

---

# Confident Multiple Choice Learning

---

Kimin Lee<sup>1</sup> Changho Hwang<sup>1</sup> KyoungSoo Park<sup>1</sup> Jinwoo Shin<sup>1</sup>

## Abstract

Ensemble methods are arguably the most trustworthy techniques for boosting the performance of machine learning models. Popular independent ensembles (IE) relying on naïve averaging/voting scheme have been of typical choice for most applications involving deep neural networks, but they do not consider advanced collaboration among ensemble models. In this paper, we propose new ensemble methods specialized for deep neural networks, called confident multiple choice learning (CMCL): it is a variant of multiple choice learning (MCL) via addressing its overconfidence issue. In particular, the proposed major components of CMCL beyond the original MCL scheme are (i) new loss, i.e., confident oracle loss, (ii) new architecture, i.e., feature sharing and (iii) new training method, i.e., stochastic labeling. We demonstrate the effect of CMCL via experiments on the image classification on CIFAR and SVHN, and the foreground-background segmentation on the iCoseg. In particular, CMCL using 5 residual networks provides 14.05% and 6.60% relative reductions in the top-1 error rates from the corresponding IE scheme for the classification task on CIFAR and SVHN, respectively.

## 1. Introduction

Ensemble methods have played a critical role in the machine learning community to obtain better predictive performance than what could be obtained from any of the constituent learning models alone, e.g., Bayesian model/parameter averaging (Domingos, 2000), boosting (Freund et al., 1999) and bagging (Breiman, 1996). Recently, they have been successfully applied to enhancing the power of many deep neural networks, e.g., 80% of

top-5 best-performing teams on ILSVRC challenge 2016 (Krizhevsky et al., 2012) employ ensemble methods. They are easy and trustworthy to apply for most scenarios. While there exists a long history on ensemble methods, the progress on developing more advanced ensembles specialized for deep neural networks has been slow. Despite continued efforts that apply various ensemble methods such as bagging and boosting to deep models, it has been observed that traditional independent ensembles (IE) which train models independently with random initialization achieve the best performance (Ciregan et al., 2012; Lee et al., 2015). In this paper, we focus on developing more advanced ensembles for deep models utilizing the concept of multiple choice learning (MCL).

The MCL concept was originally proposed in (Guzman-Rivera et al., 2012) under the scenario when inference procedures are cascaded:

- (a) First, generate a set of plausible outputs.
- (b) Then, pick the correct solution form the set.

For example, (Park & Ramanan, 2011; Batra et al., 2012) proposed human-pose estimation methods which produce multiple predictions and then refine them by employing a temporal model, and (Collins & Koo, 2005) proposed a sentence parsing method which re-ranks the output of an initial system which produces a set of plausible outputs (Huang & Chiang, 2005). In such scenarios, the goal of the first stage (a) is generating a set of plausible outputs such that at least one of them is correct for the second stage (b), e.g., human operators. Under this motivation, MCL has been studied (Guzman-Rivera et al., 2014; 2012; Lee et al., 2016), where various applications have been demonstrated, e.g., image classification (Krizhevsky & Hinton, 2009), semantic segmentation (Everingham et al., 2010) and image captioning (Lin et al., 2014b). It trains an ensemble of multiple models by minimizing the so-called *oracle loss*, only focusing on the most accurate prediction produced by them. Consequently, it makes each model specialized for a certain subset of data, not for the entire one similarly as mixture-of-expert schemes (Jacobs et al., 1991).

Although MCL focuses on the first stage (a) in cascaded scenarios and thus can produce diverse/plausible outputs, it might be not useful if one does not have a good scheme for

---

<sup>1</sup>School of Electrical Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea. Correspondence to: Jinwoo Shin <jinwoos@kaist.ac.kr>.

the second stage (b). One can use a certain average/voting scheme of the predictions made by models for (b), but MCL using deep neural networks often fails to make a correct decision since each network tends to be overconfident in its prediction. Namely, the oracle error/loss of MCL is low, but its *top-1* error rate might be very high.

**Contribution.** To address the issue, we develop the concept of confident MCL (CMCL) that does not lose any benefit of the original MCL, while its target loss and architecture are redesigned for making the second stage (b) easier. Specifically, it targets to generate a set of diverse/plausible confident predictions from which one can pick the correct one using a simple average/voting scheme. To this end, we first propose a new loss function, called *confident oracle loss*, for relaxing the overconfidence issue of MCL. Our key idea is to additionally minimize the Kullback-Leibler divergence from a predictive distribution to the uniform one in order to give confidence to non-specialized models. Then, CMCL that minimizes the new loss can be efficiently trained like the original MCL for certain classes of models including neural networks, via stochastic alternating minimization (Lee et al., 2016). Furthermore, when CMCL is applied to deep models, we propose two additional regularization techniques for boosting its performance: *feature sharing* and *stochastic labeling*. Despite the new components, we note that the training complexity of CMCL is almost same to that of MCL or IE.

We apply the new ensemble model trained by the new training scheme for several convolutional neural networks (CNNs) including VGGNet (Simonyan & Zisserman, 2015), GoogLeNet (Szegedy et al., 2015), and ResNet (He et al., 2016) for image classification on the CIFAR (Krizhevsky & Hinton, 2009) and SVHN (Netzer et al., 2011) datasets, and fully-convolutional neural networks (FCNs) (Long et al., 2015) for foreground-background segmentation on the iCoseg dataset (Batra et al., 2010). First, for the image classification task, CMCL outperforms all baselines, i.e., the traditional IE and the original MCL, in top-1 error rates. In particular, CMCL of 5 ResNet with 20 layers provides 14.05% and 6.60% relative reductions in the top-1 error rates from the corresponding IE on CIFAR-10 and SVHN, respectively. Second, for the foreground-background segmentation task, CMCL using multiple FCNs with 4 layers also outperforms all baselines in top-1 error rates. Each model trained by CMCL generates high-quality solutions by specializing for specific images while each model trained by IE does not. We believe that our new approach should be of broader interest for many deep learning tasks requiring high accuracy.

