
Supplementary Material of “Sample Selection with Uncertainty of Losses for Learning with Noisy Labels”

Anonymous Author(s)

Affiliation

Address

email

1 A Proof of Theoretical Results

2 A.1 Proof of Theorem 1

3 For the circumstance with soft truncation, $\tilde{\mu}_z = \frac{1}{n} \sum_{i=1}^n \psi(z_i)$. As suggested in [1], we can exploit
4 $\tilde{\mu}_z^-$ and $\tilde{\mu}_z^+$ such that

$$\tilde{\mu}_z^- \leq \tilde{\mu}_z \leq \tilde{\mu}_z^+, \quad (1)$$

5 to derive a bound for $\tilde{\mu}_z$. For some positive real parameter α , we define

$$r(\tilde{\mu}_z) = \sum_{i=1}^n \psi[\alpha(z_i - \tilde{\mu}_z)] = 0. \quad (2)$$

6 Let us introduce the quantity

$$r(\theta) = \frac{1}{\alpha n} \sum_{i=1}^n \psi[\alpha(z_i - \theta)]. \quad (3)$$

7 With the exponential moment inequality [5] and the C_r inequality [9], we have

$$\begin{aligned} \exp\{\alpha n r(\theta)\} &\leq \{1 + \alpha(\mu_z - \theta) + \alpha^2[\sigma^2 + (\mu_z - \theta)^2]\}^n \\ &\leq \exp\{n\alpha(\mu_z - \theta) + n\alpha^2[\sigma^2 + (\mu_z - \theta)^2]\}. \end{aligned} \quad (4)$$

8 In the same way,

$$\exp\{-\alpha n r(\theta)\} \leq \exp\{-n\alpha(\mu_z - \theta) + n\alpha^2[\sigma^2 + (\mu_z - \theta)^2]\}. \quad (5)$$

9 If we define for any $\mu_s \in \mathbb{R}$ the bounds

$$B_-(\theta) = \mu_z - \theta - \alpha[\sigma^2 + (\mu_z - \theta)^2] - \frac{\log(\epsilon^{-1})}{\alpha n} \quad (6)$$

10 and

$$B_+(\theta) = \mu_z - \theta + \alpha[\sigma^2 + (\mu_z - \theta)^2] + \frac{\log(\epsilon^{-1})}{\alpha n}. \quad (7)$$

11 From [2] (Lemma 2.2), we obtain that

$$P(r(\theta) > B_-(\theta)) \geq 1 - \epsilon \quad \text{and} \quad P(r(\theta) < B_+(\theta)) \geq 1 - \epsilon. \quad (8)$$

12 Let $\tilde{\mu}_z^-$ be the largest solution of the quadratic equation $B_-(\theta)$ and $\tilde{\mu}_z^+$ be the smallest solution of the
13 quadratic equation $B_+(\theta)$. Also, to guarantee the solution of the quadratic equation, we assume

$$4\alpha^2\sigma^2 + \frac{4\log(\epsilon^{-1})}{n} \leq 1. \quad (9)$$

14 From [2] (Theorem 2.6), we then have

$$\tilde{\mu}_z^- \geq \mu_z - \frac{\alpha\sigma^2 + \frac{\log(\epsilon^{-1})}{\alpha n}}{\alpha - 1}, \quad (10)$$

15 and

$$\tilde{\mu}_z^+ \leq \mu_z + \frac{\alpha\sigma^2 + \frac{\log(\epsilon^{-1})}{\alpha n}}{\alpha - 1}. \quad (11)$$

16 With probability at least $1-2\epsilon$, we have $\tilde{\mu}_z^- \leq \tilde{\mu}_z \leq \tilde{\mu}_z^+$. We can choose $\alpha = \frac{n}{\sigma^2}$. Then we have

$$|\tilde{\mu}_z - \mu_z| \leq \frac{\sigma^2(n + \frac{\sigma^2 \log(\epsilon^{-1})}{n^2})}{n - \sigma^2}, \quad (12)$$

17 which holds with probability at least $1-2\epsilon$.

