

ST-CoNAL: Consistency-Based Acquisition Criterion Using Temporal Self-Ensemble for Active Learning

Jae Soon Baik^[0000–0002–1508–0619], In Young Yoon^[0000–0003–3173–3947], and
Jun Won Choi^{*[0000–0002–3733–0148]}

Hanyang University, Seoul, Korea
{jsbaik, inyoungyoon}@spa.hanyang.ac.kr, junwchoi@hanyang.ac.kr

Abstract. Modern deep learning has achieved great success in various fields. However, it requires the labeling of huge amounts of data, which is expensive and labor-intensive. Active learning (AL), which identifies the most informative samples to be labeled, is becoming increasingly important to maximize the efficiency of the training process. The existing AL methods mostly use only a single final fixed model for acquiring the samples to be labeled. This strategy may not be good enough in that the structural uncertainty of a model for given training data is not considered to acquire the samples. In this study, we propose a novel acquisition criterion based on temporal self-ensemble generated by conventional stochastic gradient descent (SGD) optimization. These self-ensemble models are obtained by capturing the intermediate network weights obtained through SGD iterations. Our acquisition function relies on a *consistency measure* between the student and teacher models. The student models are given a fixed number of temporal self-ensemble models, and the teacher model is constructed by averaging the weights of the student models. Using the proposed acquisition criterion, we present an AL algorithm, namely *student-teacher consistency-based AL* (ST-CoNAL). Experiments conducted for image classification tasks on CIFAR-10, CIFAR-100, Caltech-256, and Tiny ImageNet datasets demonstrate that the proposed ST-CoNAL achieves significantly better performance than the existing acquisition methods. Furthermore, extensive experiments show the robustness and effectiveness of our methods.

Keywords: Active Learning · Consistency · Temporal Self-Ensemble · Image Classification .

1 Introduction

Deep neural networks (DNNs) require a large amount of training data to optimize millions of weights. In particular, for supervised-learning tasks, labeling of training data by human annotators is expensive and time-consuming. The labeling cost can be a major concern for machine learning applications, which requires

* Corresponding Author

a collection of real-world data on a massive scale (e.g., autonomous driving) or the knowledge of highly trained experts for annotation (e.g., automatic medical diagnosis). Active learning (AL) is a promising machine learning framework that maximizes the efficiency of labeling tasks within a fixed labeling budget [35].

This study focuses on the pool-based AL problem, where the data instances to be labeled are selected from a pool of unlabeled data. In a pool-based AL method, the decision to label a data instance is based on a *sample acquisition function*. The acquisition function, $a(x, f)$, takes the input instance x and the currently trained model f and produces a score to decide if x should be labeled. Till date, various types of AL methods have been proposed [9–11, 20, 35–37, 45, 46]. The predictive uncertainty-based methods [9, 10, 36, 45] used well-studied theoretic measures such as entropy and mutual information. Recently, representation-based methods [20, 34, 37, 46] have been widely used as a promising AL approach to exploit high-quality representation of DNNs. However, the acquisition used for these methods rely on a single trained model f , failing to account for model uncertainty arising given a limited labeled dataset. To solve this problem, several AL methods [2, 35] attempted to utilize an ensemble of DNNs and design acquisition functions based on them. However, these methods require significant computational costs to train multiple networks.

When the amount of labeled data is limited, semi-supervised learning (SSL) is another promising machine learning approach to improve performance with low labeling costs. SSL improves the model performance by leveraging a large number of unlabeled examples [31]. *Consistency regularization* is one of the most successful approach to SSL [11, 22, 30, 39]. In a typical semi-supervised learning, the model is trained using the consistency-regularized loss function $\mathcal{E}_{ce} + \lambda \mathcal{E}_{con}$, where \mathcal{E}_{ce} denotes the cross-entropy loss and \mathcal{E}_{con} denotes the consistency-regularized loss. Minimization of \mathcal{E}_{con} regularizes the model to produce consistent predictions over the training process, improving the performance for a given task. *IT* model [22] applied a random perturbation to the input of the DNN and measured the consistency between the model outputs. Mean Teacher (MT) [39] produced the temporal self-ensemble through SGD iterations and measured the consistency between the model being trained and the teacher model obtained by taking the exponential moving average (EMA) of the self-ensemble. These methods successfully regularized the model to produce consistent predictions while using perturbed predictions on unlabeled samples.

The objective of this study is to improve the sample acquisition criterion for pool-based AL methods. Inspired by the consistency regularization for SSL, we build a new sample acquisition criterion that measures consistency between multiple ensemble models obtained during the training phase. The proposed acquisition function generates the temporal self-ensemble by sampling the models at the intermediate check-points of the weight trajectory formed by the SGD optimization. This provides a better acquisition performance and eliminates the additional computational cost required by previous ensemble-based AL methods. In [1, 17, 27], the aforementioned method has shown to produce good and diverse self-ensembles, where were used to improve the inference model via *stochastic*




Fig. 1: **Acquisition criterion:** Both labeled and unlabeled samples are represented in the feature space in the binary classification problem. *Low consistency samples* produce the predictions which significantly differ between the student and teacher models. *Low confidence samples* are found near the decision boundaries of the student models. The proposed ST-CoNAL evaluates the consistency measure for each data sample and selects those with the highest consistency score.

weight averaging (SWA) [1, 17, 27]. We derive the acquisition criterion based on the temporal self-ensemble models used in SWA. We present the AL method, referred to as *student-teacher consistency-based AL* (ST-CoNAL), which measures the consistency between the student and teacher models. The self-ensemble model constitutes a student model, and a teacher model is formed by taking an equally-weighted average (EWA) of the parameters of the student models. Treating the output of the teacher model as a desired supervisory signal, ST-CoNAL measures the Kullback–Leibler (KL) divergence of each teacher-student output pairs. The acquisition function of ST-CoNAL acquires the samples to be labeled that yield the highest inconsistency.

