainty), such as the measurement noise. Our proposed method can estimate both types of uncertainty in a unified framework, and it plays a significant role in achieving robustness and estimating the SDN patterns.

To this end, we utilize a mixture-of-experts model for classification tasks named mixture logit networks (MLN) and present an effective training method to achieve both robustness and explainability by revealing the label corruption process. Although a mixture-of-experts method was first presented in the 80s, we show its effectiveness on a robust learning framework with simple modifications and show that it can also estimate the noise distributions as well. We first present an uncertainty estimation method for the MLN that can distinguish two different types of predictive uncertainties, epistemic (model) uncertainty and aleatoric (data)

* Corresponding author.

E-mail addresses: baro0906@korea.ac.kr (J. Park), sungjoon-choi@korea.ac.kr (S. Choi).

uncertainty. Then, the estimated uncertainty is utilized for the uncertainty-aware regularization method. Intuitively speaking, unlike a single deterministic model (e.g., a ResNet), using the MLN allows us to model multi-modal (and possibly noisy) target distributions, which plays a crucial role in achieving both robustness as well as explainability. Furthermore, we present an evaluation scheme on SDN settings, which gives information about the collective outliers and label noise distribution of sets.

The main contributions of this work are threefold. 1) We propose a simple yet effective robust learning method leveraging a mixture-of-experts model on various noise settings. 2) The proposed method can not only robustly learn from noisy data but can also discover the underlying set-dependent noise pattern (i.e., the noise transition matrix) as well as the two types of predictive uncertainties (i.e., aleatoric and epistemic uncertainty) within the dataset. 3) Finally, we present a novel evaluation scheme for validating the set-dependent corruption pattern estimation performance.

2. Related work

In the context of robust learning, the label noise patterns can be roughly categorized into two groups, Class-Conditional Noise (CCN) and Instance-Dependent noise (IDN) settings, based on which information the label corruptions are made. Note that the IDN setting is much more practical as it is more natural to assume that the label noise pattern is a function of inputs. Furthermore, the IDN setting can inherently incorporate the CCN setting. Our proposed method can cope with both CCN and IDN settings. While most of the robust learning literature focuses on simply estimating the clean target distribution, a number of works have been recently made on achieving both robustness and the ability to estimate the label noise patterns.

One possible approach is to extract a clean subset, then utilize the set to learn the clean target signal. Co-Teaching [6] utilizes two separate networks (a teacher network and a student network) by teaching the student network using the teacher network. Co-Teaching+ [8] extends Co-Teaching by further leveraging the disagreement strategy. JoCoR [9] is based on Co-Teaching+, which uses a joint loss function of minimizing cross-entropy while maximizing the agreement between two networks to achieve better robustness. DivideMix [24] leverages a hybrid approach by fitting the loss using a GMM to divide the dataset into the clean labeled set and noisy unlabeled set, then utilizes MixMatch[2] with divided sets for training two separate networks. WarPI [3] achieves its robustness by adaptively rectifying the training procedure for the classification network within the meta-learning scenario.

The other approach is learning the label transition matrix (noise patterns) from noisy data to estimate the clean-class posterior. F-correction [7] estimates the noise transition matrix and applies it to loss function correction. Dual-T [17] incorporates a matrix factorization method to avoid directly estimating a noisy class posterior without any anchor points (i.e., clean data). Total variation regularization [18] effectively regularizes the predicted probability to be more distinguishable by restricting the total variance distance resulting in a better estimation of the noise transition matrix.

The IDN setting is more practical than the CCN setting in that it is more natural to assume that the annotator gets confused by ambiguous instances leading to mislabeling. Only recently, a few papers have incorporated confidence estimation and noise transition matrix prediction to handle confusing instances. However, most of the work only uses the confidence estimation to robustly learn clean target distribution in instance-dependent noise settings. Cheng et al. [19] present confidence regularization to prevent overfitting in multi-class classification problems[19]. is further extended in [21] by designing a sample sieve method to get clean in-

stances from the noisy dataset, using confidence regularized cross-entropy loss. The confidence regularized method gives information about the corruption of each instance but does not provide information about label noise patterns. Another approach is to estimate the noise transition matrix instance-wise to correct the loss function. Part-dependent noise (i.e., PDN) was introduced by Xia et al. [22] which approximates the transition matrix by the combination of transition matrices for each instance. Yang et al.[23] first collect a predicted clean set to learn the noise transition matrix and then train a classifier with a corrected loss function based on the estimated noise transition matrix.

Perhaps, the most similar setting to ours is [10] which introduced a confidence-scored Instance-Dependent Noise (IDN) setting; a label noise is given based on prior information about confidence score by annotators. It uses both confident estimation and noise transition matrix estimation for robust learning. Then, the model utilizes the corrected loss function using the confidence-based noise transition matrix. However, the proposed Set-Dependent Noise (SDN) setting differs from the confidence-scored Instance-Dependent Noise (IDN) setting in that the annotators can divide the set by the ambiguity measure of each instance, making it more intuitive and straightforward. Furthermore, the proposed method directly estimates both corruption information and the noise distribution without the necessity of training multiple networks. The categorization of robust learning papers is shown in Table 1.

The mixture-of-experts models have been widely used in robust learning [5,27,28]. SsSMM [27] incorporates a student-teacher model similar to MentorNet [6], but employs a finite mixture models for student networks, updating via an EM algorithm [29] in semi-supervised manner. For robust learning for language domains, Irie et al. [28] proposed a recurrent adaptive mixture model to represent diverse outputs. ChoiceNet [5] utilizes a mixture density network to model the correlated outputs where the correlation between the target and noisy distributions is estimated in an end-to-end manner. Our proposed method is also based on a mixture-of-experts model; however, a novel uncertainty-aware regularization method is presented.

3. Problem formulation

3.1. Training data generation process

In this paper, we focus on the classification task of finding a mapping from an instance \mathbf{x} (e.g., an RGB image) to an output \mathbf{y} (e.g., an one-hot vector) where the input \mathbf{x} and the output \mathbf{y} are sampled from the input distribution $p(\mathbf{x})$ and the clean target distribution $p(\mathbf{y}|\mathbf{x})$. We assume that some noise patterns can be induced to both input and output where we denote $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ as the corrupted input and output, respectively. The input noise pattern, $\tilde{\mathbf{x}} \sim p(\tilde{\mathbf{x}}|\mathbf{x})$, can be adding more blur to the instance so that the resulting image is obscured or applying artificial manipulation to the image (e.g., CutMix [30]).

Roughly speaking, the output corruption process can be divided into twofold: the Class-Conditional Noise (CCN) and the Instance-Dependent Noise (IDN) settings. For the CCN setting, it is assumed that the training label information is corrupted via a single label transition matrix $T \in \mathbb{R}^{C \times C}$ where C is the number of classes and $[T]_{ij} = p(\mathbf{y}_j|\mathbf{y}_i)$ is the probability of a label i being shifted to a label j . For example, we can simply select a certain portion of the training data and shuffle the labels uniformly randomly or shift the labels by assigning label 1 to 2, label 2 to 3, and so forth. In the robust learning literature, the formal and the latter are often referred to as symmetric and asymmetric noise patterns, respectively. On the other hand, the IDN assumes that the noise pattern is a function of an instance (i.e., $T(\mathbf{x}) \in \mathbb{R}^{C \times C}$). However, as a single

Table 1
Categorization of robust learning papers.

	Class-Conditional Noise (CCN)	Instance-Dependent Noise (IDN)
Robust Learning	[5,6,8,9,14–16,24]	[10,19]
+ Noise Transition Matrix Prediction	[17,25]	[22,26], Ours

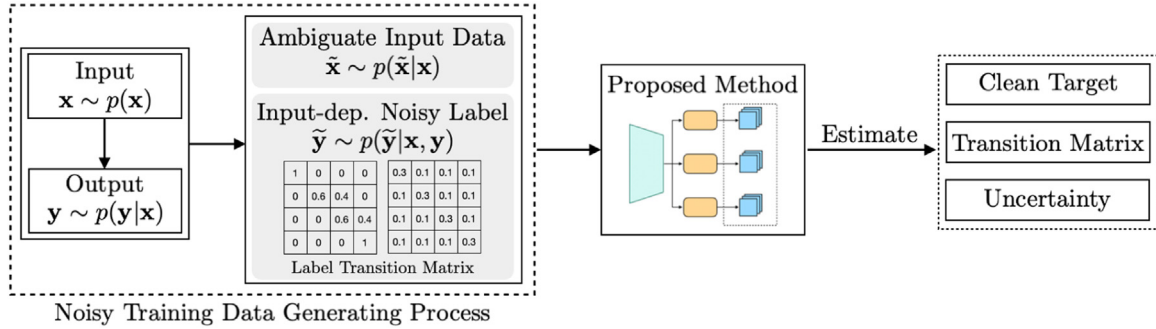


Fig. 1. General process of the proposed method.

instance can only have a single target label, it would be unrealistic to have the whole label transition matrix T per instance.

Throughout this paper, a Set-Dependent Noise (SDN) setting is utilized where we assume that the training dataset is partitioned into subsets where each subset contains its own label corruption matrices. This assumption is rather more practical in that it is more natural to assume that the annotators will be more likely to make mistakes on a specific subset consisting of hard instances. We would like to stress that our proposed method can estimate the label noise patterns in both CCN and SDN settings without the necessity of additional clean data.

3.2. Robust learning and corruption pattern estimation

The main objectives of the proposed method are twofold: the first is to robustly learn the underlying clean target signal out of noisy training data, and the other is to gain the explainability of the prediction via estimating the label corruption information as well as the predictive uncertainty. Specifically, we disentangle the total uncertainty into aleatoric and epistemic uncertainty similar to Kendall et al. [31] and will be explained in the next section. Aleatoric uncertainty corresponds to the *irreducible* part of the uncertainty, which is inherent in the data generation process (e.g., measurement noise). On the other hand, epistemic uncertainty captures the model uncertainty, which may *reduce* as we have more training data.

With respect to the label corruption information, we estimate the Set-Dependent Noise (SDN) pattern of the training dataset without the necessity of a clean validation dataset. Note that the SDN inherently handles the CCN as it can simply condition the whole data. Specifically, we estimate the label transition matrix conditioned on the subset of training of test data where the corruption rates and the noise patterns, symmetric or asymmetric, can be estimated from the transition matrix.

4. Proposed method

We present a robust learning method via a mixture-of-experts model for a classification task named mixture logit networks (MLN) and a noise pattern estimation method utilizing the outputs of the MLN. To fully utilize the multiple mixtures, we further propose an uncertainty-aware regularization method. We empirically show that this regularization method plays an influential role

in achieving both robustness and explainability. The intuition behind leveraging the mixture model is that, when given corrupted training data, the noise pattern will give rise to the discrepancy of the prediction outputs, where a single deterministic model (e.g., a ResNet) often fails to correctly capture the clean target signal. However, as a mixture model, when adequately trained, can better capture the inconsistent output patterns (including both clean and noisy distributions), it not only can robustly learn the underlying target distribution but also can model the noise patterns injected in the data generating process. The overall process of the proposed method is illustrated in Fig. 1.

