

A APPENDIX

A.1 RoCL_{Base} (NO CURRICULUM) IN SECTION 2 AND FIGURE 2, 10

Algorithm 2 RoCL_{Base} (no curriculum)

```

1: input:  $\{(x_i, y_i)\}_{i=1}^n, h(\cdot; \eta), \ell(\cdot, \cdot), f(\cdot; \theta), T_{0:K}; \gamma \in [0, 1]$ 
2: initialize:  $\theta_0, l_0(i) = c_0(i) = 0 \forall i \in [n], T_{-1} = 0$ 
3: for  $k \in \{0, \dots, K\}$  do
4:   for  $t' \in \{1, \dots, T_k\}$  do
5:      $t \leftarrow t' + T_{k-1}$ ;
6:      $S_t \leftarrow [n]$ ;
7:     if  $k \% 2 = 0$  then
8:        $\theta_t \leftarrow \theta_{t-1} + h(\nabla_{\theta} \frac{1}{n} \sum_{i=0}^n \ell_t(i); \eta)$ ; {supervised learning using given labels}
9:       Update  $l_{t+1}(i)$  by Eq. (1); {update EMA loss}
10:    else
11:       $\theta_t \leftarrow \theta_{t-1} + h(\nabla_{\theta} \frac{1}{n} \sum_{i=0}^n \zeta_t(i); \eta)$ ; {self-supervised learning using pseudo labels}
12:      Update  $c_{t+1}(i)$  by Eq. (3); {update EMA consistency loss}
13:    end if
14:  end for
15: end for

```

Note the EMA metrics in line 9 and line 12 are not used for training in RoCL_{Base} . They have been updated and recorded for the purpose of empirical study presented in Section 2

A.2 ADDITIONAL EXPERIMENTS




Figure 2: RoCL (Algorithm 1) vs. RoCL_{Base} without any curriculum (Algorithm 2 in Appendix) during the training of ResNet34 on CIFAR10 containing 60% symmetric noises on labels.




Figure 3: RoCL (Algorithm 1) vs. RoCL_{Base} without any curriculum (Algorithm 2 in Appendix) during the training of ResNet34 on CIFAR10 containing 80% symmetric noises on labels.




Figure 4: RoCL (Algorithm 1) vs. RoCL_{Base} without any curriculum (Algorithm 2 in Appendix) during the training of ResNet34 on CIFAR100 containing 60% symmetric noises on labels.




Figure 5: RoCL (Algorithm 1) vs. RoCL_{Base} without any curriculum (Algorithm 2 in Appendix) during the training of ResNet34 on CIFAR100 containing 80% symmetric noises on labels.

Table 6: Extended version of Table 3 with two more baselines: NFL+RCE and NCE+MAE.

Dataset	CIFAR10			CIFAR100		
Noise Rate	40%	60%	80%	40%	60%	80%
MD-DYR-SH	92.3	86.1	74.1	70.1	59.5	39.5
MentorNet	91.2	74.2	60.0	68.5	61.2	35.5
MentorMix	94.2	91.3	81.0	71.3	64.6	41.2
O2U-net	90.3	-	43.4	69.2	-	39.4
RoG+D2L	87.0	78.0	-	64.9	40.6	-
PENCIL	-	-	-	69.12 ± 0.62	57.79 ± 3.86	fail
GCE	87.62 ± 0.26	82.70 ± 0.23	67.92 ± 0.60	62.64 ± 0.33	54.04 ± 0.56	29.60 ± 0.51
SCE	85.34 ± 0.07	80.07 ± 0.02	53.81 ± 0.27	53.69 ± 0.07	41.47 ± 0.04	15.00 ± 0.04
NFL+MAE	83.81 ± 0.06	76.36 ± 0.31	45.23 ± 0.52	58.18 ± 0.08	46.10 ± 0.50	24.78 ± 0.82
NFL+RCE	86.05 ± 0.12	79.78 ± 0.13	55.06 ± 1.08	58.20 ± 0.31	46.30 ± 0.45	25.16 ± 0.55
NCE+MAE	84.19 ± 0.43	77.61 ± 0.05	49.62 ± 0.72	59.22 ± 0.36	48.06 ± 0.34	25.50 ± 0.76
NCE+RCE	86.02 ± 0.09	79.78 ± 0.50	52.71 ± 1.90	59.48 ± 0.56	47.12 ± 0.62	25.80 ± 1.12
RoCL (ours)	94.55 ± 0.12	92.06 ± 0.23	85.76 ± 0.26	74.64 ± 0.43	66.79 ± 0.58	47.24 ± 0.75

Table 7: Extended version of Table 4 with two more baselines: NFL+RCE and NCE+MAE.

Dataset	CIFAR10			CIFAR100		
Noise Rate	20%	30%	40%	20%	30%	40%
PENCIL	92.43	91.84	91.01	74.70 ± 0.56	72.52 ± 0.38	63.61 ± 0.23
Bootstrap	86.57 ± 0.08	84.86 ± 0.05	79.76 ± 0.07	63.44 ± 0.35	63.18 ± 0.35	62.08 ± 0.22
F-correct	89.09 ± 0.47	86.79 ± 0.36	83.55 ± 0.58	42.46 ± 2.16	38.13 ± 2.97	34.44 ± 1.93
GCE	86.07 ± 0.31	80.78 ± 0.21	74.98 ± 0.32	59.99 ± 0.83	53.99 ± 0.29	41.49 ± 0.79
SCE	83.92 ± 0.07	79.70 ± 0.27	78.20 ± 0.03	58.22 ± 0.47	49.85 ± 0.91	42.19 ± 0.19
NFL+MAE	86.81 ± 0.32	83.91 ± 0.34	77.16 ± 0.10	63.10 ± 0.22	56.19 ± 0.61	43.51 ± 0.42
NFL+RCE	88.73 ± 0.29	85.74 ± 0.22	79.27 ± 0.43	63.12 ± 0.41	54.72 ± 0.38	42.97 ± 1.03
NCE+MAE	86.44 ± 0.23	83.98 ± 0.52	78.23 ± 0.42	62.38 ± 0.60	58.02 ± 0.48	47.22 ± 0.30
NCE+RCE	88.56 ± 0.17	85.58 ± 0.44	79.59 ± 0.40	62.68 ± 0.79	57.82 ± 0.41	46.79 ± 0.96
RoCL (ours)	95.38 ± 0.21	94.19 ± 0.28	92.31 ± 0.35	80.03 ± 0.34	77.59 ± 0.45	73.28 ± 0.83