18 We exploit the lower bound and let $\epsilon = \frac{1}{2t}$. Then we have

$$\ell_s^* = \tilde{\mu}_s - \frac{\sigma^2(t + \frac{\sigma^2 \log(2t)}{t^2})}{n_t - \sigma^2}, \quad (13)$$

19 where n_t denotes the number of times that the example was selected in the time intervals.

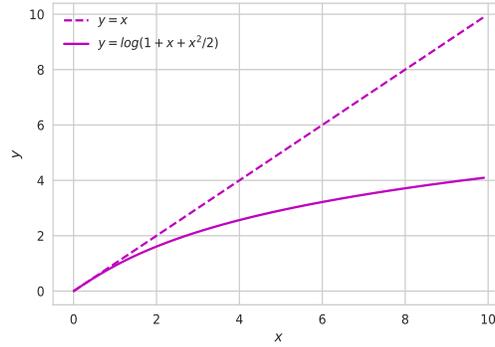


Figure 1: The illustration of the influence function for the soft estimator.

20 Here, we provide the graph of the used influence function for the soft estimator, which explains the
 21 mechanism of the function $y = \log(1 + x + x^2/2)$ more clearly. The illustration is presented in
 22 Figure 1. As can be seen, when x is large and may be an outlier, the influence function can reduce its
 23 negative impact for mean estimation. Therefore, we exploit such an influence function for robust
 24 mean estimation, which brings better classification performance.

25 A.2 Proof of Theorem 2

26 **Lemma 1** ([11]). *Let $Z_n = \{z_1, \dots, z_n\}$ be a (not necessarily time homogeneous) Markov chain*
 27 *with mean μ_z , taking values in a Polish state space $\Lambda_1 \times \dots \times \Lambda_n$, with a mixing time $\tau(v)$ (for*
 28 *$0 \leq v \leq 1$). Let*

$$\tau_{\min} = \inf_{0 \leq v < 1} \tau(v) \cdot \left(\frac{2-v}{1-v} \right)^2. \quad (14)$$

29 For some $\eta \in \mathbb{R}^+$, suppose that $f : \Lambda \rightarrow \mathbb{R}$ satisfies the following inequality:

$$f(a) - f(b) \leq \sum_{i=1}^n \eta \mathbb{1}[a_i \neq b_i], \quad (15)$$

30 for every $a, b \in \Lambda$. Then for any $\epsilon \geq 0$, we have

$$P(|f(Z_n) - \mathbb{E}f(Z_n)| \geq \epsilon) \leq 2 \exp\left(\frac{-2\epsilon^2}{\eta^2 \tau_{\min}}\right). \quad (16)$$

31 The detailed definition of the mixing time for the Markov chain can be found in [11, 12]. Let f be the
 32 mean function. Following the prior work on mean estimation [7, 4, 3, 10], without loss of generality,
 33 we assume $\mu_z = 0$ for the underlying true distribution, and $|z_i|$ is upper bounded by Z . Then we can
 34 set η to $4Z/n$ for Eq. (15). Combining the above analyses, we can revise Eq. (16) as follows:

$$P\left(\left|\frac{1}{n}\sum_{i=1}^n z_i\right| \geq \frac{2Z}{n}\sqrt{2\tau_{\min}\log\frac{2}{\epsilon_1}}\right) \leq \epsilon_1, \quad (17)$$

35 and

$$P\left(\max_{i \in [n]} |z_i| \geq \frac{2Z}{n}\sqrt{2\tau_{\min}\log\frac{2n}{\epsilon_2}}\right) \leq \epsilon_2, \quad (18)$$