The MLN architecture is illustrated in Fig. 2. Suppose that the number of mixtures is K , then the MLN outputs consist of mixture weights $\{\pi_k\}_{k=1}^K$, logits $\{\mu_k\}_{k=1}^K$ where $\mu_k \in \mathbb{R}^C$ and C is the number of classes, and Mixture standard deviations (Mixture STD) $\{\sigma_k\}_{k=1}^K$. Note that only the uppermost layer is modified. Hence the total number of parameters does not change significantly.

4.1. Uncertainty estimation using the MLN

We first present ways to estimate two types of uncertainties with the MLN: epistemic (model) uncertainty and aleatoric (data) uncertainty. We denote σ_e as epistemic uncertainty and σ_a as aleatoric uncertainty. First, epistemic uncertainty is computed as follows.

$$\sigma_e^2 = \sum_{j=1}^K \left(\sum_{c=1}^C \pi_j \left\| \mu_j^{(c)}(\mathbf{x}) - \sum_{k=1}^K \pi_k \mu_k^{(c)}(\mathbf{x}) \right\|^2 \right) \quad (1)$$

where μ_k^c is logit of label c in k th mixture.

On the other hand, epistemic uncertainty (σ_e) indicates how much the model is uncertain about its prediction. (1) corresponds to the weighted average variance of each mixture's predicted logits, which can be seen as disagreements between $\{\mu_k\}_{k=1}^K$. On the other hand, aleatoric uncertainty (σ_a) is computed as follows.

$$\sigma_a^2 = \sum_{k=1}^K \pi_k \sigma_k(\mathbf{x}) \quad (2)$$

Aleatoric uncertainty captures noise inherent in observation and how much the model is uncertain about its data. (2) indicates the weighted average of each mixture's predicted STD of the given input. Mixture STD $\{\sigma_k\}_{k=1}^K$ denotes predicted noise by mixtures,

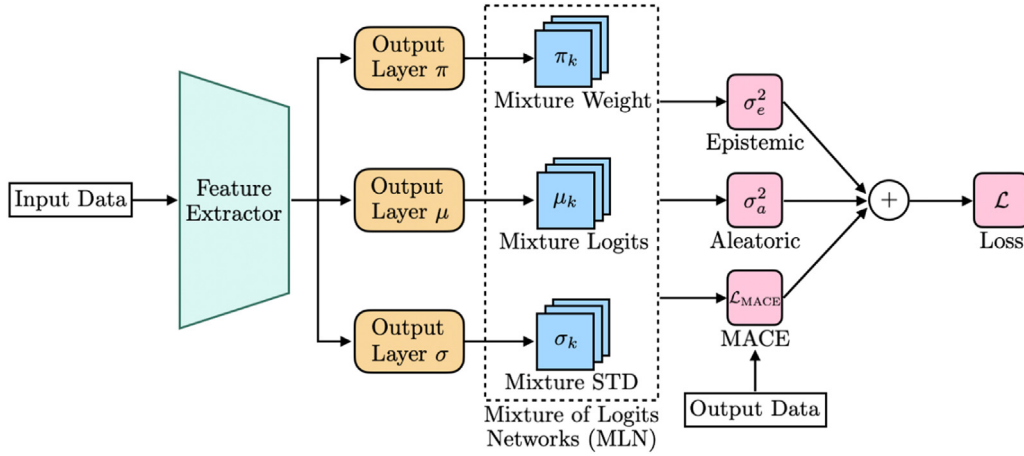


Fig. 2. The proposed mixture of logit network architecture.

also can be used as attenuation factor for loss function, similar to Kendal et al. [31].

4.2. Mixture of the attenuated losses

We present a Mixture of the Attenuated Cross-Entropy (MACE) loss for effectively training the MLN. We denote the target as y_i , which can either be clean or noisy (i.e., \tilde{y}_i) depending on the dataset. The proposed loss function consists of cross-entropy loss divided by the standard deviation of each mixture (i.e., loss attenuation) and then a weighted summation of the attenuated loss for each mixture. Mixture standard deviation $\{\sigma_k\}_{k=1}^K$ corresponds to the expected measurement noise of each instance.

The MACE loss function is defined as follows:

$$\mathcal{L}_{MACE} = \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \pi_k(\mathbf{x}_i) \frac{l(\boldsymbol{\mu}_k(\mathbf{x}_i), y_i)}{\sigma_k(\mathbf{x}_i)} \quad (3)$$

where $l(\boldsymbol{\mu}_k(\mathbf{x}_i), y_i)$ is the cross entropy loss.

If an input is ambiguous or corrupted, it becomes more likely to make a false prediction. Then σ_k will increase to reduce the overall loss function of the prediction. As a result, this attenuated factor prevents the overfitting of the model to the corrupted dataset, making the proposed model more robust.

4.3. Uncertainty-aware regularization method

We observe that there could exist two possible problems when training only with the original MACE loss function: a simple increment of the mixture standard deviation will minimize the loss and the insufficient usage of mixtures. To resolve these issues, we propose a novel regularization method utilizing predictive uncertainty measures. Let us first present the proposed loss function utilizing the uncertainty measures:

$$\mathcal{L}(\mathcal{D}) = \mathcal{L}_{MACE} - \lambda_1 \sigma_e + \lambda_2 \sigma_a. \quad (4)$$

The first problem is that $\{\sigma_k\}_{k=1}^K$ in (3) will simply grow to minimize the total loss. To prevent this, we need the regularization term, $\{\sigma_k\}_{k=1}^K$. This is inspired by Kendall et al. [31], where the loss function is based on the Gaussian likelihood and is the sum of attenuated regression loss and regularization of $\{\sigma_k\}_{k=1}^K$. We present a parameter λ_2 as a weight.

Next, it is known that the mixture-of-experts model is prone to use only one or few components, incapable of capturing the various data distributions. Moreover, we observe that different mixtures are easily agreeable with each other, which is inappropriate to represent a multi-modal distribution. Due to this property, the

whole model often fails to learn both clean and corrupted data distribution, leaving the model vulnerable to noisy data. To this end, we regularize epistemic uncertainty to be large, which naturally incentivizes to utilize more mixtures where λ_1 is the weight parameter.

We illustrate effect of this regularizer in Fig. 3. The synthetic dataset consists of two-dimensional inputs (i.e., x_0 and x_1) for a binary classification problem where we assign label 0 to instances on the upper moon and label 1, otherwise. We corrupt the label information by flipping the labels at the rate of 30% and training the MLN with three mixtures. In the first column, we illustrate the clean half-moon dataset, noisy half-moon dataset, and decision boundary trained by the clean half-moon dataset. The second and third columns present the effect of the regularizer. The first row presents the estimated flipping rate, computed from (14), which will be discussed in the later section. The result indicates that the proposed regularizer helps the better prediction of the noise distribution. The second row of Fig. 3 shows that the output of each mixture disagrees with each other, making a better representation. Furthermore, the third row illustrates that the proposed regularizer smooths the decision boundary in the presence of outliers.

4.4. Corruption pattern estimation

We further use the output of the MLN to gain information about the noise corruption pattern. This can be done by estimating the noise transition matrix. The noise transition matrix $T_{ij}(\mathbf{x})$ [17,25] indicates the probabilities of the clean labels flipping to noisy labels. The notation represents the probability that the instance \mathbf{x} with the clean label $y = i$ will have a noisy label $\tilde{y} = j$. Formally, the noise transition matrix is defined as follows.

$$T_{ij}(\mathbf{x}) = P(\tilde{y} = j | y = i, \mathbf{x}) \quad (5)$$

Capturing multi-modality is one of the strengths that the mixture model possesses. This property leads the MLN to model a multi-modal distribution on a noisy instance, representing both the clean and noisy label distribution. As the label corruption patterns can be regarded as a multi-modal distribution, we introduce an auxiliary random variable z to estimate the noise transition matrix.

$$T_{ij}(\mathcal{D}_l; z) = \frac{1}{|\mathcal{D}_l(i)|} \sum_{\mathbf{x} \in \mathcal{D}_l(i)} P(\tilde{y} = j | \mathbf{x}) \quad (6)$$

here, \mathcal{D}_l denotes the set indexed l and $\mathcal{D}_l(i) = \{\mathbf{x} | y = i, (\mathbf{x}, y) \in \mathcal{D}_l\}$. For SDN setting set index can be 0 (clean) or 1 (ambiguous) and

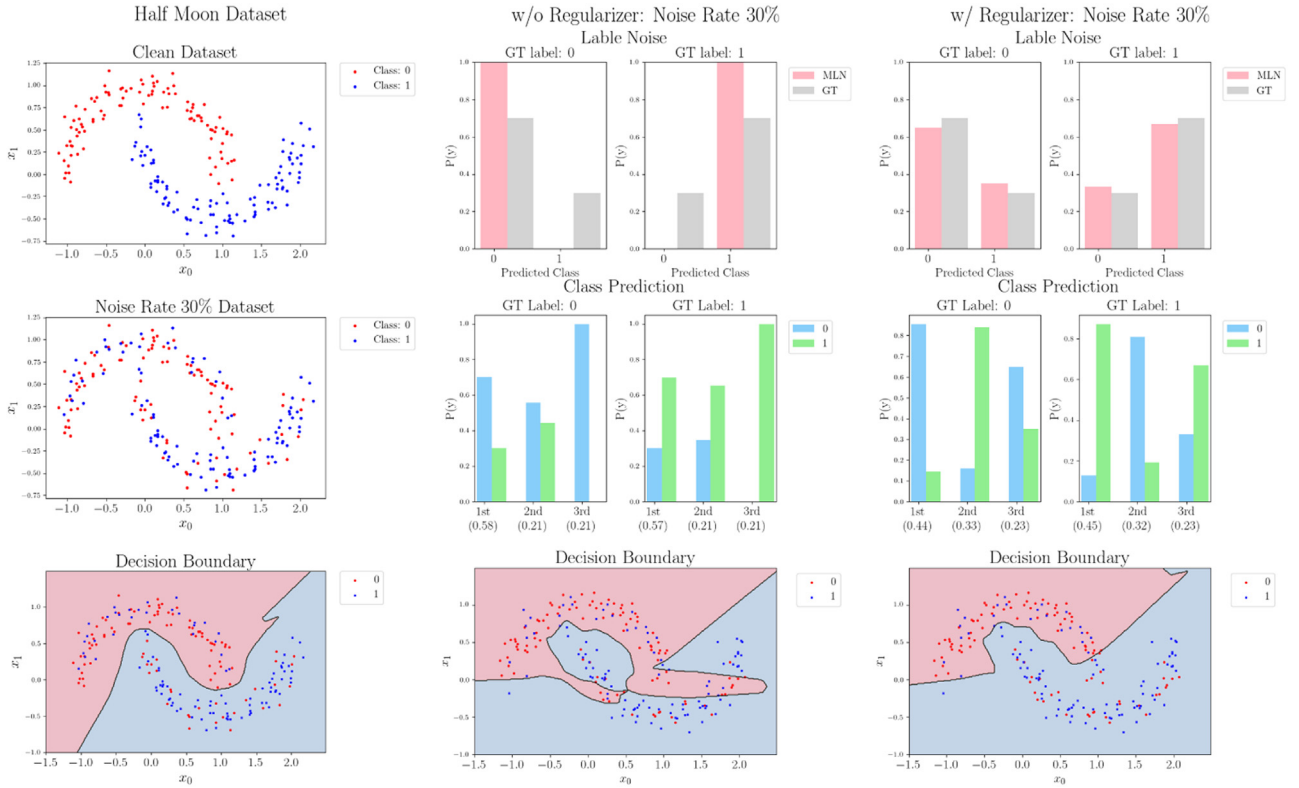


Fig. 3. Effect of the Regularizer On half-moon dataset.

