We appreciate positive and constructive comments, and address the main concerns raised by the reviewers below.

Effects of K' and M [R4] Figure 1 shows the effects of K' and M on CUB-200 with open-set uniform noise ($\eta = 0.5$). Our models outperform the baseline even with a fairly small number of meta-class sets regardless of K'. In particular, the model with $K' = 2$ works best among the ones with the same number of parameters, which are depicted by a diamond marker on each line plot, since it reduces noise level most effectively. We observe the same tendency in the experiments with other noise settings.

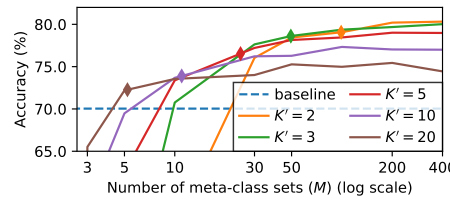

Figure 1: Accuracy by varying K' and M.

Relation to ECOC [R4] Unlike ECOC that constructs codewords deterministically, it is not straightforward to show a theoretical guarantee of our method based on an iterative clustering algorithm in a high-dimensional latent feature space. However, our approach increases the number of partitions exponentially by simply adding meta-class sets, which makes a small number of meta-class sets ($M = 20$) large enough to decipher all 200 classes on CUB-200 through a combination of binary meta-class sets ($K = 2$). It is a valuable suggestion to adopt the intuition behind ECOC for meta-class set configuration. It effectively decreases the number of meta-class sets to 15 from 20 for deciphering in practice. Moreover, when we employ the column separation idea proposed in the ECOC paper and reduce the correlation between meta-class sets, our approach shows additional accuracy gains by approximately 1% point in average on CUB-200.

Optimality of meta-class sets [R5] The ideal meta-class sets should be compact and uncorrelated, and these properties affect the final accuracy of our algorithm. Our reward function (Eq. (6) of the main paper) guides to identify the meta-class sets with such desirable characteristics while there is no direct supervision for correlation between meta-class sets. However, it turns out that the meta-class sets given by our RL-based method are significantly less correlated compared to the random selection. We will consider a more direct feedback based on the measured correlation while a simple approach has already been discussed in "Relation to ECOC [R5]". Thanks for the insightful comment.

Issues related to RL [R1] Our meta-class set search algorithm relies on a noisy validation set held out from the training set, which is also adopted in the experiments on WebVision. The search cost by RL is just as much as the cost for training a classifier. To reduce the computational complexity, we employ a simple two-layer perceptron and optimize the agent based on the in-batch validation accuracy instead of computing the accuracy on the entire validation set.

Ablation study on meta-class set construction methods [R1] As shown in Table 1, the combinatorial classifier outperforms the baseline (Standard) even with randomly constructed meta-class sets (Random). Our clustering-based algorithm (Clustering) brings additional improvement especially on the datasets with high noise level while employing the meta-class set search technique (Clustering+RL_Search) boosts the accuracy even further.

Table 1: Accuracies [%] of baseline and proposed models with different meta-class set configurations on CUB-200.

	Clean	Open-Uniform		Open-Nearest		Clean	Closed-Uniform		Closed-Nearest	
	$\eta = 0$	$\eta = 0.25$	$\eta = 0.50$	$\eta = 0.25$	$\eta = 0.50$	$\eta = 0$	$\eta = 0.25$	$\eta = 0.50$	$\eta = 0.25$	$\eta = 0.50$
Standard	80.57	73.37	70.04	77.14	75.45	79.58	63.65	42.35	65.21	47.70
Random	82.05	**80.05**	78.66	78.74	76.75	79.82	69.36	46.98	67.12	48.55
Clustering	82.80	79.38	79.19	79.28	77.95	81.36	71.75	51.90	68.35	52.00
Clustering+RL_Search	**83.12**	79.72	**79.98**	**79.36**	**78.35**	**81.62**	**72.43**	**54.52**	**68.79**	**52.43**

Use of pretrained ImageNet [R1] Although we partly agree to the concern about the use of the pretrained network on ImageNet, we also believe that the pretrained model is a commodity and is widely used for the initialization of many CNNs. Also, our experiment on WebVision does not rely on the pretrained model but the CNNs trained from scratch; it illustrates the representation learning capability of the proposed method.

Results on CIFAR-100 [R1] Table 2 presents the results from the methods tested on CIFAR-100 with the ResNet-50 backbone model, where the proposed approach (CombCls) achieves the largest accuracy gains with respect to the baseline (Standard) and a combination with Co-teaching (CombCls+) improves accuracies further.

Table 2: Results on CIFAR-100 with open-set noise ($\eta = 0.25$).

	Uniform	Nearest
Standard	73.51	73.88
q-loss	73.91	74.41
Co-teaching	75.84	76.30
CombCls	76.30	76.69
CombCls+	**78.57**	**78.39**

Qualitative examples in meta-class sets [R1] The binary separations of input images given by meta-class sets look reasonable; each meta-class captures either one or more attributes, related classes, or common visual appearances, which are crucial cues to identify the fine-grained categories. It is difficult to include specific cases in the rebuttal, but we will present interesting examples in the supplementary file if our paper is accepted.

Final vs. best accuracy [R1] The results from all algorithms are given by the models identified using a clean validation set. This is because the learning curves of individual methods may be different and reporting accuracies at a particular epoch may be unfair. We also note that our approach still outperforms others even when we fix the number of epochs in a wide range for comparisons. We will present more detailed results if our paper is accepted.

Others [All] We will supplement the missing details and results in the final manuscript if our paper is accepted.