36 for $\epsilon_1 > 0$ and $\epsilon_2 > 0$. If we remove the potential outliers Z_{n_o} from Z_n . Therefore, we have

$$\begin{aligned} \left|\frac{1}{n-n_o}\sum_{z_i \in Z_n \setminus Z_{n_o}} - \mu_z\right| &= \frac{1}{n-n_o}\left|\sum_{z_i \in Z_n} - \sum_{z_i \in Z_{n_o}}\right| \\ &\leq \frac{1}{n-n_o}\left(\left|\sum_{z_i \in Z_n}\right| + \left|\sum_{z_i \in Z_{n_o}}\right|\right) \\ &\leq \frac{1}{n-n_o}\left(\left|\sum_{z_i \in Z_n}\right| + n_o \max_{i \in [n]} |z_i|\right) \\ &\leq \frac{1}{n-n_o}\left(2Z\sqrt{2\tau_{\min}\log\frac{2}{\epsilon_1}} + \frac{2Zn_o}{n}\sqrt{2\tau_{\min}\log\frac{2n}{\epsilon_2}}\right), \end{aligned} \quad (19)$$

37 which holds with probability at least $1 - \epsilon_1 - \epsilon_2$.

38 For our task, we exploit the concentration inequality. Let $\epsilon_1 = \epsilon_2 = \frac{1}{2t}$, and the losses be bounded by
 39 L . Next we can obtain

$$\begin{aligned} |\tilde{\mu}_h - \mu| &\leq \frac{2L}{t-t_o}\left(\sqrt{2\tau_{\min}\log(4t)} + \frac{t_o}{t}\sqrt{4\tau_{\min}\log(4t)}\right) \\ &= \frac{2\sqrt{2\tau_{\min}}L(t+\sqrt{2}t_o)}{(t-t_o)t}\sqrt{\log(4t)} \end{aligned} \quad (20)$$

40 with the probability at least $1 - \frac{1}{t}$. In practice, it is easy to identify the value of L . For example, we
 41 can training deep networks on noisy datasets to observe the loss distributions. Then, we exploit the
 42 lower bound such that

$$\ell_h^* = \tilde{\mu}_h - \frac{2\sqrt{2\tau_{\min}}L(t+\sqrt{2}t_o)}{(t-t_o)\sqrt{t}}\sqrt{\frac{\log(4t)}{n_t}} \quad (21)$$

43 for sample selection.

44 B Complementary Experimental Analyses

	# of training	# of testing	# of class	size
<i>MNIST</i>	60,000	10,000	10	$28 \times 28 \times 1$
<i>F-MNIST</i>	60,000	10,000	10	$28 \times 28 \times 1$
<i>CIFAR-10</i>	50,000	10,000	10	$32 \times 32 \times 3$
<i>CIFAR-100</i>	50,000	10,000	100	$32 \times 32 \times 3$

Table 1: Summary of synthetic datasets used in the experiments.

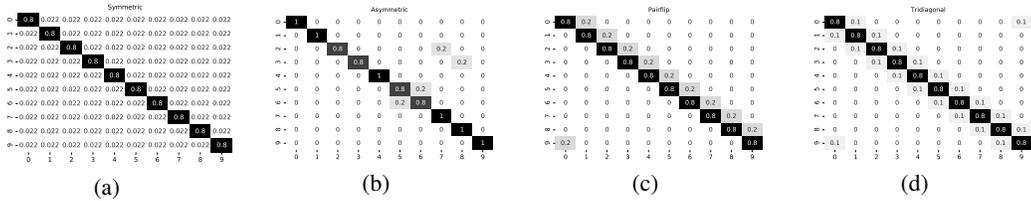


Figure 2: Synthetic class-dependent transition matrices used in our experiments on *MNIST*. The noise rate is set to 20%.

45 B.1 The Details of Datasets and Generating Noisy Labels

46 For the details of datasets, the important statistics of the used datasets are summarized in Table 1.

47 For the details of generating noisy labels, we exploit both *class-dependent* and *instance-dependent*
 48 *label noise* which include five types of synthetic label noise to verify the effectiveness of the proposed
 49 method. Here, we describe the details of the noise setting as follows:

50 (1). Class-dependent label noise:

51 • Symmetric noise: this kind of label noise is generated by flipping labels in each class uniformly to
 52 incorrect labels of other classes.