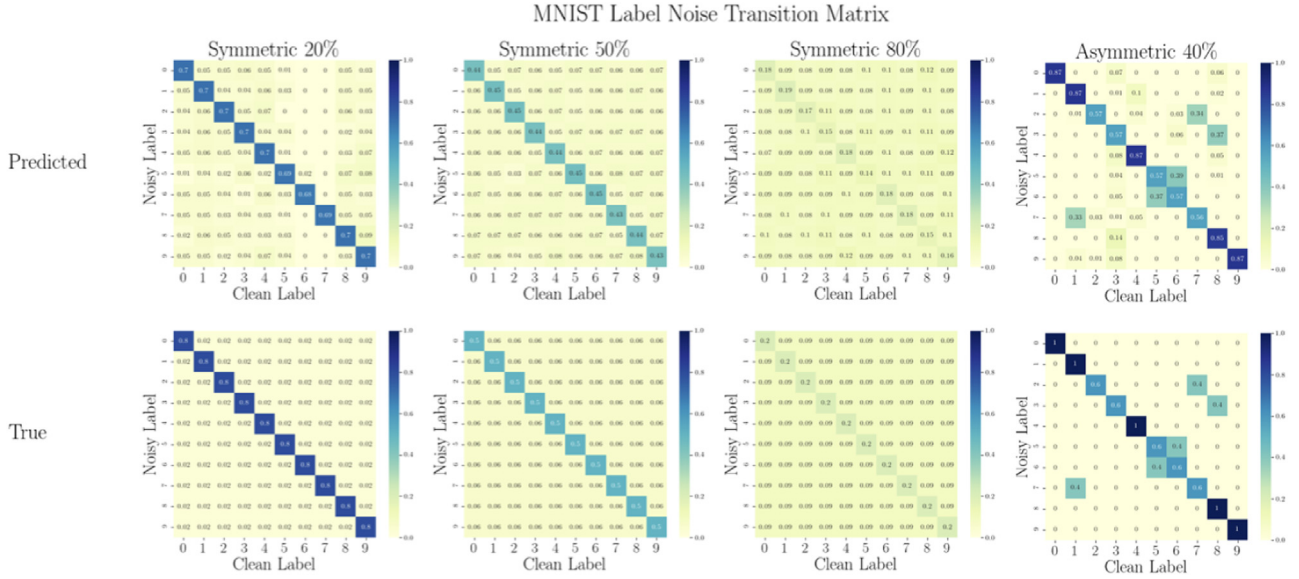


Fig. 4. Noise Transition Matrix on MNIST.

for CCN set index will be 0 (total).

$$T_{ij}(\mathcal{D}_l; z) = \frac{1}{|\mathcal{D}_l(i)|} \sum_{\mathbf{x} \in \mathcal{D}_l(i)} \sum_{k=1}^K P(z=k)P(\tilde{y}=j|z=k, \mathbf{x}) \quad (7)$$

where K is total number of mixtures, z is latent variable, and $P(z=k)$ is weight of component distribution denoted as π_k above.

Starting from (7), we denote the $\hat{P}(\mathbf{y}|\mathbf{x})$ as the soft-max output vector approximating $P(\mathbf{y}|\mathbf{x})$ by parametrized model and estimator for the noise transition matrix as \hat{T} . In addition, since we cannot observe clean y, we assume y as Bayes optimal la-

bels, the class labels that maximize the clean class posteriors $\hat{f}(\mathbf{x}) := \operatorname{argmax}_y \hat{P}(y|\mathbf{x})$. Furthermore, we define a subset $\hat{\mathcal{D}}_l(i)$ with the Bayes optimal label index i in the set $\hat{\mathcal{D}}_l$, as we do not have any prior knowledge about the clean y.

$$\hat{\mathcal{D}}_l(i) = \{\mathbf{x} | \hat{f}(\mathbf{x}) = i, \mathbf{x} \in \mathcal{D}_l\} \quad (8)$$

Then, the (7) can be modified as follows.

$$\hat{T}_{ij}(\mathcal{D}_l; z) = \frac{1}{|\hat{\mathcal{D}}_l(i)|} \sum_{\mathbf{x} \in \hat{\mathcal{D}}_l(i)} \sum_{k=1}^K \hat{P}(z=k) \hat{P}(\tilde{y}=j|z=k, \mathbf{x}) \quad (9)$$

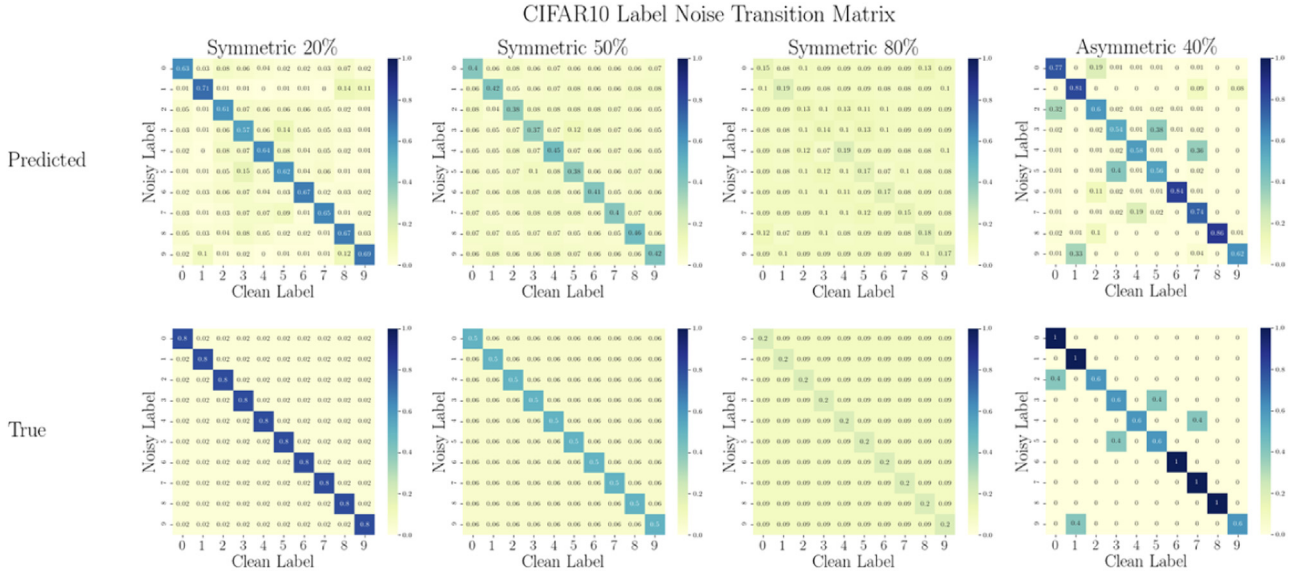


Fig. 5. Noise Transition Matrix on CIFAR10.

The Lemma 4.1 shows (9) is a valid transition matrix, since the row-wise sum of predicted transition matrix becomes one.

Lemma 4.1. *The row wise sum of the proposed transition matrix estimation in (9) becomes one.*

$$\sum_{j=1}^C \hat{T}_{ij}(D_i; z) = 1 \quad (10)$$

Proof. *The row wise sum of the proposed estimated noise transition matrix becomes as follows.*

$$\sum_{j=1}^C \hat{T}_{ij}(D_i; z) = \sum_{j=1}^C \frac{1}{|\hat{D}_i(i)|} \sum_{\mathbf{x} \in \hat{D}_i(i)} \sum_{k=1}^K \hat{P}(z = k) \hat{P}(\tilde{y} = j | z = k, \mathbf{x}) \quad (11)$$

Since j is not dependent to $|\hat{D}_i(i)|$ and $\hat{P}(z = k)$, the equation can be rewritten as follows.

$$\sum_{j=1}^C \hat{T}_{ij}(D_i; z) = \frac{1}{|\hat{D}_i(i)|} \sum_{\mathbf{x} \in \hat{D}_i(i)} \sum_{k=1}^K \hat{P}(z = k) \sum_{j=1}^C \hat{P}(\tilde{y} = j | z = k, \mathbf{x}) \quad (12)$$

By the definition of the categorical distribution and mixture weight, $\sum_{j=1}^C \hat{P}(\tilde{y} = j | z = k, \mathbf{x}) = 1$ and $\sum_{k=1}^K \hat{P}(z = k) = 1$.

$$\sum_{j=1}^C \hat{T}_{ij}(D_i; z) = \frac{1}{|\hat{D}_i(i)|} \sum_{\mathbf{x} \in \hat{D}_i(i)} 1 = 1 \quad (13)$$

As we define the set $\hat{D}_i(i)$ as (8), the proposed noise transition matrix holds $\sum_{j=1}^C \hat{T}_{ij}(D_i; z) = 1$. \square

As the confidence of the softmax-output decreases on the noisy dataset, the predicted transition matrix often suffers from being too smooth. To better estimate the noise transition matrix, inspired by Liang et al. [32], we apply temperature scaling for the softmax activation. We set the temperature to zero, which makes the softmax function an indicator function, driving to the predicted noise transition matrix inherent to the confidence score. \hat{P}_{scaled} is defined as follows.

$$\hat{P}_{scaled}(\tilde{y} = j | z = k, \mathbf{x}) = \mathbb{I}_j(\arg\max_c \hat{P}(\tilde{y} = c | z = k, \mathbf{x})) \quad (14)$$

The scaled transition matrix is defined by replacing \hat{P} to \hat{P}_{scaled} in (9).

5. Experiments

In this section, we present experimental results of validating the robustness of the proposed method. We first describe the implementation details for experiment settings, including datasets, corruption patterns, and hyperparameters. Next, we present the results in the CCN setting and compare them with benchmarks. Furthermore, we utilize the estimated uncertainty measures to distinguish the collective outliers in the SDN settings, where the noise transition matrices of each partition are estimated and compared with the ground truth.

5.1. Implementation details

Class-Conditional Noise (CCN) Setting We first construct a Class-Conditional Noise (CCN) dataset with clean instances and noisy labels whose corruption rate is solely a function of class information. We evaluate the proposed method on four different datasets, MNIST, CIFAR10, and CIFAR100. These datasets are popularly used for evaluating the robustness of the image classification algorithms. Following JoCoR [9], we conduct experiments on four different label corruption patterns: Symmetry-20%, Symmetry-50%, Symmetry-80%, and Asymmetry-40%.

Set-Dependent Noise (SDN) setting

A Set-Dependent Noise (SDN) setting is utilized where the dataset is partitioned into two subsets: clean set and ambiguous set. In particular, we define an ambiguous set as a set containing a pair of corrupted instances and noisy labels. We experiment on two different datasets: Dirty-MNIST and Dirty-CIFAR10. Dirty-MNIST, proposed by Mukhoti et.al. [33] is formed as the union of MNIST set and Ambiguous-MNIST set. Ambiguous MNIST contains corrupted instances where it has multiple plausible labels but contains only one GT label. To conduct set-dependent noise on the Dirty-MNIST dataset, we added label noise on the Ambiguous-MNIST set. We validate on four different label noise patterns: Symmetry-20%, Symmetry-50%, Symmetry-80%, and Asymmetry-40%. We define the Dirty-CIFAR10 dataset, which contains half of the original CIFAR10 dataset and the other half ambiguated with

Table 2
MNIST Test Accuracy in CCN setting.

Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
Noise Adaptation [1]	89.0 ± 0.05	98.75 ± 0.05	14.08 ± 0.59	60.08 ± 0.01
F-correction [7]	98.37±0.28	95.70±0.6	85.33±2.28	95.19±1.29
Co-teaching [6]	99.08±0.04	98.19±0.09	85.26±0.11	96.69±0.70
Co-teaching+ [8]	99.00±0.10	98.83 ±0.11	86.24±0.16	98.65 ±0.18
JoCoR [9]	99.20 ±0.07	98.76±0.06	86.00±0.15	98.46±0.21
MLN (ours)	98.97±0.01	98.32±0.01	93.49 ± 0.21	97.13±0.10

the CutMix [30] method. We choose the CutMix method to maintain the scheme that samples on ambiguous sets should have multiple possible labels but has one GT label. Again, to form the SDN setting, we added label noise on the ambiguous set.

Hyperparameters

We use a three-layer CNN for MNIST and a seven-layer CNN for both CIFAR10 and CIFAR100 following JoCoR [9]. For the Clothing1M dataset, we use ResNet50 as a backbone. We set the batch size as 128 and use an Adam optimizer with the learning rate 10^{-3} and train the model with 200 epochs for CIFAR10, CIFAR100, and 20 epochs for MNIST. The learning rate is decayed 0.2 times for every ten epochs for CIFAR10, and CIFAR100 and 0.2 decay rate for every five epochs on the MNIST dataset. On Clothing1M, we use a SGD optimizer with a learning rate of 10^{-3} , decaying 0.1 times for every 30 epochs. The weight decay rate is set to 10^{-3} , and we train the model for 80 epochs. We set the minimum of $\{\sigma_k\}_{k=1}^K$ as one and maximum as ten by using a sigmoid function except for the IDN setting and MLN+MixUp on Section 5.2.3. In these settings, where $\{\sigma_k\}_{k=1}^K$ is set the minimum as 0.1 and maximum as 1. In addition, we set the number of the mixtures to be 20 for all experiments, which should be large enough to cover all the noise distribution. In addition, when using MixUp [4] augmentation, we tune α as four. Furthermore, we set regularizer hyperparameters as $\lambda_1 = 1, \lambda_2 = 1$ except for CIFAR100, where we scale the parameter to $\lambda_1 = 0.1, \lambda_2 = 1$, and the IDN setting where $\lambda_1 = 10$. Regularizer parameters are selected using cross-validation results. On cross-validation, we assume there exists a small clean set and use 10% of the clean test set as the validation set.

5.2. Class-Conditional noise (CCN) setting

This section first shows the robust learning performance of MNIST, CIFAR10, and CIFAR100 in CCN settings and compares them with supervised-learning methods. Next, we prove the validity of the estimated noise transition matrix on the CCN setting. Furthermore, as state-of-the-art models deploy semi-supervised methods to deal with noisy labels, we combine the proposed method with a semi-supervised method to observe the effectiveness of the proposed method.

5.2.1. Robust learning accuracy

We conduct robust learning experiments with the Class-Conditional Noise (CCN) setting to investigate the performance of the MLN. We evaluate the test accuracy on four datasets with four different noise patterns and compare with Noise Adaptation [1], F-correction [34], Co-teaching [6], Co-teaching+ [8] and JoCoR [9]. The test accuracy on MNIST is shown in Table 2. The proposed method outperforms on the Symmetry-80% setting, and on other noise settings, it is compatible with the compared methods. However, the results on CIFAR10 in Table 3 show that our method outperforms the compared methods on Symmetry-80% and Asymmetry-40%, with the second-best performance on other noise patterns. Furthermore, Table 4 presents the test accuracy on CIFAR100 dataset. The MLN outperforms on Symmetry-80% and Asymmetry-40% noise patterns and performs second-best on the

rest of the noise patterns. We would like to note that the proposed method shows its strengths in heavy corruptions, such as Symmetry-80% and Asymmetry-40%. The proposed method shows a significant performance margin on large corruption rates, such as Symmetry-80% and Asymmetry-40%.

5.2.2. Noise transition matrix estimation

In this section, we evaluate the noise transition matrix estimation on the Class-Conditional Noise (CCN) setting. We estimate the noise transition matrix using (14), which is an anchor-free method that does not require a clean validation set. We evaluate the noise transition matrix on MNIST and CIFAR10 datasets with Symmetry-20%, Symmetry-50%, Symmetry-80%, Asymmetry-40%, Dual-40%, and Tridiagonal-60% noise patterns¹.

We report the average total variation (ATV) and Kendall Tau rank distance (KTD) [35] to evaluate the transition matrix estimation, which are defined as follows:

$$\text{Average total variance} = \frac{1}{C} \sum_{i=1}^C \frac{1}{2} \sum_{j=1}^C |T_{ij} - \hat{T}_{ij}| \tag{15}$$

$$\text{Average Kendall Tau Distance} = \frac{1}{C} \sum_{i=1}^C \sum_{j,k} \bar{K}_{j,k}(\mathbf{t}_i, \hat{\mathbf{t}}_i) \tag{16}$$

where \mathbf{t}_i denotes for i -th row vector of transition matrix T and $\hat{\mathbf{t}}_i$ for i -th row vector of \hat{T} .

The function \bar{K} for two arbitrary vectors \mathbf{t}_1 and \mathbf{t}_2 is defined as follows:

$$\bar{K}_{j,k}(\mathbf{t}_1, \mathbf{t}_2) = \begin{cases} 0 & \text{if } j \text{ and } k \text{ are in same order of ranking} \\ 1 & \text{if not} \end{cases} \tag{17}$$

The total variance is an average l1 norm between two matrices, which denotes the absolute difference between the estimated and ground-truth matrices. Kendall-tau rank distance is defined as a metric that counts the number of pairwise disagreements between two ranking lists, which means comparing the ranking of two matrices.

We measure ATV and KTD with comparison to Noise Adaptation [1]. We set compared method to Noise Adaptation since it estimates the noise transition matrix per image, which can be further expanded to Set-Dependent Noise (SDN) settings on Section 5.3.3. The result is shown in Table 5, indicating that the proposed method outperforms the compared method except for CIFAR10 Asymmetry-40% with ATV. However, the proposed method outperforms on KTD measure, meaning that the proposed method captures the rank of the ground-truth noise transition matrix more precisely.

5.2.3. Comparison with semi-supervised method

For combating noisy labels, the state-of-the-art methods deploy semi-supervised methods with additional augmentation methods,

¹ We call Dual for two mislabeled classes and Tridiagonal for three mislabeled classes

Table 3
CIFAR10 Test Accuracy in CCN setting.

Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
Noise Adaptation [1]	82.03 ± 0.1	39.88 ± 0.21	10.0 ± 0.0	72.44 ± 0.1
F-correction [7]	68.74±0.20	42.71±0.42	15.88±0.42	70.60±0.40
Co-teaching [6]	78.23±0.27	71.30±0.13	26.58±2.22	73.78±0.22
Co-teaching+ [8]	78.71±0.34	57.05±0.54	24.19±2.74	68.84±0.20
JoCoR [9]	85.73±0.19	79.41±0.25	27.78±3.06	76.36±0.49
MLN (ours)	84.20±0.05	77.88±0.07	41.83±0.10	76.62±0.07

Table 4
CIFAR100 Test Accuracy in CCN setting.

Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
Noise Adaptation [1]	30.14 ± 0.09	2.83 ± 0.05	1.0 ± 0.0	22.23 ± 0.1
F-correction [7]	37.95±0.10	24.98±1.82	2.10±2.23	25.94±0.44
Co-teaching [6]	43.73±0.16	34.96±0.50	15.15±0.46	28.35±0.25
Co-teaching+ [8]	49.27±0.03	40.04±0.70	13.44±0.37	33.62±0.39
JoCoR [9]	53.01±0.04	43.49±0.46	15.49±0.98	32.79±0.35
MLN (ours)	51.60±0.08	42.22±0.07	19.88±0.14	36.36±0.10

Table 5
Evaluation of Noise Transition Matrix ATV in CCN setting (average total variation)(x100), KTD(Kendall-Tau distance).

	Noise Rate	Noise Adaptation [1]		MLN(Ours)	
		ATV	KTD	ATV	KTD
MNIST	Symmetry-20%	22.19	0.4472	15.40	0.4472
	Symmetry-50%	48.67	0.4472	8.46	0.4472
	Symmetry-80%	3.41	0.4472	2.33	0.4472
	Asymmetry-40%	13.36	0.5237	10.19	0.5349
CIFAR10	Symmetry-20%	23.15	0.4472	20.71	0.4472
	Symmetry-50%	34.36	0.4472	9.69	0.4472
	Symmetry-80%	6.86	0.4472	5.85	0.4472
	Asymmetry-40%	14.18	0.5236	16.81	0.4948

e.g., Dividemix [24]. On the other hand, the proposed method focuses on robust architecture and loss function to deal with this problem. As mentioned in Xia et al. [20], to make the comparison fair, we combine our method with semi-supervised learning. Making the architecture and loss function robust to label noise in a semi-supervised framework will prevent the model from learning the wrong target signal when the estimated clean set is still noisy.

Instead of a linear classification head, we added a Mixture of Logit network heads to predict the label. To fit per-sample loss distribution, we set the selection function as follows

$$\ell(\theta) = \{\ell_i\}_{i=1}^N = \left\{ \sum_{c=1}^C p(\tilde{y} = c|\mathbf{x}) \log(\hat{p}(y = c|\mathbf{x}, z = k; \theta)) \right\} \quad (18)$$

where $k = \arg \max_k p(z = k|\mathbf{x}; \theta)$.

In addition, for making the pseudo-label, we replace $p_{model}(x; \theta)$ as $\hat{p}(y = c|\mathbf{x}, z = k; \theta)$ where $k = \arg \max_k p(z = k|\mathbf{x}; \theta)$. For updating the model, we replace cross-entropy loss in labeled

set with the mixture cross-entropy loss function with epistemic uncertainty regularizer. The labeled loss is as follows:

$$\mathcal{L}_{\mathcal{X}} = -\frac{1}{|\mathcal{X}^l|} \sum_{(\mathbf{x}, y) \in \mathcal{X}^l} \left[\sum_{k=1}^K \hat{p}(z = k) \sum_{c=1}^C y^c \log(\hat{p}(y = c|\mathbf{x}, z = k; \theta) - \lambda \cdot \sigma_e) \right] \quad (19)$$

where σ_e is same as Eq. 1 and we set $\lambda = 1$. We did not learn the $\{\sigma_i\}_{i=1}^K$ in this architecture. This is because as the DivideMix [24] framework first has a warm-up phase with a noisy set and then divides the set with an estimated clean labeled set, learning $\{\sigma_i\}_{i=1}^K$ will be unstable as the training set becomes relatively clean after warm-up phase during training.