Figure 6: Illustration of Eq. (9) and visualization of our choice for $g(\cdot)$ and the resulted τ_t when $T = 50$. We use $g(\cdot) = \tanh(\cdot)$ (which can be other “S”-shape functions) and $\sigma = 0.95$ in our experiments. Here, we map the points on the black curve in the left plot to the points on the red curve in the right plot. Each gray point on the bottom of the left plot is from the T evenly spaced x-coordinates between the x-interval $[g^{-1}(-\sigma), g^{-1}(\sigma)]$. We scale them to the T t-coordinates in the bottom of the right plot (i.e., $t = 1, 2, \dots, 50$), which associates with T τ_t values represented by the red points between $[\tau_1, \tau_T]$ on the red curve.




Figure 7: **Ablation study:** RoCL vs. its variants during the training of ResNet34 on **CIFAR10** containing **60% symmetric noises** on labels.




Figure 8: **Ablation study:** RoCL vs. its variants during the training of ResNet34 on **CIFAR10** containing **80% symmetric noises** on labels.




Figure 9: **Ablation study:** RoCL vs. its variants during the training of ResNet34 on **CIFAR100** containing **60% symmetric noises** on labels.




Figure 10: **Ablation study:** RoCL vs. its variants during the training of ResNet34 on **CIFAR100** containing **80% symmetric noises** on labels.

We present a more detailed analysis of the ablation study results with explanations of the observed phenomenons below.

- Most variants (except RoCL_{Base} , no RandAugment, and no ClassBalance) have similar performance as the original RoCL and perform better than or competitive with the SoTA results achieved by MentorMix. The differences compared to original RoCL become smaller under the lower noise rate setting (60%). RoCL_{Base} uses all data for training in each step without applying any curriculum, showing that our proposed curriculum is the most critical component of RoCL in achieving the appealing improvements. Note RoCL_{Base} already outperforms most methods in Table 3, which verifies the effectiveness of multi-episode training that alternates between supervised learning with the given labels and self-supervision with the pseudo labels.
- Removing RandAugment degrades the performance, especially when the noise rate is very high (e.g., 80%) because strong data augmentations are required by the self-supervision and the EMA consistency loss in RoCL, while trivial data augmentations can result in error accumulation or over-confidence in pseudo labels and inaccurate EMA consistency loss. The self-supervision aims to encourage the model output consistency over different augmentations of the same sample. Without augmentations with sufficient variations, self-supervision reduces to reinforcing the same outputs on similar samples and thus carries little information and can even magnify/accumulate errors (if any) in the original outputs. Also, the EMA consistency loss cannot generate meaningful consistency measures if computed on the same data or its trivial augmentations. Note a strong data augmentation is not always beneficial in all noisy label learning methods since it can increase the uncertainty in the presence of wrong labels, making the detection of clean data and noise correction more challenging. For example, we tried applying RandAugment to MentorMix (using

the official implementations of both) but observed inferior performance compared to the results using its original data augmentations.

- Class balance regularization is useful for the very high noise rate setting (80%), in which a wrong label may dominate the learning on a mini-batch by a large chance. However, when the noise rate is not that high (e.g., 60% on CIFAR10), removing it results in better performance.
- Although Mix-Up has been proved effective in previous methods, and for this reason, we followed MentorMix by starting with a relatively strong Mix-Up ($\alpha = 8.0$) and then gradually reducing it to $\alpha = 0.2$. In the ablation study, we find that completely removing Mix-Up significantly improves performance. Mix-Up is helpful when applied to mix a clean label with a noisy label since the latter can be mediated with the former and thus softened. However, this is rarely the case for RoCL since RoCL either mainly learns from clean data or wrongly-labeled data with correct pseudo labels, and the transition between the two phases is short. When applied to two correct labels/pseudo labels, Mix-Up weakens each label’s confidence, and we may lose information from the inter-class probabilities in the soft pseudo labels.
- Replacing weighted sampling with top-k selection (“no RandSampling”) or replacing EMA metrics with instantaneous metrics (“no EMA metrics”) causes less degeneration on the final test accuracies. However, they are important to the early-stage exploration and accurate estimation of EMA metrics on less-visited samples. In Figure 7-10, these two variants usually suffer from low accuracy and convergence speed during early stages. The only exception is “no RandSampling” in Figure 10, which performs better than the original RoCL. A possible reason is that the randomness brought by high uniform label noises already bring sufficient randomness for exploration.
- Replacing $p_t(i)$, $q_t(i)$ or both with uniform probabilities over all samples reduces the final test accuracies in all cases, e.g., the degradation is significant on CIFAR100 with 80% noise. In Figure 7-10, we can see that by setting $q_t(i) = 1/n$ results in less degradation than the other two. This is due to the more accurate pseudo labels generated for more data (even the ones with larger EMA consistency loss) as training proceeds. Moreover, since we are conservative in setting λ_T and τ_T , the performance is not very sensitive to wrong pseudo labels.