53 • Asymmetric noise : this kind of label noise is generated by flipping labels within a set of
 54 similar classes. In this paper, for *MNIST*, flipping $2 \rightarrow 7$, $3 \rightarrow 8$, $5 \leftrightarrow 6$. For *F-MNIST*, flip-
 55 ping TSHIRT \rightarrow SHIRT, PULLOVER \rightarrow COAT, SANDALS \rightarrow SNEAKER. For *CIFAR-10*, flipping
 56 TRUCK \rightarrow AUTOMOBILE, BIRD \rightarrow AIRPLANE, DEER \rightarrow HORSE, CAT \leftrightarrow DOG. For *CIFAR-100*,
 57 the 100 classes are grouped into 20 super-classes, and each has 5 sub-classes. Each class is then
 58 flipped into the next within the same super-class.

59 • Pairflip noise: the noise flips each class to its adjacent class.

60 • Tridiagonal noise: the noise corresponds to a spectral of classes where adjacent classes are easier
 61 to be mutually mislabeled, unlike the unidirectional pair flipping. It can be implemented by two
 62 consecutive pair flipping transformations in the opposite direction.

63 (2). Instance-dependent label noise:

64 • Instance noise: the noise is quite realistic, where the probability that an instance is mislabeled
 65 depends on its features. We generate this type of label noise to validate the effectiveness of the
 66 proposed method as did in [13].

67 We use synthetic noisy *MNIST* as an example and plot the noise transition matrices in Figure 2. The
 68 noise rate is set to 20%.

69 B.2 Comparison with Other Types of Baselines

70 As we focus on the sample selection approach in learning with noisy labels, in the main paper
 71 (Section 3.1), we fairly compare our methods with the baselines which also focus on sample selection.
 72 Here, we evaluate other types of baselines. We exploit APL [8] and CDR [14], which add implicit
 73 regularization from different perspectives. The experiments are conducted on *MNIST* and *F-MNIST*.
 74 Other experimental settings are the same as those in the main paper. The experimental results are
 75 provided in Table 2 and 3, which show that the proposed methods can outperform them with respect
 76 to classification performance.

77 B.3 Experiments on Synthetic *CIFAR-100*

78 For *CIFAR-100*, we use a 7-layer CNN structure from [17, 16]. Other experimental settings are the
 79 same as those in the experiments on *MNIST*, *F-MNIST*, and *CIFAR-10*. The results are provided in
 80 Table 4. We can see the proposed method outperforms all the baselines.

Noise type	Sym.		Asym.		Pair.		Trid.		Ins.	
	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
Method/Noise ratio	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
APL	98.76	94.92	98.63	88.65	98.66	68.44	98.93	76.44	97.63	87.90
	± 0.06	± 0.31	± 0.05	± 1.72	± 0.10	± 2.95	± 0.04	± 3.04	± 0.73	± 1.94
CDR	94.77	92.16	96.73	91.05	93.25	71.02	94.06	70.28	93.17	77.45
	± 0.17	± 0.73	± 0.19	± 0.76	± 0.90	± 3.89	± 0.92	± 4.01	± 0.96	± 3.04

Table 2: Test accuracy (%) on *MNIST* over the last ten epochs.

Noise type	Sym.		Asym.		Pair.		Trid.		Ins.	
	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
Method/Noise ratio	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
APL	91.73	89.06	90.13	80.34	90.22	78.54	90.84	86.53	90.96	85.55
	± 0.20	± 0.41	± 0.17	± 0.63	± 0.80	± 4.33	± 0.22	± 0.76	± 0.77	± 2.86
CDR	85.62	71.83	89.78	79.05	85.72	69.07	86.75	73.63	85.92	73.14
	± 0.96	± 1.37	± 0.41	± 1.39	± 0.65	± 2.31	± 1.19	± 2.82	± 1.43	± 3.12

Table 3: Test accuracy on *F-MNIST* over the last ten epochs.