**Organization.** In Section 2, we introduce necessary backgrounds for multiple choice learning and the corresponding loss function. We describe the proposed loss and the corre-

sponding training scheme in Section 3. Section 4 provides additional techniques for the proposed ensemble model. Experimental results are reported in Section 5.

## 2. Preliminaries

### 2.1. Multiple Choice Learning

In this section, we describe the basic concept of multiple choice learning (MCL) (Guzman-Rivera et al., 2014; 2012). Throughout this paper, we denote the set  $\{1, \dots, n\}$  by  $[n]$  for positive integer  $n$ . The MCL scheme is a type of ensemble learning that produces diverse outputs of high quality. Formally, given a training dataset  $\mathcal{D} = \{(\mathbf{x}_i, y_i) \mid i \in [N], \mathbf{x}_i \in \mathcal{X}, y_i \in \mathcal{Y}\}$ , we consider an ensemble of  $M$  models  $f$ , i.e.,  $(f_1, \dots, f_M)$ . For some task-specific loss function  $\ell(y, f(\mathbf{x}))$ , the oracle loss over the dataset  $\mathcal{D}$  is defined as follows:

$$L_O(\mathcal{D}) = \sum_{i=1}^N \min_{m \in [M]} \ell(y_i, f_m(\mathbf{x}_i)), \quad (1)$$

while the traditional independent ensemble (IE) loss is

$$L_E(\mathcal{D}) = \sum_{i=1}^N \sum_{m \in [M]} \ell(y_i, f_m(\mathbf{x}_i)). \quad (2)$$

If all models have the same capacity and one can obtain the (global) optimum of the IE loss with respect to the model parameters, then all trained models should produce the same outputs, i.e.,  $f_1 = \dots = f_M$ . On the other hand, the oracle loss makes the most accurate model optimize the loss function  $\ell(y, f(\mathbf{x}))$  for each data  $\mathbf{x}$ . Therefore, MCL produces diverse outputs of high quality by forcing each model to be specialized on a part of the entire dataset.

Minimizing the oracle loss (1) is harder than minimizing the independent ensemble loss (2) since the min function is a non-continuous function. To address the issue, (Guzman-Rivera et al., 2012) proposed an iterative block coordinate decent algorithm and (Dey et al., 2015) reformulated this problem as a submodular optimization task in which ensemble models are trained sequentially in a boosting-like manner. However, when one considers an ensemble of deep neural networks, it is challenging to apply these methods since they require either costly retraining or sequential training. Recently, (Lee et al., 2016) overcame this issue by proposing a stochastic gradient descent (SGD) based algorithm. Throughout this paper, we primarily focus on ensembles of deep neural networks and use the SGD algorithm for optimizing the oracle loss (1) or its variants.

### 2.2. Oracle Loss for Top-1 Choice

The oracle loss (1) used for MCL is useful for producing diverse/plausible outputs, but it is often inappropriate for ap-

	(a) Multiple choice learning (MCL)					(b) Confident MCL (CMCL)					(c) Independent ensemble (IE)				
Airplane	0.0%	0.0%	93.6%	0.0%	0.0%	95.8%	0.0%	4.4%	12.2%	2.2%	86.6%	85.5%	86.4%	85.7%	86.0%
Automobile	0.0%	0.0%	96.1%	0.0%	0.0%	0.0%	0.0%	0.8%	98.6%	9.0%	90.7%	90.3%	90.5%	90.6%	90.5%
Bird	99.9%	0.0%	0.0%	0.0%	0.0%	0.1%	0.3%	2.4%	4.1%	94.0%	75.4%	75.9%	74.5%	76.3%	76.5%
Cat	0.0%	0.0%	95.6%	0.0%	0.0%	94.5%	2.6%	0.0%	0.0%	0.2%	68.5%	66.5%	66.1%	67.1%	67.1%
Deer	0.0%	0.0%	0.0%	97.5%	0.0%	0.0%	23.6%	1.2%	98.7%	4.5%	85.8%	86.3%	86.1%	86.1%	86.2%
Dog	0.0%	97.0%	0.0%	0.0%	0.0%	15.8%	8.0%	2.9%	4.7%	91.7%	76.3%	75.6%	77.5%	75.0%	76.5%
Frog	0.0%	0.0%	0.0%	0.0%	97.7%	7.1%	0.9%	99.2%	2.7%	0.0%	90.1%	90.7%	90.3%	91.4%	90.6%
Horse	0.0%	0.0%	0.0%	0.0%	97.2%	0.0%	0.0%	98.1%	0.0%	0.0%	87.3%	86.9%	86.6%	86.3%	87.2%
Ship	0.0%	0.0%	0.0%	97.2%	0.0%	0.0%	97.3%	0.0%	0.0%	0.0%	91.6%	91.6%	91.4%	91.7%	90.7%
Truck	0.0%	97.4%	0.0%	0.0%	0.0%	0.5%	96.1%	0.0%	0.0%	28.0%	90.4%	89.3%	89.8%	90.0%	90.0%
	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5

Figure 1. Class-wise test set accuracy of each ensemble model trained by various ensemble methods. One can observe that most models trained by MCL and CMCL become specialists for certain classes while they are generalized in case of traditional IE.