The clean test set accuracies in CIFAR10 and CIFAR100 datasets are shown in Table 6. First of all, we observe that using additional data augmentation and regularizer like MixUp [4] would lead to a performance increase. In addition, on the CIFAR10 dataset, we found out that the hybrid method of DivideMix and MLN outperforms naive DivideMix with the gap of 1.66%, 3.58%, 11.67%, 6.08% for each noise setting. In the CIFAR100 dataset, the proposed method outperforms DivideMix except for the Symmetry-50% noise rate. In Table 7, we measure the AUROC of the partitioning clean labeled set and noisy unlabeled set during training. We observe that when the noise is heavy, DivideMix fails to divide the clean set. If we use DivideMix with MLN, then the model tends to learn robustly even when the estimated clean sets are still dirty. We found out that this property leads to better performance.

5.3. Set-Dependent noise (SDN) setting

In this section, we show the effectiveness of the MLN on SDN settings. We use the Dirty-MNIST and Dirty-CIFAR10 datasets explained in Fig. 6. These datasets contain a clean set and an am-

Table 6
Comparison with SSL.

	Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
CIFAR10	DivideMix [24]	85.06 ± 0.09	85.08 ± 0.09	53.24 ± 0.41	76.32 ± 0.17
	MLN (Ours)	84.20 ± 0.05	77.88 ± 0.07	41.83 ± 0.10	76.72 ± 0.11
	MLN (Ours) + MixUp [4]	87.77±0.08	84.59±0.09	42.38±0.21	75.35±0.12
	MLN + DivideMix [24]	86.82 ± 0.16	88.76 ± 0.11	64.91 ± 0.3	82.40 ± 0.21
CIFAR100	DivideMix [24]	62.29 ± 0.17	58.17 ± 0.16	40.28 ± 0.17	46.83 ± 0.21
	MLN (Ours)	51.60 ± 0.08	42.22±0.07	19.88±0.14	36.36±0.10
	MLN (Ours) + MixUp [4]	55.15±0.22	46.75±0.17	26.17±0.21	41.38±0.19
	MLN (Ours) + DivideMix [24]	62.70 ± 0.05	57.68 ± 0.13	40.55 ± 0.24	49.22 ± 0.34

Table 7
AUROC of Dividing Train set on CIFAR10.

Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
DivideMix [24]	0.979	0.982	0.635	0.872
MLN + DivideMix	0.987	0.977	0.909	0.934

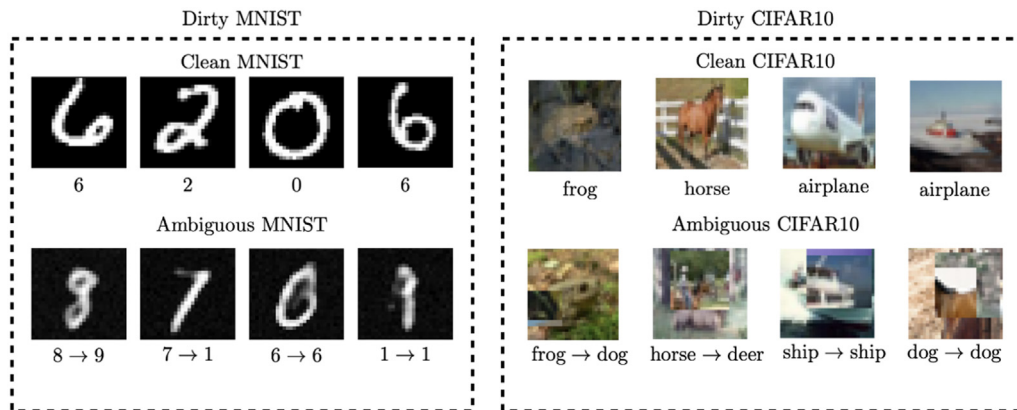


Fig. 6. Example of Set-Dependent Noise Dataset. The arrows denote the label corruption.

Table 8
Dirty-MNIST Test Accuracy in SDN setting.

Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
Noise Adaptation [1]	79.17 ± 0.02	89.31 ± 0.03	89.14 ± 0.02	79.12 ± 0.02
F-correction [7]	99.17 ± 0.04	99.14 ± 0.04	99.18 ± 0.04	99.26 ± 0.02
Co-teaching [6]	99.32 ± 0.02	99.04 ± 0.03	90.93 ± 0.57	98.68 ± 0.07
Co-teaching+ [8]	98.67 ± 0.06	99.00 ± 0.06	99.06 ± 0.07	98.78 ± 0.08
JoCoR [9]	98.81 ± 0.06	98.13 ± 0.01	98.09 ± 0.0	98.77 ± 0.03
DivideMix [24]	98.11 ± 0.01	98.46 ± 0.01	98.81 ± 0.0	97.94 ± 0.01
MLN (ours)	99.26 ± 0.01	99.45 ± 0.01	99.40 ± 0.02	99.31 ± 0.01

Table 9
Dirty-CIFAR10 Test Accuracy in SDN setting.

Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
Noise Adaptation [1]	72.35 ± 0.06	62.94 ± 0.13	72.29 ± 0.13	72.5 ± 0.09
F-correction [7]	81.55 ± 0.27	77.93 ± 0.4	79.74 ± 0.24	82.73 ± 0.15
Co-teaching [6]	87.64 ± 0.12	83.09 ± 0.13	56.08 ± 0.14	81.33 ± 0.11
Co-teaching+ [8]	86.02 ± 0.11	84.89 ± 0.16	71.55 ± 0.06	85.96 ± 0.19
JoCoR [9]	87.75 ± 0.08	82.72 ± 0.08	48.19 ± 0.07	87.46 ± 0.05
DivideMix [24]	80.10 ± 0.11	84.46 ± 0.15	85.34 ± 0.15	73.72 ± 0.1
MLN (ours)	86.24 ± 0.06	86.59 ± 0.07	87.45 ± 0.05	87.79 ± 0.04

ambiguous set, where an ambiguous set is composed of ambiguous instances with corrupted labels. We first show that the robust learning performance on the SDN setting outperforms the previous method, including a semi-supervised method with an additional regularizer like DivideMix [24]. Next, we validate proposed uncertainty measure can partition the clean and ambiguous set. Third, we evaluate the estimated noise transition matrix as well. Finally, we experiment on the Instance-Dependent Noise setting, which can be seen as an extreme case of the Set-Dependent Noise setting where the number of the set is the same as the number of instances.

5.3.1. Robust learning accuracy

First, we evaluate the accuracy of clean test set of Dirty-MNIST dataset and Dirty-CIFAR10 dataset, shown in Table 8 and 9 respectively. We observe that the proposed method outperforms compared methods except for Symmetry-20% noise rate, where the noise ratio is small. We would like to emphasize that the proposed method works better than the semi-supervised method like DivideMix [24] even without any data augmentation like MixUp

[4] or MixMatch [2]. This is because, in the SDN setting, the model has to be robust not only to label corruption but also to noisy input as well. MLN obtains its robustness on both label corruption and input noise by mixture-of-experts architecture and its uncertainty. We would like to emphasize that MLN has its strength in robust learning on the corruption of output is dependent on *quality* of input. SDN setting can be more practical in that it is more natural to assume that the annotators will be more likely to make mistakes on a specific subset consisting of hard instances.

5.3.2. Partitioning sets

Next, we show the ability of the proposed method to partition the dataset leveraging aleatoric uncertainty estimated from the MLN. In particular, the collective outliers, partitions of training data with corrupted labels, are well captured via aleatoric uncertainty.

We measure AUROC of partitioning two sets, compared with Noise Adaptation [1] and state-of-the-art method DivideMix [24]. Metric for partitioning set on Noise Adaptation [1] is KL divergence between identity matrix and instance wise estimated noise transi-

Table 10
Measured AUROC for Partitioning Sets.

	Noise Rate	Noise Adaptation [1]	DivideMix [24]	MLN (Ours)
Dirty MNIST	Symmetry-20%	0.3549	0.8859	0.9895
	Symmetry-50%	0.5778	0.9161	1.0000
	Symmetry-80%	0.5544	0.9269	1.0000
	Asymmetry-40%	0.3901	0.8714	0.9895
Dirty CIFAR10	Symmetry-20%	0.5396	0.6017	0.9061
	Symmetry-50%	0.8353	0.6378	0.9924
	Symmetry-80%	0.9932	0.6751	0.9985
	Asymmetry-40%	0.4500	0.5085	0.7117

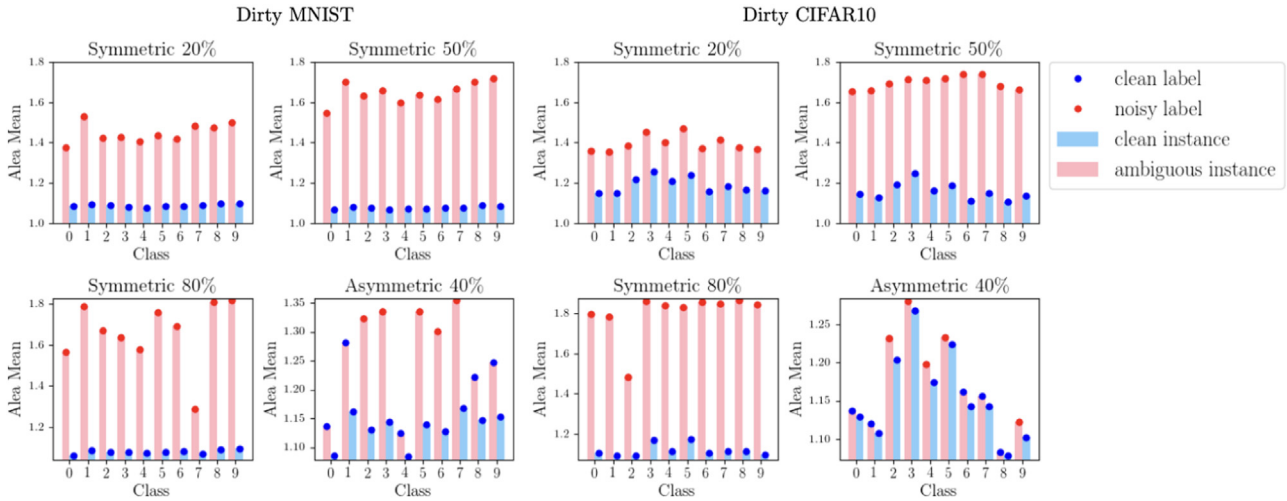


Fig. 7. Aleatoric uncertainty along each label on Dirty MNIST and Dirty CIFAR10.

tion matrix, since the estimated noise transition matrix of clean instances will be close to identity matrix if trained properly. In addition, we use soft-max entropy as a metric on DivideMix [24]. The result are shown in Table 10. We have found out that the proposed aleatoric uncertainty measure outperforms all of the compared methods, indicating that aleatoric uncertainty from MLN is valid for capturing collective outliers in the noisy datasets.

Fig. 7 shows the average of aleatoric uncertainty for both clean and ambiguous sets among each class. The result first shows that the aleatoric uncertainty is higher in ambiguous instances. Furthermore, shown in the symmetric noise cases, a heavy label corruption rate leads to higher aleatoric uncertainty. Finally, Dirty-MNIST with Asymmetry-40% label noise case shows that aleatoric uncertainty increases in the corrupted labels compared to clean labels in ambiguous instances. This demonstrates that both corruption in instances and corruption in the label are related to aleatoric uncertainty.