Noise type	Sym.		Asym.		Pair.		Trid.		Ins.	
	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
Method/Noise ratio	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
S2E	44.59	25.78	42.18	26.81	42.99	26.96	43.16	27.72	43.13	27.12
	± 0.32	± 5.44	± 1.73	± 2.25	± 1.54	± 2.48	± 0.93	± 3.56	± 0.67	± 3.86
MentorNet	43.15	37.62	41.03	28.27	40.06	27.17	42.20	31.74	40.54	33.09
	± 0.42	± 0.89	± 0.22	± 0.41	± 0.37	± 0.92	± 0.30	± 0.88	± 0.69	± 1.53
Co-teaching	45.17	40.95	42.76	30.27	42.50	30.07	44.41	34.96	42.23	35.87
	± 0.25	± 0.52	± 0.34	± 0.33	± 0.39	± 0.17	± 0.41	± 0.35	± 0.52	± 1.47
SIGUA	42.03	40.53	36.67	26.71	36.48	26.73	39.21	32.69	39.19	33.51
	± 0.33	± 0.49	± 0.25	± 0.42	± 0.37	± 0.33	± 0.40	± 0.36	± 0.32	± 0.43
JoCor	45.93	41.56	42.89	29.19	42.12	30.12	44.98	34.23	44.28	35.60
	± 0.21	± 0.57	± 0.37	± 1.42	± 0.35	± 0.65	± 0.27	± 1.13	± 0.59	± 0.99
CNLCU-S	46.09	42.11	43.06	30.47	43.08	30.33	45.19	35.49	44.80	36.23
	± 0.29	± 0.70	± 0.28	± 0.37	± 0.92	± 0.74	± 0.90	± 1.30	± 0.70	± 0.49
CNLCU-H	46.27	42.05	43.21	30.55	43.25	30.79	45.02	35.24	45.02	36.17
	± 0.38	± 0.87	± 0.93	± 0.72	± 0.75	± 0.86	± 1.06	± 0.93	± 1.07	± 1.54

Table 4: Test accuracy (%) on *CIFAR-100* over the last ten epochs. The best two results are in bold.

81 B.4 Experiments for Ablation Study

82 We conduct the ablation study to analyze the sensitivity of the length of time intervals. The results
83 are shown in Figure. 3 and 4. As we can see, the proposed method, i.e., CNLCU-S and CNLCU-H
84 are robust to the choices of hyperparameters.

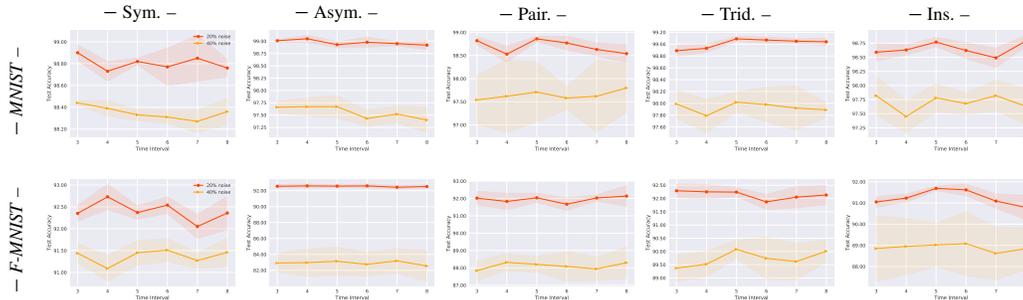


Figure 3: Illustrations of the hyperparameter sensitivity for the proposed CNLCU-S. The error bar for standard deviation in each figure has been shaded.