plications requiring a single choice, i.e., top-1 error. This is because ensembles of deep neural networks tend to be overconfident in their predictions, and it is hard to judge a better solution from their outputs. To explain this in more detail, we evaluate the performance of ensembles of convolutional neural networks (CNNs) for the image classification task on the CIFAR-10 dataset (Krizhevsky & Hinton, 2009). We train ensembles of 5 CNNs (two convolutional layers followed by a fully-connected layer) using MCL. We also train the models using traditional IE which trains each model independently under different random initializations. Figure 1 summarizes the class-wise test set accuracy of each ensemble member. In the case of MCL, most models become specialists for certain classes (see Figure 1(a)), while they are generalized in the case of traditional IE as shown in Figure 1(c). However, as expected, each model trained by MCL significantly outperforms for its specialized classes than that trained by IE. For choosing a single output, similar to (Wan et al., 2013; Ciregan et al., 2012), one can average the output probabilities from ensemble members trained by MCL, but the corresponding top-1 classification error rate is often very high (e.g., see Table 1 in Section 5). This is because each model trained by MCL is overconfident for its non-specialized classes. To quantify this, we also compute the entropy of the predictive distribution on the test data and use this to evaluate the quality of confidence/uncertainty level. Figure 2(a) reports the entropy extracted from the predictive distribution of one of ensemble models trained by MCL. One can observe that it has low entropy as expected for its specialized classes (i.e., classes that the model has a test accuracy higher than 90%). However, even for non-specialized classes, it also has low entropy. Due to this, with respect to top-1 error rates, simple averaging of models trained by MCL performs much worse than that of IE. Such issue typically occurs in deep neural networks since it is well known that they are poor at quantifying predictive uncertainties, and tend to be easily overconfident (Nguyen et al., 2015).

### 3. Confident Multiple Choice Learning

#### 3.1. Confident Oracle Loss

In this section, we propose a modified oracle loss for relaxing the issue of MCL described in the previous section. Suppose that the  $m$ -th model outputs the predictive distribution  $P_{\theta_m}(y | \mathbf{x})$  given input  $\mathbf{x}$ , where  $\theta_m$  denotes the model parameters. Then, we define the confident oracle loss as the following integer programming variant of (1):

$$L_C(\mathcal{D}) = \min_{v_i^m} \sum_{i=1}^N \sum_{m=1}^M \left( v_i^m \ell(y_i, P_{\theta_m}(y | \mathbf{x}_i)) + \beta (1 - v_i^m) D_{KL}(\mathcal{U}(y) \| P_{\theta_m}(y | \mathbf{x}_i)) \right) \quad (3a)$$

$$\text{subject to} \quad \sum_{m=1}^M v_i^m = 1, \quad \forall i, \quad (3b)$$

$$v_i^m \in \{0, 1\}, \quad \forall i, m \quad (3c)$$

where  $D_{KL}$  denotes the Kullback-Leibler (KL) divergence,  $\mathcal{U}(y)$  is the uniform distribution,  $\beta$  is a penalty parameter, and  $v_i^m$  is a flag variable to decide the assignment of  $\mathbf{x}_i$  to the  $m$ -th model. By minimizing the KL divergence from the predictive distribution to the uniform one, the new loss forces the predictive distribution to be closer to the uniform one, i.e., zero confidence, on non-specialized data, while those for specialized data still follow the correct one. For example, for classification tasks, the most accurate model for each data is allowed to optimize the classification loss, while others are forced to give less confident predictions by minimizing the KL divergence. We remark that although we optimize the KL divergence only for non-specialized data, one can also do it even for specialized data to regularize each model (Pereyra et al., 2017).

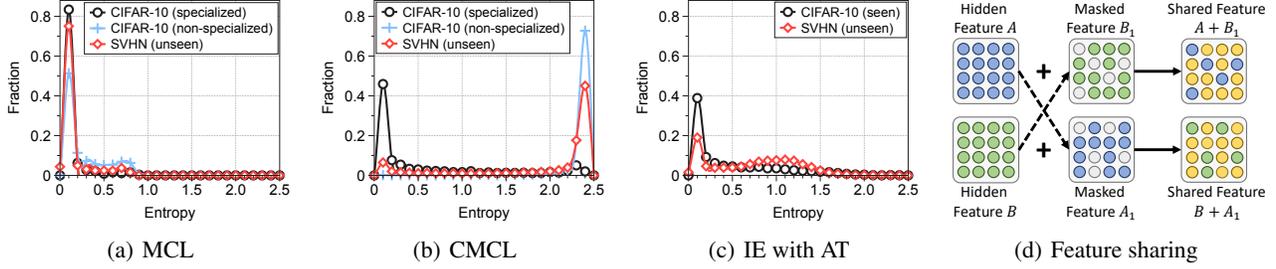


Figure 2. Histogram of the predictive entropy of model trained by (a) MCL (b) CMCL and (c) IE on CIFAR-10 and SVHN test data. In the case of MCL and CMCL, we separate the classes of CIFAR-10 into specialized (i.e., classes that model has a class-wise test accuracy higher than 90%) and non-specialized (others) classes. In the case of IE, we follow the proposed method by (Lakshminarayanan et al., 2016): train an ensemble of 5 models with adversarial training (AT) and measure the entropy using the averaged probability, i.e., averaging output probabilities from 5 models. (d) Detailed view of feature sharing between two models. Grey units indicate that they are currently dropped. Masked features passed to a model are all added to generate the shared features.

### 3.2. Stochastic Alternating Minimization for Training

In order to minimize the confident oracle loss (3) efficiently, we use the following procedure (Guzman-Rivera et al., 2012), which optimizes model parameters  $\{\theta_m\}$  and assignment variables  $\{v_i^m\}$  alternatively:

1. **Fix  $\{\theta_m\}$  and optimize  $\{v_i^m\}$ .**

Under fixed model parameters  $\{\theta_m\}$ , the objective (3a) is decomposable with respect to assignments  $\{v_i^m\}$  and it is easy to find optimal  $\{v_i^m\}$ .

2. **Fix  $\{v_i^m\}$  and optimize  $\{\theta_m\}$ .**

Under fixed assignments  $\{v_i^m\}$ , the objective (3a) is decomposable with respect to model parameters  $\{\theta_m\}$ , and it requires each model to be trained independently.