5.3.3. Noise transition matrix estimation

We then examine the ability of our method to estimate noise transition matrices per each group where we partition the training data into multiple sets using the predicted aleatoric uncertainty in Section 5.3.2. We first evaluate Average total variance and Kendall-Tau distance on *estimated ambiguous set* comparing with Noise-Adaptation [1] and DivideMix [24] same as Section 5.3.2. For Noise-Adaptation [1], we average the estimated noise transition matrix on the estimated set. For DivideMix [24] we average the softmax output to obtain the estimated noise transition matrix since the softmax output can be seen as a confidence score that can be interpreted as a noise ratio. We cannot use a confusion matrix for this setting because the noise transition matrix is estimated on an unseen test set. The noise transition matrix is defined as follows:

$$T_{i,j}^{\text{DivideMix}} = \hat{P}(y = j | \mathbf{x}, f(\mathbf{x}) = i) \tag{20}$$

where $f(\mathbf{x})$ is a prediction, \hat{P} is a soft-max output.

The experimental results shown in Table 11. Except for Symmetry-20% on the Dirty-MNIST dataset, the proposed method outperforms the compared method. The gap increases as the noise ratio increases since other methods have a tendency to have over-confident estimates even if the noise ratio is high.

Furthermore, the estimated noise transition matrix is illustrated in Fig. 8 and 9 on the Dirty-MNIST and Dirty-CIFAR10 datasets respectively. Each quarter of the figure denotes a single experiment for each corruption pattern, with the upper and lower rows showing the predicted noise transition matrix and the ground truth, respectively. Here, clean labels on the ambiguous set denote the ground-truth label of each instance after ambiguating the instances. We can see that our proposed method is able to correctly estimate the noise transition matrix for both clean and ambiguous sets in terms of a row-wise ranking manner. Furthermore, Fig. 9 suggests that our proposed method can also capture the noise-induced in inputs (i.e., CutMix [30]). In other words, the images of cats are cut-mixed with the images of dogs (and vice versa), and these corruption patterns are well captured by the noise transition matrix.

5.3.4. Instance-dependent noise

We further experiment on the Instance-Dependent Noise setting, which can be seen as an extreme version of the Set-Dependent Noise (SDN) setting where the number of the set is the same as the number of instances. We measure clean test accuracy on noise ratio of 20% and 40% on MNIST and CIFAR10 datasets, respectively. The result are shown in Table 12.

We observe that MLN with MixUp augmentations has a significant performance increase compared to MLN without any ad-

Table 11
 Evaluation of Noise Transition Matrix on SDN setting only with estimated ambiguous set Average total variation(x100)(Kendall-Tau distance).

	Method	Symmetry-20%	Symmetry-50%	Symmetry-80%	Asymmetry-40%
Dirty MNIST	Noise Adaptation [1]	25.68 (0.4472)	45.83 (0.4472)	73.11 (0.4472)	25.42 (0.5164)
	DivideMix [24]	14.00 (0.4472)	21.12 (0.4472)	46.88 (0.4472)	33.02 (0.4875)
	MLN (Ours)	22.43 (0.4472)	13.70 (0.4472)	7.99 (0.4472)	21.63 (0.5326)
Dirty CIFAR10	Noise Adaptation [1]	29.22 (0.4472)	49.29 (0.4472)	66.26 (0.4472)	20.15 (0.5230)
	DivideMix [24]	17.74 (0.4472)	19.06 (0.4472)	46.12 (0.4472)	33.60 (0.5237)
	MLN (Ours)	21.14 (0.4472)	14.59 (0.4472)	7.94 (0.4472)	18.93 (0.5090)

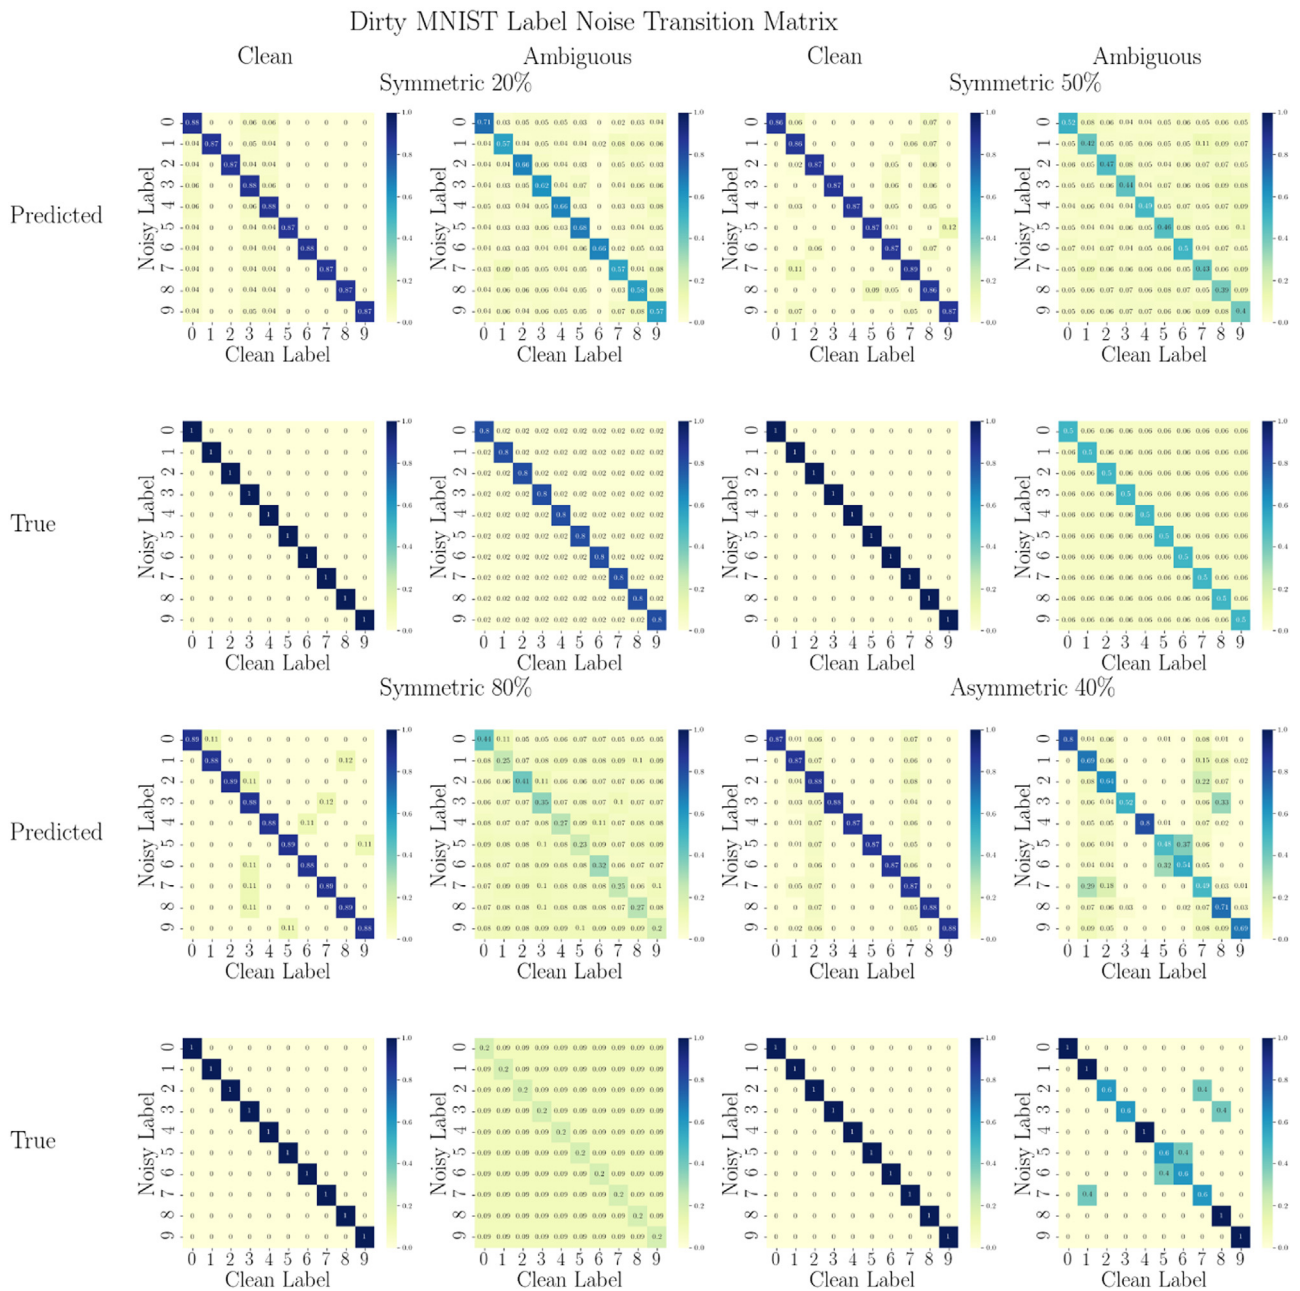


Fig. 8. Noise Transition Matrix on Dirty-MNIST.