85 Note that in this paper, we concern uncertainty from *two aspects*, i.e., the uncertainty about small-loss
86 examples and the uncertainty about large-loss examples. Here, we conduct ablation study to show the
87 effect of removing different components to provide insights into what makes the proposed methods
88 successful. The experiments are conducted on *MNIST* and *F-MNIST*. Other experimental settings

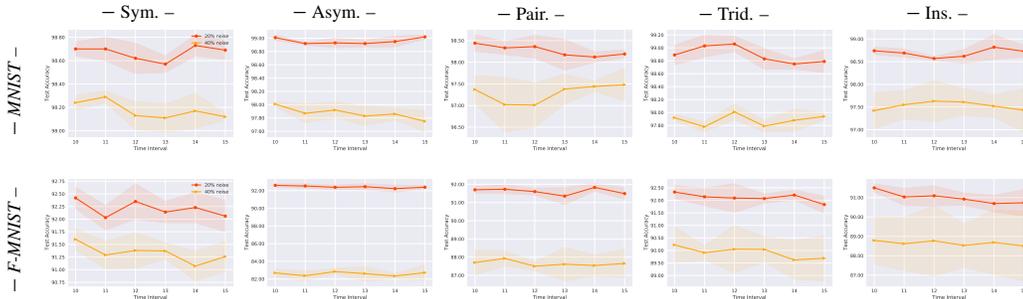


Figure 4: Illustrations of the hyperparameter sensitivity for the proposed CNLCU-H. The error bar for standard deviation in each figure has been shaded.

89 are the same as those in the main paper (Section 3.1). Note that we employ two networks to teach
 90 each other following [6]. Therefore, when we do not consider uncertainty in sample selection, the
 91 proposed methods will *reduce to* the baseline Co-teaching [6].

92 To study the effect of concerning uncertainty about small-loss examples, we remove the concerns
 93 about large-loss examples, i.e., the network is not encouraged to choose the less selected examples for
 94 updates. We express such a setting as “without concerning about large-loss examples” (abbreviated
 95 as *w/o cl*). To study the effect of concerning uncertainty about large-loss examples, we remove the
 96 concerns about small-loss examples, i.e., we only exploit the predictions of the current network. We
 97 express such a setting as “without concerning about small-loss examples” (abbreviated as *w/o cs*).
 98 Besides, we express the setting which directly uses non-robust mean as Co-teaching-M.

99 The experimental results of ablation study are provided in Table 5 and 6. As can be seen, both aspects
 100 of uncertainty concerns can improve the robustness of models. Therefore, combining two uncertainty
 101 concerns, we can better combat noisy labels. In addition, robust mean estimation is superior to the
 102 non-robust mean in learning with noisy labels.

Noise type	Sym.		Asym.		Pair.		Trid.		Ins.	
Method/Noise ratio	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
CNLCU-S	98.82 ±0.03	98.31 ±0.05	98.93 ±0.06	97.67 ±0.22	98.86 ±0.06	97.71 ±0.64	99.09 ±0.04	98.02 ±0.17	98.77 ±0.08	97.78 ±0.25
CNLCU-S <i>w/o cl</i>	98.02 ±0.08	96.83 ±0.29	98.50 ±0.04	96.25 ±0.13	98.22 ±0.13	96.08 ±0.75	98.64 ±0.31	97.25 ±0.24	98.17 ±0.20	97.13 ±0.40
CNLCU-S <i>w/o cs</i>	98.15 ±0.20	97.12 ±0.22	98.36 ±0.07	96.39 ±0.48	98.04 ±0.24	96.12 ±0.68	98.74 ±0.05	97.30 ±0.52	98.11 ±0.15	97.32 ±0.43
CNLCU-H	98.70 ±0.06	98.24 ±0.06	99.01 ±0.04	98.01 ±0.03	98.44 ±0.19	97.37 ±0.32	98.89 ±0.15	97.92 ±0.05	98.74 ±0.16	97.42 ±0.39
CNLCU-H <i>w/o cl</i>	98.06 ±0.13	96.92 ±0.23	98.39 ±0.04	96.51 ±0.57	97.04 ±0.87	95.62 ±0.93	98.33 ±0.47	97.41 ±0.92	98.01 ±0.20	96.15 ±0.28
CNLCU-H <i>w/o cs</i>	98.19 ±0.22	97.05 ±0.49	98.76 ±0.59	97.17 ±0.60	97.26 ±1.19	96.31 ±0.25	98.29 ±0.17	97.65 ±0.92	98.34 ±0.36	96.49 ±0.48
Co-teaching-M	97.72 ±0.08	97.78 ±0.32	98.27 ±0.03	95.42 ±0.42	96.22 ±0.10	95.01 ±0.65	97.92 ±0.14	96.64 ±0.77	98.02 ±0.04	96.03 ±0.57
Co-teaching	97.53 ±0.12	95.62 ±0.30	98.25 ±0.08	95.08 ±0.43	96.05 ±0.96	94.16 ±1.37	98.05 ±0.06	96.18 ±0.85	97.96 ±0.09	95.02 ±0.39

Table 5: Test accuracy (%) on *MNIST* over last ten epochs.