The above scheme iteratively assigns each data to a particular model and then independently trains each model only using its assigned data. Even though it monotonically decreases the objective, it is still highly inefficient since it requires training each model multiple times until assignments  $\{v_i^m\}$  converge. To address the issue, we propose deciding assignments and update model parameters to the gradient directions once per each batch, similarly to (Lee et al., 2016). In other words, we perform a single gradient-update on parameters in Step 2, without waiting for their convergence to a (local) optimum. In fact, (Lee et al., 2016) show that such stochastic alternating minimization works well for the oracle loss (1). We formally describe a detailed training procedure as the ‘version 0’ of Algorithm 1, and we will introduce the alternative ‘version 1’ later. This direction is complementary to ours, and we do not explore in this paper.

---

#### Algorithm 1 Confident MCL (CMCL)

---

**Input:** Dataset  $\mathcal{D} = \{(\mathbf{x}_i, y_i) \mid \mathbf{x}_i \in \mathcal{X}, y_i \in \mathcal{Y}\}$  and penalty parameter  $\beta$

**Output:** Ensemble of  $M$  trained models

---

**repeat**

Let  $\mathcal{U}(y)$  be a uniform distribution

Sample random batch  $\mathcal{B} \subset \mathcal{D}$

**for**  $m = 1$  **to**  $M$  **do**

Compute the loss of the  $m$ -th model:

$$L_i^m \leftarrow \beta \sum_{\hat{m} \neq m} D_{KL}(\mathcal{U}(y) \parallel P_{\theta_{\hat{m}}}(y \mid \mathbf{x}_i)) + \ell(y_i, P_{\theta_m}(y_i \mid \mathbf{x}_i)), \quad \forall (\mathbf{x}_i, y_i) \in \mathcal{B}$$

**end for**

**for**  $m = 1$  **to**  $M$  **do**

**for**  $i = 1$  **to**  $|\mathcal{B}|$  **do**

**if** the  $m$ -th model has the lowest loss **then**

Compute the gradient of the training loss  $\ell(y_i, P_{\theta_m}(y_i \mid \mathbf{x}_i))$  w.r.t  $\theta_m$

**else**

/\* version 0: exact gradient \*/

Compute the gradient of the KL divergence  $\beta D_{KL}(\mathcal{U}(y) \parallel P_{\theta_m}(y \mid \mathbf{x}_i))$  w.r.t  $\theta_m$

/\* version 1: stochastic labeling \*/

Compute the gradient of the cross entropy loss  $-\beta \log P_{\theta_m}(\hat{y}_i \mid \mathbf{x}_i)$  using  $\hat{y}_i$  w.r.t  $\theta_m$  where  $\hat{y}_i \sim \mathcal{U}(y)$

**end if**

**end for**

Update the model parameters

**end for**

**until** convergence

---

### 3.3. Effect of Confident Oracle Loss

Similar to Section 2.2, we evaluate the performance of the proposed training scheme using 5 CNNs for image classification on the CIFAR-10 dataset. As shown in Figure 1(b), ensemble models trained by CMCL using the exact gradient (i.e., version 0 of Algorithm 1) become specialists for certain classes. For specialized classes, they show the similar performance compared to the models trained by MCL, i.e., minimizing the oracle loss (1), which considers only specialization (see Figure 1(a)). For non-specialized classes, ensemble members of CMCL are not overconfident, which makes it easy to pick a correct output via simple voting/averaging. We indeed confirm that each model trained by CMCL has not only low entropy for its specialized classes, but also exhibits high entropy for non-specialized classes as shown in Figure 2(b).

We also evaluate the quality of confidence/uncertainty level on unseen data using SVHN (Netzer et al., 2011). Somewhat surprisingly, each model trained by CMCL only using CIFAR-10 training data exhibits high entropy for SVHN test data, whereas models trained by MCL and IE are overconfident on it (see Figure 2(a) and 2(c)). We emphasize that our method can produce confident predictions significantly better than the proposed method by (Lakshminarayanan et al., 2016), which uses the averaged probability of ensemble models trained by IE to obtain high quality uncertainty estimates (see Figure 2(c)).

## 4. Regularization Techniques

In this section, we introduce advanced techniques for reducing the overconfidence and improving the performance.

### 4.1. Feature Sharing

We first propose a feature sharing scheme that stochastically shares the features among member models of CMCL to further address the overconfidence issue. The primary reason why deep learning models are overconfident is that they do not always extract general features from data. For examples, assume that some deep model only trains frogs and roses for classifying them. Although there might exist many kinds of features on their images, the model might make a decision based only on some specific features, e.g., colors. In this case, ‘red’ apples can be classified as rose with high confidence. Such an issue might be more severe in CMCL (and MCL) compared to IE since members of CMCL are specialized to certain data. To address the issue, we suggest the feature ensemble approach that encourages each model to generate meaningful abstractions from rich features extracted from other models.

Formally, consider an ensemble of  $M$  neural networks with  $L$  hidden layers. We denote the weight matrix for layer

$\ell$  of model  $m \in [M]$  and  $\ell$ -th hidden feature of model  $m$  by  $\mathbf{W}_m^\ell$  and  $\mathbf{h}_m^\ell$ , respectively. Instead of sharing the whole units of a hidden feature, we introduce random binary masks determining which units to be shared with other models. We denote the mask for layer  $\ell$  from model  $n$  to  $m$  as  $\sigma_{nm}^\ell \sim \text{Bernoulli}(\lambda)$ , which has the same dimension with  $\mathbf{h}_n^\ell$  (we use  $\lambda = 0.7$  in all experiments). Then, the  $\ell$ -th hidden feature of model  $m$  with sharing  $(\ell - 1)$ -th hidden features is defined as follows:

$$\mathbf{h}_m^\ell(\mathbf{x}) = \phi \left( \mathbf{W}_m^\ell \left( \mathbf{h}_m^{\ell-1}(\mathbf{x}) + \sum_{n \neq m} \sigma_{nm}^\ell \star \mathbf{h}_n^{\ell-1}(\mathbf{x}) \right) \right),$$

where  $\star$  denotes element-wise multiplication and  $\phi$  is the activation function. Figure 2(d) illustrates the proposed feature sharing scheme in an ensemble of deep neural networks. It makes each model learn more generalized features by sharing the features among them. However, one might expect that it might make each model overfitted due to the increased number of parameters that induces a single prediction, i.e., the statistical dependencies among outputs of models increase, which would hurt the ensemble effect. In order to handle this issue, we introduce the randomness in sharing across models in a similar manner to DropOut (Srivastava et al., 2014) using the random binary masks  $\sigma$ . In addition, we propose sharing features at lower layers since sharing the higher layers might overfit the overall networks more. For example, in all experiments with CNNs in this paper, we commonly apply feature sharing for hidden features just before the first pooling layer. We also remark that such feature sharing strategies for better generalization have also been investigated in the literature for different purposes (Misra et al., 2016; Rusu et al., 2016).