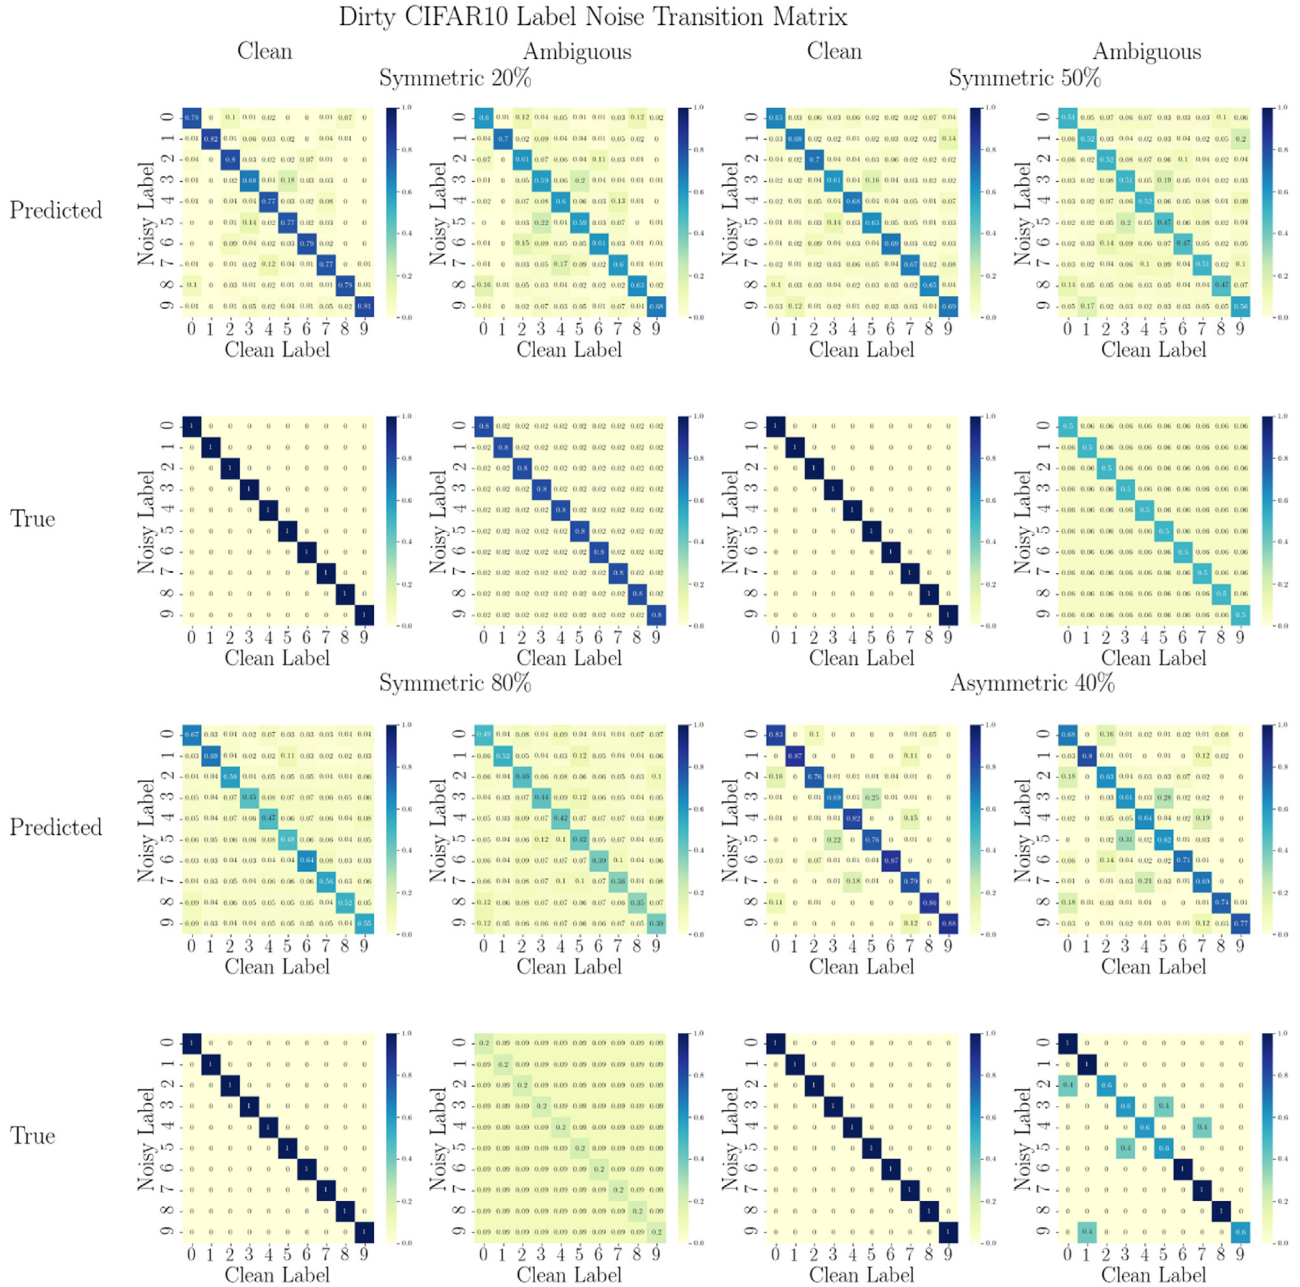


Fig. 9. Noise Transition Matrix on Dirty-CIFAR10.

Table 12 Instance-Dependent Noise (IDN) Setting.

Noise Rate	MNIST		CIFAR10	
	20%	40%	20%	40%
Noise Adaptation [1]	99.24 ± 0.03	91.03 ± 0.86	68.9 ± 0.2	45.8 ± 0.14
F-correction [7]	89.65 ± 1.54	68.66 ± 0.68	71.84 ± 2.74	48.58 ± 1.43
Co-teaching [6]	98.03 ± 0.31	95.31 ± 0.49	82.5 ± 0.11	61.33 ± 0.28
Co-teaching+ [8]	98.57 ± 0.14	98.58 ± 0.07	85.79 ± 0.16	33.76 ± 0.11
JoCoR [9]	98.94 ± 0.11	98.46 ± 0.25	85.45 ± 0.14	55.88 ± 0.23
DivdMix [24]	99.05 ± 0.03	98.97 ± 0.03	83.77 ± 0.27	61.03 ± 2.01
MLN (ours)	98.36 ± 0.0	92.65 ± 0.01	77.18 ± 0.23	55.38 ± 0.11
MLN + MixUp (ours)	98.58 ± 0.03	96.44 ± 0.12	85.09 ± 0.58	61.44 ± 0.16

Table 13
Clothing 1M Test Accuracy.

Noise Adaptation [1]	F-correction [7]	Co-teaching [6]	Co-teaching+ [8]	JoCoR [9]	MLN (ours)
67.3 ± 0.12	68.0 ± 0.10	70.33 ± 0.12	68.85 ± 0.35	71.92 ± 0.14	71.56 ± 0.05




Fig. 10. Clothing1M test set samples with low aleatoric uncertainty (left) and high aleatoric uncertainty (right). Pred denotes the predicted label and GT for the ground-truth label. 'A' denotes aleatoric uncertainty, and 'E' denotes epistemic uncertainty.

ditional augmentations. Although the proposed method does not outperform on noise settings except for CIFAR10 40%, the gap is small. At MNIST 20%, the gap is 0.66%, 2.53% for 40% and 0.60% for CIFAR 20%.

5.4. Experiments on real-world dataset

In this section, We evaluate the MLN on a real-world dataset, i.e., Clothing1M. We first show the test accuracy compared with related work. Then, we qualitatively show the proposed uncertainty can divide the whole test set into clean and ambiguous sets.

We measure the test set accuracy, shown in Table 13. The proposed method works second-best compared to other methods with a 0.36% gap. However, some instances on test set shown in Fig. 10 and 11 are still ambiguous and noisy. From this, we would like to claim that this minor gap is not significant in measuring the performance of robust learning since there still exists some noisy instances and labels in the test set.

We then sample instances from the test set, with high uncertainty instances and low uncertainty. We sample top 15 and low 15 instances for aleatoric uncertainty and epistemic uncertainty respectively, shown in Fig. 10 and 11. We first observe that set with low aleatoric uncertainty and with low epistemic uncertainty has a clean background with a relatively clean image in common. However, both a set with high aleatoric and a set with epistemic uncertainty contains some wrongly predicted instances and relatively complicated images. We observe that low aleatoric uncertainty can

capture the images containing a single object and are well centered. In addition, images with multiple clothes, i.e., wearing T-shirts inside the jacket on the fourth row in the first column, tend to have high aleatoric uncertainty. On the other hand, instances that can be interpreted as two categories, i.e., shirted style dress on the first row in the second column, tends to report high epistemic uncertainty.

5.5. Ablation studies

In this section, we conduct ablation studies on MLN. We first show how the performance of MLN varies with uncertainty regularizers (λ_1, λ_2) on the Dirty-CIFAR10 dataset. Then, we explore how other uncertainty measures perform in partitioning sets on SDN settings.

5.5.1. Effects of λ

We show how performance varies on different uncertainty regularizer (λ_1, λ_2), on Dirty-CIFAR10 dataset with Symmetry-50% and Asymmetry-40% noise ratio. The result is shown in Table 14, the value NaN happens when there do not exist any instances with the particular row. We found out that decreasing λ_1 slightly increases the accuracy of the model with the cost of degrading the performance of noise transition matrix estimation. On the other hand, increasing λ_1 may harm the model to capture the clean target signal. In addition, decreasing λ_2 decrease the accuracy of the model slightly and fails to estimate the noise distribution. Increasing λ_2




Fig. 11. Clothing1M test set samples with low epistemic uncertainty (left) and high epistemic uncertainty (right). Pred denotes the predicted label and GT for the ground-truth label. 'A' denotes aleatoric uncertainty, and 'E' denotes epistemic uncertainty.

Table 14
Ablation Study on λ on SDN CIFAR10.

λ_1	λ_2	Symmetry 50%			Asymmetry 40%		
		ACC	ATV	AUROC	ACC	ATV	AUROC
$\lambda_1 = 1$	$\lambda_2 = 1$	86.59 ± 0.07	14.59	0.9924	87.79 ± 0.04	18.93	0.7117
$\lambda_1 = 0.1$	$\lambda_2 = 1$	87.15 ± 0.07	49.38	0.9964	87.34 ± 0.05	20.04	0.7385
$\lambda_1 = 10$	$\lambda_2 = 1$	76.35 ± 0.11	NaN	0.9844	77.83 ± 0.06	NaN	0.6665
$\lambda_1 = 1$	$\lambda_2 = 0.1$	85.53 ± 0.05	26.79	0.5828	86.74 ± 0.04	63.20	0.6206
$\lambda_1 = 1$	$\lambda_2 = 10$	85.42 ± 0.03	39.41	0.9515	87.2 ± 0.07	19.75	0.5284

also drops the model's accuracy slightly and fails to partition a clean and ambiguous set. As a result, we propose best choice is $\lambda_1 = 1, \lambda_2 = 1$.

5.5.2. Partitioning set with different uncertainties in SDN

There exist different uncertainty measures than the ones explained in (1) and (2), for example, it can be measured by max-softmax [36], softmax entropy, or from the entropy of mixture weights. The Pi-entropy denotes the entropy of the mixture weights. As bigger entropy can be interpreted as a lower weight of target distribution, this can be seen as the uncertainty of the inputs. In this experiment, we assume the size of two sets is the same and set threshold as the median of the uncertainty measures to partition these sets. The uncertainty measures are as follows.

$$\text{max-softmax} = 1 - \max_c \mu_k^{(c)}(x) \quad \text{where} \quad k = \underset{i}{\text{argmax}} \pi_i(x) \quad (21)$$

$$\text{softmax-entropy} = - \sum_c \mu_k^{(c)}(x) \log(\mu_k^{(c)}(x))$$

where $k = \underset{i}{\text{argmax}} \pi_i(x)$ (22)

$$\text{pi-entropy} = - \sum_k \pi_k(x) \log(\pi_k(x)) \quad (23)$$

As the observation noise is an exemplar case of aleatoric uncertainty, instance corruption patterns can be captured by aleatoric uncertainty. Table 15 reports the aleatoric uncertainty measure can partition clean and ambiguous sets on symmetric noise compared to other uncertainty measures except for Symmetry-20% noise rate in Dirty-CIFAR10 and Asymmetry-40% noise rate in Dirty-MNIST, but with a small gap. We observe that the measure AUROC increases as the noise rate increases, indicating higher noise on the label can also increase uncertainty.

5.6. Limitations

In this section, we will discuss the limitations of the proposed method. While combining the proposed method with the semi-supervised method result in an improvement in the classification performance, the proposed method does not perform as effectively as the semi-supervised method on CCN settings. In addition, the gap of clean test set accuracy is small with JoCoR [9] in the CCN setting except for symmetric 80% noise rate. However, we observe that further combining the small-loss selection method with small

Table 15
Measured AUROC over different uncertainty measures.

	Noise Rate	Aleatoric	Epistemic	π Entropy	Max Softmax	Softmax Entropy
Dirty MNIST	Symmetry-20%	0.9998	0.9986	0.9935	0.9930	0.9958
	Symmetry-50%	1.0000	0.9999	0.9996	0.9957	0.9994
	Symmetry-80%	1.0000	1.0000	0.9993	0.9959	0.9988
	Asymmetry-40%	0.9895	0.9605	0.9581	0.9898	0.9909
Dirty CIFAR10	Symmetry-20%	0.9061	0.6718	0.7614	0.9026	0.9130
	Symmetry-50%	0.9924	0.8913	0.9600	0.9588	0.9729
	Symmetry-80%	0.9985	0.9896	0.9906	0.9449	0.9765
	Asymmetry-40%	0.7117	0.5747	0.6364	0.7100	0.7047

modification or additional augmentation like MixUp [4] can improve the performance. We would like to claim that the proposed method is complementary to the small-loss selection methods to gain further robustness when the selected small-loss set is still noisy. Furthermore, although the proposed method does not outperforms state-of-the-art method on CCN and IDN settings, we would like to claim that SDN setting are more practical than CCN or IDN, because the annotator tends to mislabel the corrupted or ambiguous data.