103 C Complementary Explanation for Network Structures

104 Table 7 describes the 9-layer CNN [6] used on *MNIST*, *F-MNIST*, and *CIFAR-10*. Table 8 describes
 105 the 9-layer CNN [17] used on *CIFAR-100*. Here, LReLU stands for Leaky ReLU [15]. The slopes of
 106 all LReLU functions in the networks are set to 0.01. Note that that the 7/9-layer CNN is a standard
 107 and common practice in weakly supervised learning. We decided to use these CNNs, since then the
 108 experimental results are directly comparable with previous approaches in the same area, i.e., learning
 109 with noisy labels.

Noise type	Sym.		Asym.		Pair.		Trid.		Ins.	
	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
Method/Noise ratio	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
CNLCU-S	92.37 ±0.15	91.45 ±0.28	92.57 ±0.15	83.14 ±1.77	92.04 ±0.26	88.20 ±0.44	92.24 ±0.17	90.08 ±0.34	91.69 ±0.10	89.02 ±1.02
CNLCU-S <i>w/o cl</i>	91.77 ±0.35	89.40 ±0.26	91.25 ±0.30	72.93 ±2.63	91.53 ±0.17	87.31 ±0.59	91.31 ±0.52	89.50 ±0.32	91.09 ±0.13	88.45 ±0.57
CNLCU-S <i>w/o cs</i>	91.85 ±0.33	90.76 ±0.28	91.94 ±0.09	80.99 ±2.74	91.28 ±0.20	87.31 ±0.72	91.39 ±0.07	89.29 ±0.51	90.98 ±0.43	88.73 ±0.62
CNLCU-H	92.42 ±0.21	91.60 ±0.19	92.60 ±0.18	82.69 ±0.43	91.70 ±0.18	87.70 ±0.69	92.33 ±0.26	90.22 ±0.71	91.50 ±0.21	88.79 ±1.22
CNLCU-H <i>w/o cl</i>	91.70 ±0.04	90.05 ±0.31	91.08 ±0.06	71.35 ±2.30	91.03 ±0.29	87.22 ±0.72	91.59 ±0.07	90.01 ±0.24	90.80 ±0.27	88.31 ±1.09
CNLCU-H <i>w/o cs</i>	91.82 ±0.13	90.92 ±0.42	92.45 ±0.25	80.73 ±1.63	91.21 ±0.17	87.49 ±0.32	92.08 ±0.13	89.72 ±0.24	91.21 ±0.38	88.62 ±0.73
Co-teaching-M	91.33 ±0.18	89.05 ±0.73	91.14 ±0.90	71.03 ±3.73	90.85 ±0.61	86.95 ±0.19	91.50 ±0.46	89.18 ±0.44	90.74 ±1.06	88.25 ±0.92
Co-teaching	91.48 ±0.10	88.80 ±0.29	91.03 ±0.14	68.07 ±4.58	90.77 ±0.23	86.91 ±0.71	91.24 ±0.11	89.18 ±0.36	90.60 ±0.12	87.90 ±0.45

Table 6: Test accuracy (%) on *F-MNIST* over last ten epochs.

Table 7: CNN on *MNIST*, *F-MNIST*, and *CIFAR-10*.