### 4.2. Stochastic Labeling

For more efficiency in minimizing the confident oracle loss, we also consider a noisy unbiased estimator of gradients of the KL divergence with Monte Carlo samples from the uniform distribution. The KL divergence from the predictive distribution to the uniform distribution can be written as follows:

$$\begin{aligned} D_{KL}(\mathcal{U}(y) \parallel P_\theta(y | \mathbf{x})) &= \sum_y \mathcal{U}(y) \log \frac{\mathcal{U}(y)}{P_\theta(y | \mathbf{x})} \\ &= \sum_y \mathcal{U}(y) \log \mathcal{U}(y) - \sum_y \mathcal{U}(y) \log P_\theta(y | \mathbf{x}). \end{aligned}$$

Hence, the gradient of the above KL divergence with respect to the model parameter  $\theta$  becomes

$$\nabla_\theta D_{KL}(\mathcal{U}(y) \parallel P_\theta(y | \mathbf{x})) = -\mathbb{E}_{\mathcal{U}(y)}[\nabla_\theta \log P_\theta(y | \mathbf{x})].$$

From the above, we induce the following noisy unbiased estimator of gradients with Monte Carlo samples from the uniform distribution:

$$-\mathbb{E}_{\mathcal{U}(y)}[\nabla_{\theta} \log P_{\theta}(y | \mathbf{x})] \simeq -\frac{1}{S} \sum_s \nabla_{\theta} \log P_{\theta}(y^s | \mathbf{x}),$$

where  $y^s \sim \mathcal{U}(y)$  and  $S$  is the number of samples. This random estimator takes samples from the uniform distribution  $\mathcal{U}(y)$  and constructs estimates of the gradient using them. In other words,  $\nabla_{\theta} \log P_{\theta}(y^s | \mathbf{x})$  is the gradient of the cross entropy loss under assigning a random label to  $\mathbf{x}$ . This stochastic labeling provides efficiency in implementation/computation and stochastic regularization effects. We formally describe detailed procedures, as the version 1 of Algorithm 1.

## 5. Experiments

We evaluate our algorithm for both classification and foreground-background segmentation tasks using CIFAR-10 (Krizhevsky & Hinton, 2009), SVHN (Netzer et al., 2011) and iCoseg (Batra et al., 2010) datasets. In all experiments, we compare the performance of CMCL with those of traditional IE and MCL using deep models. We provide the more detailed experimental setups including model architectures in the supplementary material.<sup>1</sup>

### 5.1. Image Classification

**Setup.** The CIFAR-10 dataset consists of 50,000 training and 10,000 test images with 10 image classes where each image consists of  $32 \times 32$  RGB pixels. The SVHN dataset consists of 73,257 training and 26,032 test images.<sup>2</sup> We pre-process the images with global contrast normalization and ZCA whitening following (Ian J. Goodfellow & Bengio, 2013; Zagoruyko & Komodakis, 2016), and do not use any data augmentation. Using these datasets, we train various CNNs, e.g., VGGNet (Simonyan & Zisserman, 2015), GoogLeNet (Szegedy et al., 2015), and ResNet (He et al., 2016). Similar to (Zagoruyko & Komodakis, 2016), we use the softmax classifier, and train each model by minimizing the cross-entropy loss using the stochastic gradient descent method with Nesterov momentum.

For evaluation, we measure the top-1 and oracle error rates on the test dataset. The top-1 error rate is calculated by averaging output probabilities from all models and predicting the class of the highest probability. The oracle error rate is the rate of classification failure over all outputs of individual ensemble members for a given input, i.e., it measures whether none of the members predict the correct class for

<sup>1</sup>Our code is available at <https://github.com/chhwang/cmcl>.

<sup>2</sup>We do not use the extra SVHN dataset for training.

Ensemble Method	Feature Sharing	Stochastic Labeling	Oracle Error Rate	Top-1 Error Rate
IE	-	-	10.65%	15.34%
MCL	-	-	4.40%	60.40%
CMCL	-	-	4.49%	15.65%
	✓	-	5.12%	14.83%
	✓	✓	<b>3.32%</b>	<b>14.78%</b>

Table 1. Classification test set error rates on CIFAR-10 using various ensemble methods.

an input. While a lower oracle error rate suggests higher diversity, a lower oracle error rate does not always bring a higher top-1 accuracy as this metric does not reveal the level of overconfidence of each model. By collectively measuring the top-1 and oracle error rates, one can grasp the level of specialization and confidence of a model.

**Contribution by each technique.** Table 1 validates contributions of our suggested techniques under comparison with other ensemble methods IE and MCL. We evaluate an ensemble of five simple CNN models where each model has two convolutional layers followed by a fully-connected layer. We incrementally apply our optimizations to gauge the stepwise improvement by each component. One can note that CMCL significantly outperforms MCL in the top-1 error rate even without feature sharing or stochastic labeling while it still provides a comparable oracle error rate. By sharing the 1st ReLU activated features, the top-1 error rates are improved compared to those that employ only confident oracle loss. Stochastic labeling further improves both error rates. This implies that stochastic labeling not only reduces computational burdens but also provides regularization effects.