6. Conclusion

We have introduced an uncertainty-aware robust learning framework by leveraging a mixture of the experts' model named mixture logit networks (MLN). The MLN can estimate two different types of uncertainty, epistemic and aleatoric, where the predictive uncertainty is further utilized to define a novel regularization method. We showed that the MLN could represent multimodal distributions, making the model not only robust to outliers but also able to estimate noise patterns. In addition, we found out that combining the proposed method with a semi-supervised small loss selection method can lead to further improvement. In particular, we presented a Set-Dependent Noise (SDN) learning problem where multiple corruption patterns exist per partition and proposed a novel validation scheme for estimating the corruption patterns. To tackle this problem, we leveraged aleatoric uncertainty to detect the corrupted partition and estimated the SDN patterns using the multi-modal target distribution computed from the MLN. We would like to note that uncertainty estimation on the robust learning framework plays a significant role in providing information about the corruption of each instance. The current evaluation scheme for the SDN setting relies on two assumptions: the collective outliers can be separated via the estimated aleatoric uncertainty, and a particular label transition matrix exists per each partition. One promising future research direction could be examining our proposed method to real-world datasets without applying artificial noises to both inputs and outputs.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgement

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by

the Korea government(MSIT) (No. 2019-0-00079, Artificial Intelligence Graduate School Program (Korea University)), and No. 2022-0-00612, Geometric and Physical Commonsense Reasoning based Behavior Intelligence for Embodied AI).

References

- [1] J. Goldberger, E. Ben-Reuven, Training deep neural-networks using a noise adaptation layer, in: Proc. of the International Conference on Learning Representations (ICLR), 2017.
- [2] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, C. Raffel, Mix-match: a holistic approach to semi-supervised learning, in: Proc. of the Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [3] H. Sun, C. Guo, Q. Wei, Z. Han, Y. Yin, Learning to rectify for robust learning with noisy labels, Pattern Recognit 124 (2022) 108467.
- [4] H. Zhang, M. Cisse, Y.N. Dauphin, D. Lopez-Paz, Mixup: beyond empirical risk minimization, in: Proc. of the International Conference on Learning Representations (ICLR), 2017.
- [5] S. Choi, S. Hong, K. Lee, S. Lim, Task agnostic robust learning on corrupt outputs by correlation-guided mixture density networks, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 3872–3881.
- [6] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, M. Sugiyama, Co-teaching: robust training of deep neural networks with extremely noisy labels, in: Proc. of the Conference on Neural Information Processing Systems (NeurIPS), 2018.
- [7] G. Patrini, A. Rozza, A.K. Menon, R. Nock, L. Qu, Making deep neural networks robust to label noise: A loss correction approach, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)-, 2017, pp. 1944–1952.
- [8] X. Yu, B. Han, J. Yao, G. Niu, I. Tsang, M. Sugiyama, How does disagreement help generalization against label corruption? in: Proc. of the International Conference on Machine Learning (ICML), PMLR, 2019, pp. 7164–7173.
- [9] H. Wei, L. Feng, X. Chen, B. An, Combating noisy labels by agreement: a joint training method with co-regularization, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 13726–13735.
- [10] A. Berthon, B. Han, G. Niu, T. Liu, M. Sugiyama, Confidence scores instance-dependent label-noise learning possible, in: Proc. of the International Conference on Machine Learning (ICML), PMLR, 2021, pp. 825–836.
- [11] J. Cheng, T. Liu, K. Ramamohanarao, D. Tao, Learning with bounded instance and label-dependent label noise, in: Proc. of the International Conference on Machine Learning (ICML), PMLR, 2020, pp. 1789–1799.
- [12] J. Bootkrajang, J. Chaijarwanich, Towards instance-dependent label noise-tolerant classification: a probabilistic approach, Pattern Analysis and Applications 23 (1) (2020) 95–111.
- [13] S. Karamcheti, R. Krishna, L. Fei-Fei, C.D. Manning, Mind your outliers! investigating the negative impact of outliers on active learning for visual question answering, Annual Meeting for the Association of Computational Linguistics (ACL-IJCNLP), 2021.
- [14] X. Wang, Y. Hua, E. Kodirov, N.M. Robertson, IMAE for noise-robust learning: Mean absolute error does not treat examples equally and gradient magnitude's variance matters, Arxiv:1903.12141.
- [15] X. Ma, H. Huang, Y. Wang, S. Romano, S. Erfani, J. Bailey, Normalized loss functions for deep learning with noisy labels, in: Proc. of the International Conference on Machine Learning (ICML), PMLR, 2020, pp. 6543–6553.
- [16] Y. Lyu, I.W. Tsang, Curriculum loss: Robust learning and generalization against label corruption, in: Proc. of the International Conference on Learning Representations (ICLR), 2020.
- [17] Y. Yao, T. Liu, B. Han, M. Gong, J. Deng, G. Niu, M. Sugiyama, Dual t: Reducing estimation error for transition matrix in label-noise learning, in: Proc. of the Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [18] Y. Zhang, G. Niu, M. Sugiyama, Learning noise transition matrix from only noisy labels via total variation regularization, in: Proc. of the International Conference on Machine Learning (ICML), 2021.
- [19] H. Cheng, Z. Zhu, X. Li, Y. Gong, X. Sun, Y. Liu, Learning with instance-dependent label noise: a sample sieve approach, in: Proc. of the International Conference on Learning Representations (ICLR), 2020.

- [20] X. Xia, T. Liu, B. Han, M. Gong, J. Yu, G. Niu, M. Sugiyama, Sample selection with uncertainty of losses for learning with noisy labels, in: Proc. of the International Conference on Learning Representations (ICLR), 2022.
- [21] H. Cheng, Z. Zhu, X. Li, Y. Gong, X. Sun, Y. Liu, Learning with instance dependent label noise: a sample sieve approach, in: Proc. of the International Conference on Learning Representations (ICLR), 2021.
- [22] X. Xia, T. Liu, B. Han, N. Wang, M. Gong, H. Liu, G. Niu, D. Tao, M. Sugiyama, Part-dependent label noise: Towards instance-dependent label noise, in: Proc. of the Conference on Neural Information Processing Systems (NeurIPS), Vol. 33, 2020.
- [23] S. Yang, E. Yang, B. Han, Y. Liu, M. Xu, G. Niu, T. Liu, Estimating instance-dependent label-noise transition matrix using DNNS, Arxiv:2105.13001
- [24] J. Li, R. Socher, S.C. Hoi, DivideMix: learning with noisy labels as semi-supervised learning, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020.
- [25] X. Xia, T. Liu, N. Wang, B. Han, C. Gong, G. Niu, M. Sugiyama, Are anchor points really indispensable in label-noise learning? in: Proc. of the Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [26] Q. Wang, B. Han, T. Liu, G. Niu, J. Yang, C. Gong, Tackling instance-dependent label noise via a universal probabilistic model, Arxiv:2101.05467
- [27] W. Shao, Z. Ge, Z. Song, J. Wang, Semisupervised robust modeling of multi-mode industrial processes for quality variable prediction based on student's t mixture model, IEEE Trans. Ind. Inf. 16 (5) (2019) 2965–2976.
- [28] K. Irie, S. Kumar, M. Nirschl, H. Liao, RADMM: recurrent adaptive mixture model with applications to domain robust language modeling, in: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2018, pp. 6079–6083.
- [29] G.J. McLachlan, T. Krishnan, The EM Algorithm and Extensions, vol. 382, John Wiley & Sons, 2007.
- [30] S. Yun, D. Han, S.J. Oh, S. Chun, J. Choe, Y. Yoo, CutMix: regularization strategy to train strong classifiers with localizable features, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 6023–6032.
- [31] A. Kendall, Y. Gal, What uncertainties do we need in Bayesian deep learning for computer vision? in: Proc. of the Conference on Neural Information Processing Systems (NIPS), 2017.
- [32] S. Liang, Y. Li, R. Srikant, Enhancing the reliability of out-of-distribution image detection in neural networks, in: Proc. of the International Conference on Learning Representations (ICLR), 2018.
- [33] J. Mukhoti, A. Kirsch, J. van Amersfoort, P.H. Torr, Y. Gal, Deterministic neural networks with appropriate inductive biases capture epistemic and aleatoric uncertainty, bibinfo:arXiv:2102.11582.
- [34] G. Patrini, A. Rozza, A.K. Menon, R. Nock, L. Qu, Making deep neural networks robust to label noise: a loss correction approach, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1944–1952, 2017.
- [35] M.G. Kendall, A new measure of rank correlation, Biometrika 30 (1–2) (1938) 81–93.
- [36] D. Hendrycks, K. Gimpel, A baseline for detecting misclassified and out-of-distribution examples in neural networks, in: Proc. of the International Conference on Learning Representations (ICLR), 2017.
- [37] R. Ekambaram, S. Fefilatye, M. Shreve, K. Kramer, L.O. Hall, D.B. Goldgof, R. Kasturi, Active cleaning of label noise, Pattern Recognit 51 (2016) 463–480.
- [38] J. Bootkrajang, A. Kabán, Learning kernel logistic regression in the presence of class label noise, Pattern Recognit 47 (11) (2014) 3641–3655.
- [39] Z. Lu, L. Wang, Noise-robust semi-supervised learning via fast sparse coding, Pattern Recognit 48 (2) (2015) 605–612.

Jeongeun Park received the B.S. degree in mechanical engineering from Korea University, Seoul, Korea in 2021. Currently, she is working towards the Ph.D. degree at the Department of Artificial Intelligence, Seoul, Korea University. Her current research interests include uncertainty estimation and robust learning algorithms.

Seungyoun Shin is currently an undergraduate student in Computer Science and Engineering at Dongguk University. He has been doing his internship at Robot Intelligence Lab at Korea University. His field of interests are uncertainty and computer vision in machine learning.

Sangheum Hwang received B.S. and Ph.D. degrees in industrial engineering from KAIST, Daejeon, Korea, in 2005 and 2012, respectively. Currently, he is an Assistant Professor in the Department of Industrial and Information Systems Engineering, Seoul National University of Science and Technology, Seoul, Korea. His research interests are in the areas of statistical learning algorithms, and deep learning and its practical applications.

Sungjoon Choi received the B.S. and Ph.D. degrees in electrical and computer engineering from Seoul National University, Seoul, Korea, in 2012 and 2018, respectively. Currently, he is an assistant professor with the Department of Artificial Intelligence, Korea University, Seoul, Korea. In 2018, he was a Research Scientist at Kakao Brain, Pangyo, Korea. From 2018 to 2020, he was a Post-Doctoral Associate at Disney Research, Los Angeles, CA, USA. His current research interests include nonparametric Bayesian methods, uncertainty assessments, and robust machine learning algorithms with applications to Robotics.