CNN on <i>MNIST</i>	CNN on <i>F-MNIST</i>	CNN on <i>CIFAR-10</i>
28×28 Gray Image	28×28 Gray Image	32×32 RGB Image
	3×3 conv, 128 LReLU	
	3×3 conv, 128 LReLU	
	3×3 conv, 128 LReLU	
	2×2 max-pool	
	dropout, $p = 0.25$	
	3×3 conv, 256 LReLU	
	3×3 conv, 256 LReLU	
	3×3 conv, 256 LReLU	
	2×2 max-pool	
	dropout, $p = 0.25$	
	3×3 conv, 512 LReLU	
	3×3 conv, 256 LReLU	
	3×3 conv, 128 LReLU	
	avg-pool	
dense 128→10	dense 128→10	dense 128→10

Table 8: CNN on *CIFAR-100*.

CNN on <i>CIFAR-100</i>
32×32 RGB Image
3×3 conv, 64 ReLU
3×3 conv, 64 ReLU
2×2 max-pool
3×3 conv, 128 ReLU
3×3 conv, 128 ReLU
2×2 max-pool
3×3 conv, 196 ReLU
3×3 conv, 196 ReLU
2×2 max-pool
dense 256→100

References

- [1] Olivier Catoni. Challenging the empirical mean and empirical variance: a deviation study. In *Annales de l'IHP Probabilités et statistiques*, volume 48, pages 1148–1185, 2012.
- [2] Peng Chen, Xinghu Jin, Xiang Li, and Lihu Xu. A generalized catoni’s m -estimator under finite α -th moment assumption with $\alpha \in (1, 2)$. *arXiv preprint arXiv:2010.05008*, 2020.
- [3] Ilias Diakonikolas and Daniel M Kane. Recent advances in algorithmic high-dimensional robust statistics. *arXiv preprint arXiv:1911.05911*, 2019.
- [4] Ilias Diakonikolas, Daniel M Kane, and Ankit Pensia. Outlier robust mean estimation with subgaussian rates via stability. *arXiv preprint arXiv:2007.15618*, 2020.
- [5] Evarist Giné, Rafał Łatała, and Joel Zinn. Exponential and moment inequalities for u -statistics. In *High Dimensional Probability II*, pages 13–38. Springer, 2000.
- [6] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, pages 8527–8537, 2018.
- [7] Liu Liu, Tianyang Li, and Constantine Caramanis. High dimensional robust m -estimation: Arbitrary corruption and heavy tails. *arXiv preprint arXiv:1901.08237*, 2019.

- 126 [8] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey.
127 Normalized loss functions for deep learning with noisy labels. In *ICML*, pages 6543–6553,
128 2020.
- 129 [9] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*.
130 MIT Press, 2018.
- 131 [10] Laura Niss and Ambuj Tewari. What you see may not be what you get: Ucb bandit algorithms
132 robust to ϵ -contamination. In *UAI*, pages 450–459, 2020.
- 133 [11] Daniel Paulin et al. Concentration inequalities for markov chains by marton couplings and
134 spectral methods. *Electronic Journal of Probability*, 20, 2015.
- 135 [12] Gareth O Roberts, Jeffrey S Rosenthal, et al. General state space markov chains and mcmc
136 algorithms. *Probability surveys*, 1:20–71, 2004.
- 137 [13] Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Mingming Gong, Haifeng Liu, Gang Niu,
138 Dacheng Tao, and Masashi Sugiyama. Part-dependent label noise: Towards instance-dependent
139 label noise. In *NeurIPS*, 2020.
- 140 [14] Xiaobo Xia, Tongliang Liu, Bo Han, Chen Gong, Nannan Wang, Zongyuan Ge, and Yi Chang.
141 Robust early-learning: Hindering the memorization of noisy labels. In *ICLR*, 2021.
- 142 [15] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations
143 in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.
- 144 [16] Quanming Yao, Hansi Yang, Bo Han, Gang Niu, and James Tin-Yau Kwok. Searching to exploit
145 memorization effect in learning with noisy labels. In *ICML*, pages 10789–10798, 2020.
- 146 [17] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W Tsang, and Masashi Sugiyama. How
147 does disagreement benefit co-teaching? In *ICML*, 2019.