**Overlapping.** As a natural extension of CMCL, we also consider picking  $K$  specialized models instead of having only one specialized model, which was investigated for original MCL (Guzman-Rivera et al., 2012; Lee et al., 2016). This is easily achieved by modifying the constraint (3b) as  $\sum_{m=1}^M v_i^m = K$ , where  $K$  is an overlap parameter that controls training data overlap between the models. This simple but natural scheme brings extra gain in top-1 performance by generalizing each model better. Table 2 compares the performance of various ensemble methods with varying values of  $K$ . Under the choice of  $K = 4$ , CMCL of 10 CNNs provides 9.13% relative reduction in the top-1 error rates from the corresponding IE. Somewhat interestingly, IE has similar error rates on ensembles of both 5 and 10 CNNs, which implies that the performance of CMCL might be impossible to achieve using IE even if one increases the number of models in IE.

**Large-scale CNNs.** We now evaluate the performance of our ensemble method when it is applied to larger-scale CNN models for image classification tasks on CIFAR-10

Confident Multiple Choice Learning

Ensemble Method	$K$	Ensemble Size $M = 5$		Ensemble Size $M = 10$	
		Oracle Error Rate	Top-1 Error Rate	Oracle Error Rate	Top-1 Error Rate
IE	-	10.65%	15.34%	9.26%	15.34%
	1	4.40%	60.40%	<b>0.00%</b>	76.88%
MCL	2	3.75%	20.66%	1.46%	49.31%
	3	4.73%	16.24%	1.52%	22.63%
	4	5.83%	15.65%	1.82%	17.61%
CMCL	1	<b>3.32%</b>	14.78%	1.96%	14.28%
	2	3.69%	<b>14.25% (-7.11%)</b>	1.22%	13.95%
	3	4.38%	14.38%	1.53%	14.00%
	4	5.82%	14.49%	1.73%	<b>13.94% (-9.13%)</b>

Table 2. Classification test set error rates on CIFAR-10 with varying values of the overlap parameter  $K$  explained in Section 5.1. We use CMCL with both feature sharing and stochastic labeling. Boldface values in parentheses represent the relative reductions from the best results of MCL and IE.

Model Name	Ensemble Method	CIFAR-10		SVHN	
		Oracle Error Rate	Top-1 Error Rate	Oracle Error Rate	Top-1 Error Rate
VGGNet-17	-(single)	10.65%	10.65%	5.22%	5.22%
	IE	3.27%	8.21%	1.99%	4.10%
	MCL	<b>2.52%</b>	45.58%	<b>1.45%</b>	45.30%
	CMCL	2.95%	<b>7.83% (-4.63%)</b>	1.65%	<b>3.92% (-4.39%)</b>
GoogLeNet-18	-(single)	10.15%	10.15%	4.59%	4.59%
	IE	3.37%	7.97%	1.78%	3.60%
	MCL	<b>2.41%</b>	52.03%	1.39%	37.92%
	CMCL	2.78%	<b>7.51% (-5.77%)</b>	<b>1.36%</b>	<b>3.44% (-4.44%)</b>
ResNet-20	-(single)	14.03%	14.03%	5.31%	5.31%
	IE	3.83%	10.18%	1.82%	3.94%
	MCL	<b>2.47%</b>	53.37%	1.29%	40.91%
	CMCL	2.79%	<b>8.75% (-14.05%)</b>	<b>1.42%</b>	<b>3.68% (-6.60%)</b>

Table 3. Classification test set error rates on CIFAR-10 and SVHN for various large-scale CNN models. We train an ensemble of 5 models, and use CMCL with both feature sharing and stochastic labeling. Boldface values in parentheses indicate relative error rate reductions from the best results of MCL and IE.

and SVHN datasets. Specifically, we test VGGNet (Simonyan & Zisserman, 2015), GoogLeNet (Szegedy et al., 2015), and ResNet (He et al., 2016). We share the non-linear activated features right before the first pooling layer, i.e., the 6th, 2nd, and 1st ReLU activations for ResNet with 20 layers, VGGNet with 17 layers, and GoogLeNet with 18 layers, respectively. This choice is for maximizing the regularization effect of feature sharing while minimizing the statistical dependencies among the ensemble models. For all models, we choose the best hyperparameters for confident oracle loss among the penalty parameter  $\beta \in \{0.5, 0.75, 1, 1.25, 1.5\}$  and the overlapping parameter  $K \in \{2, 3, 4\}$ . Table 3 shows that CMCL consistently outperforms all baselines with respect to the top-1 error rate while producing comparable oracle error rates to those of MCL. We also apply the feature sharing to IE as reported in Figure 4(a). Even though the feature sharing also improves the performance of IE, CMCL still outperforms IE: CMCL provides 6.11% relative reduction of the top-1

error rate from the IE with feature sharing under the choice of  $M = 10$ . We also remark that IE with feature sharing has similar error rates as the ensemble size increases, while CMCL does not (i.e., the gain is more significant for CMCL). This implies that feature sharing is more effectively working for CMCL.

## 5.2. Foreground-Background Segmentation

In this section, we evaluate if ensemble models trained with CMCL produce high-quality segmentation of foreground and background of an image with the iCoseg dataset. The foreground-background segmentation is formulated as a pixel-level classification problem with 2 classes, i.e., 0 (background) or 1 (foreground). To tackle the problem, we design fully convolutional networks (FCNs) model (Long et al., 2015) based on the decoder architecture presented in (Radford et al., 2016). The dataset consists of 38 groups of related images with pixel-level ground truth on foreground-background segmentation of each image. We only use im-

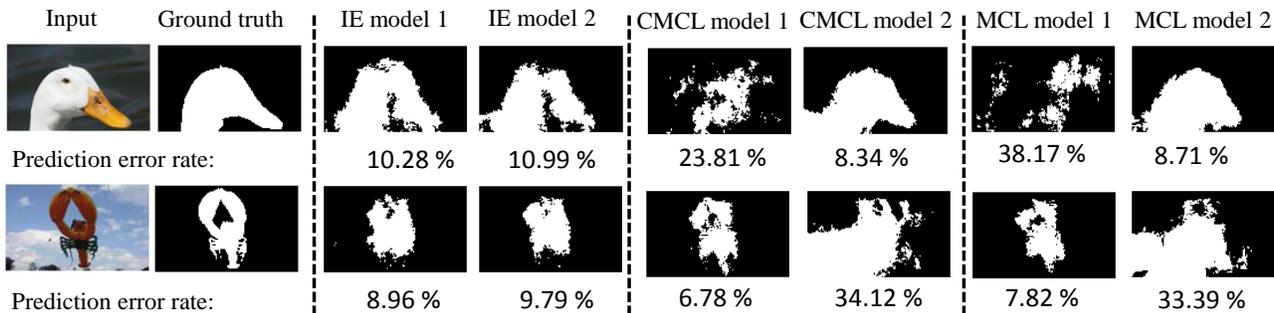


Figure 3. Prediction results of foreground-background segmentation for a few sample images. A test error rate is shown below each prediction. The ensemble models trained by CMCL and MCL generate high-quality predictions specialized for certain images.

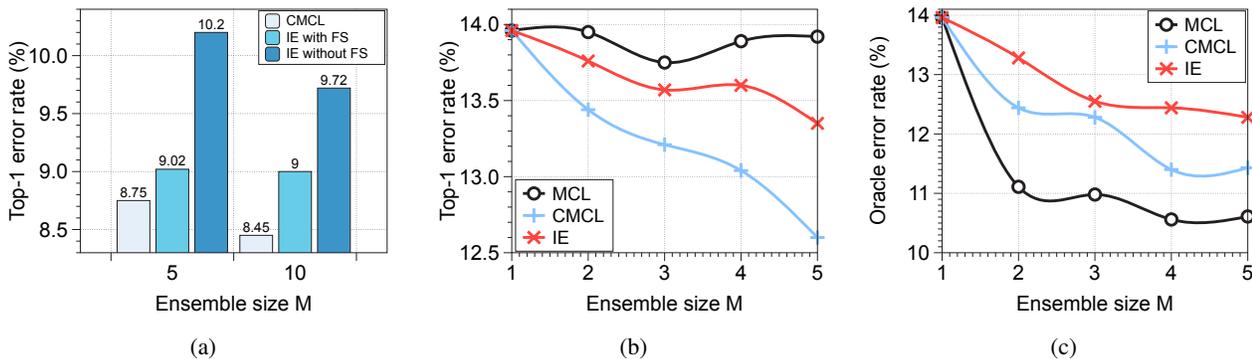


Figure 4. (a) Top-1 error rate on CIFAR-10. We train an ensemble of  $M$  ResNets with 20 layers, and apply feature sharing (FS) to IE and CMCL. (b) Top-1 error rate and (c) oracle error rate on iCoseg by varying the ensemble sizes. The ensemble models trained by CMCL consistently improves the top-1 error rate over baselines.

ages that are larger than  $300 \times 500$  pixels. For each class, we randomly split 80% and 20% of the data into training and test sets, respectively. We train on  $75 \times 125$  resized images using the bicubic interpolation (Keys, 1981). Similar to (Guzman-Rivera et al., 2012; Lee et al., 2016), we initialize the parameters of FCNs with those trained by IE for MCL and CMCL. For all experiments, CMCL is used with both feature sharing and stochastic labeling.

Similar to (Guzman-Rivera et al., 2012), we define the percentage of incorrectly labeled pixels as prediction error rate. We measure the oracle error rate (i.e., the lowest error rate over all models for a given input) and the top-1 error rate. The top-1 error rate is measured by following the predictions of the member model that has a lower pixel-wise entropy, i.e., picking the output of a more confident model. For each ensemble method, we vary the number of ensemble models and measure the oracle error rate and test error rate. Figure 4(b) and 4(c) show both top-1 and oracle error rates for all ensemble methods. We remark that the ensemble models trained by CMCL consistently improves the top-1 error rate over baselines. In an ensemble of 5 models, we find that CMCL achieve up to 6.77% relative reduction

in the top-1 error rate from the corresponding IE. As shown in Figure 3, an individual model trained by CMCL generates high-quality solutions by specializing itself in specific images (e.g., model 1 is specialized for ‘lobster’ while model 2 is specialized for ‘duck’) while each model trained by IE does not.

## 6. Conclusion

This paper proposes CMCL, a novel ensemble method of deep neural networks that produces diverse/plausible confident prediction of high quality. To this end, we address the over-confidence issues of MCL, and propose a new loss, architecture and training method. In our experiments, CMCL outperforms not only the known MCL, but also the traditional IE, with respect to the top-1 error rates in classification and segmentation tasks. The recent trend in the deep learning community tends to make models bigger and wider. We believe that our new ensemble approach brings a refreshing angle for developing advanced large-scale deep neural networks in many related applications.

## Acknowledgements

This work was supported in part by the ICT R&D Program of MSIP/IITP, Korea, under [2016-0-00563, Research on Adaptive Machine Learning Technology Development for Intelligent Autonomous Digital Companion], R0190-16-2012, [High Performance Big Data Analytics Platform Performance Acceleration Technologies Development], and by the National Research Council of Science & Technology (NST) grant by the Korea government (MSIP) (No. CRC-15-05-ETRI).

## References

- Batra, Dhruv, Kowdle, Adarsh, Parikh, Devi, Luo, Jiebo, and Chen, Tsuhan. icoseg: Interactive co-segmentation with intelligent scribble guidance. In *Computer Vision and Pattern Recognition (CVPR)*, pp. 3169–3176. IEEE, 2010.
- Batra, Dhruv, Yadollahpour, Payman, Guzman-Rivera, Abner, and Shakhnarovich, Gregory. Diverse m-best solutions in markov random fields. In *European Conference on Computer Vision (ECCV)*, pp. 1–16. Springer, 2012.
- Breiman, Leo. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- Ciregan, Dan, Meier, Ueli, and Schmidhuber, Jürgen. Multi-column deep neural networks for image classification. In *Computer Vision and Pattern Recognition (CVPR)*, pp. 3642–3649. IEEE, 2012.
- Collins, Michael and Koo, Terry. Discriminative reranking for natural language parsing. *Computational Linguistics*, 31(1):25–70, 2005.
- Dey, Debadepta, Ramakrishna, Varun, Hebert, Martial, and Andrew Bagnell, J. Predicting multiple structured visual interpretations. In *International Conference on Computer Vision (ICCV)*, pp. 2947–2955, 2015.
- Domingos, Pedro. Bayesian averaging of classifiers and the overfitting problem. In *International Conference on Machine Learning (ICML)*, volume 2000, pp. 223–230, 2000.
- Everingham, Mark, Van Gool, Luc, Williams, Christopher KI, Winn, John, and Zisserman, Andrew. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- Freund, Yoav, Schapire, Robert, and Abe, N. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612, 1999.
- Guzman-Rivera, Abner, Batra, Dhruv, and Kohli, Pushmeet. Multiple choice learning: Learning to produce multiple structured outputs. In *Advances in neural information processing systems (NIPS)*, pp. 1799–1807, 2012.
- Guzman-Rivera, Abner, Kohli, Pushmeet, Batra, Dhruv, and Rutenbar, Rob A. Efficiently enforcing diversity in multi-output structured prediction. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 2, pp. 3, 2014.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Huang, Liang and Chiang, David. Better k-best parsing. In *Proceedings of the Ninth International Workshop on Parsing Technology*, pp. 53–64. Association for Computational Linguistics, 2005.
- Ian J. Goodfellow, David Warde-Farley, Mehdi Mirza Aaron Courville and Bengio, Yoshua. Maxout networks. In *International Conference on Machine Learning (ICML)*, pp. 1319–1327, 2013.
- Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, pp. 448–456, 2015.
- Jacobs, Robert A, Jordan, Michael I, Nowlan, Steven J, and Hinton, Geoffrey E. Adaptive mixtures of local experts. *Neural computation*, 1991.
- Keys, Robert. Cubic convolution interpolation for digital image processing. *IEEE transactions on acoustics, speech, and signal processing*, 29(6):1153–1160, 1981.
- Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Krizhevsky, Alex and Hinton, Geoffrey. Learning multiple layers of features from tiny images. 2009.
- Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems (NIPS)*, pp. 1097–1105, 2012.
- Lakshminarayanan, Balaji, Pritzel, Alexander, and Blundell, Charles. Simple and scalable predictive uncertainty estimation using deep ensembles. *NIPS Workshop on Bayesian Deep Learning*, 2016.
- Lee, Stefan, Purushwalkam, Senthil, Cogswell, Michael, Crandall, David, and Batra, Dhruv. Why m heads are better than one: Training a diverse ensemble of deep networks. *arXiv preprint arXiv:1511.06314*, 2015.

- Lee, Stefan, Prakash, Senthil Purushwalkam Shiva, Cogswell, Michael, Ranjan, Viresh, Crandall, David, and Batra, Dhruv. Stochastic multiple choice learning for training diverse deep ensembles. In *Advances in neural information processing systems (NIPS)*, pp. 2119–2127, 2016.
- Lin, Min, Chen, Qiang, and Yan, Shuicheng. Network in network. In *International Conference on Learning Representations (ICLR)*, 2014a.
- Lin, Tsung-Yi, Maire, Michael, Belongie, Serge, Hays, James, Perona, Pietro, Ramanan, Deva, Dollár, Piotr, and Zitnick, C Lawrence. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, pp. 740–755. Springer, 2014b.
- Long, Jonathan, Shelhamer, Evan, and Darrell, Trevor. Fully convolutional networks for semantic segmentation. In *Computer Vision and Pattern Recognition (CVPR)*, pp. 3431–3440, 2015.
- Misra, Ishan, Shrivastava, Abhinav, Gupta, Abhinav, and Hebert, Martial. Cross-stitch networks for multi-task learning. In *Computer Vision and Pattern Recognition (CVPR)*, pp. 3994–4003, 2016.
- Netzer, Yuval, Wang, Tao, Coates, Adam, Bissacco, Alessandro, Wu, Bo, and Ng, Andrew Y. Reading digits in natural images with unsupervised feature learning. *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011(2):5, 2011.
- Nguyen, Anh, Yosinski, Jason, and Clune, Jeff. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Computer Vision and Pattern Recognition (CVPR)*, pp. 427–436, 2015.
- Park, Dennis and Ramanan, Deva. N-best maximal decoders for part models. In *International Conference on Computer Vision (ICCV)*, pp. 2627–2634. IEEE, 2011.
- Pereyra, Gabriel, Tucker, George, Chorowski, Jan, Kaiser, Łukasz, and Hinton, Geoffrey. Regularizing neural networks by penalizing confident output distributions. In *International Conference on Learning Representations (ICLR)*, 2017.
- Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- Rusu, Andrei A, Rabinowitz, Neil C, Desjardins, Guillaume, Soyer, Hubert, Kirkpatrick, James, Kavukcuoglu, Koray, Pascanu, Razvan, and Hadsell, Raia. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- Simonyan, Karen and Zisserman, Andrew. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- Srivastava, Nitish, Hinton, Geoffrey E, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1): 1929–1958, 2014.
- Szegedy, Christian, Liu, Wei, Jia, Yangqing, Sermanet, Pierre, Reed, Scott, Anguelov, Dragomir, Erhan, Dumitru, Vanhoucke, Vincent, and Rabinovich, Andrew. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- Wan, Li, Zeiler, Matthew D., Zhang, Sixin, LeCun, Yann, and Fergus, Rob. Regularization of neural networks using dropconnect. In *International Conference on Machine Learning (ICML)*, pp. 1058–1066, 2013.
- Zagoruyko, Sergey and Komodakis, Nikos. Wide residual networks. In *British Machine Vision Conference (BMVC)*, 2016.