Though ST-CoNAL was inspired by the consistency regularization between student and teacher models of the MT, these two methods are quite different in the following aspects. MT constructs the teacher model by assigning larger weights to more recent model weights obtained through SGD iterations. Due to this constraint, as training progresses, the teacher model in MT tends to be correlated with the student models, making it difficult to measure good enough consistency measure for AL acquisition. To address this problem, ST-CoNAL generates a better teacher model by taking an equally-weighted averaging (EWA) of the weights of the student models instead of EMA used in MT. Similar to previous AL methods [9, 10] that utilize ensemble models to capture the posterior distribution of model weights for a given data set, the use of student model weights allows our acquisition criterion to account for model uncertainty.

We further improve our ST-CoNAL method by adopting the principle of entropy minimization used for SSL [3, 4, 12, 24, 38, 43]. We apply one of the entropy

minimization methods, sharpening to the output of the teacher model. When the sharpened output is used for our KL divergence, our acquisition criterion can measure the uncertainty of the prediction for the given sample. Our evaluation shows that the ST-CoNAL is superior to other AL methods on various image classification benchmark datasets.

Fig.1 illustrates that these low consistency samples lie in the region of the feature space where the student models produce the predictions of a larger variation. Note that these samples are not necessarily identical to the low confidence samples located near the decision boundary specified by the teacher model. The proposed ST-CoNAL prefers the acquisition of the inconsistent samples rather than the low-confidence samples.

The main contributions of our study are summarized as follows:

- We propose a new acquisition criterion based on the temporal self-ensemble. Temporal self-ensemble models are generated by sampling DNN weights through SGD optimization. ST-CoNAL measures the consistency between these self-ensemble models and acquires the most inconsistent samples for labeling. We evaluated the performance of ST-CoNAL on four different public datasets for multi-class image classification tasks. We observe that the proposed ST-CoNAL method achieves the significant performance gains over other AL methods.
- We identified a work relevant to ours [11]. While both ours and their work aim to exploit consistency regularization for AL, our work differs from theirs in the following aspects. While CSSAL [11] relies on the input perturbation to a single fixed model, ST-CoNAL utilizes the self-ensemble models to measure the consistency. Note that the benefits of using model ensembles for AL have been demonstrated in [2] and our findings about the superior performance of ST-CoNAL over CSSAL are consistent with the results of these studies.

2 Related Work

2.1 Acquisition for AL

The acquisition methods for pool-based AL can be roughly categorized into three classes: 1) *uncertainty-based*, 2) *representation-based* and 3) *consistency-based methods*. The uncertainty-based methods [2, 9, 10, 19, 26, 41] estimate the prediction uncertainty for the given data instances. Max-entropy [26, 36], least confidence [26] and variation ratio [18] are the widely used criteria. Recently, the representation-based method [20, 34, 37, 46] significantly improved the acquisition performance. VAAL [37] performed adversarial learning to find informative unlabeled data in the task-agnostic latent representation of the data. TA-VAAL [20] is the latest state-of-the-art method, which incorporated a task-aware uncertainty-based approach to the VAAL baseline. The consistency-based methods measure the disagreement among the predictions of the ensemble models [9–11, 15]. TOD-Semi [16], *Query-by-committee* [35], and DBAL [10] are the well known criteria used in these methods.

2.2 Consistency Regularization for SSL

One of the most successful approaches to SSL is *consistency regularization*. The Π model [22,33] enforces consistency for model $f_\theta(x)$ by minimizing $d(f_\theta(x), f_\theta(\hat{x}))$, where \hat{x} denotes the data perturbed by noise and augmentation, and $d(a, b)$ measures the distance between a and b . While Π model uses the predictions as a target, the consistency loss for regularizing the consistent behavior between an ensemble prediction and the model prediction was shown to be effective. *Mean teacher* method [39] generates the weights of the teacher model with EMA of the student weights over SGD iterations. ICT [42] and MixMatch [4] enforce consistency between linearly interpolated inputs and model predictions.

2.3 Temporal Self-Ensemble via Stochastic Weight Averaging

Several studies have attempted to use temporal self-ensemble obtained through model optimization to improve the inference performance [6, 17, 22, 25]. Izmailov et al. [17] proposed SWA method, which determines an improved inference model using the equally weighted average of the intermediate model weights traversed via SGD iterates. Using cyclical learning rate scheduling [27], SGD can yield the weights in the flat regions in the loss surface. By averaging the weights found at SGD trajectory, SWA can find the weight solution that can generalize well. In [1], consistency regularization and modified cyclical learning rate scheduling assisted in finding the improved weights through SWA.

2.4 Combining AL and SSL

The combination of SSL and AL was discussed in several works [8, 10, 11, 32, 40, 47]. However, most works considered applying SSL and AL independently. Only a few works considered the joint design of SSL and AL. Zhu et al. [47] used a Gaussian random field on a weighted graph to combine SSL and AL. Additionally, the authors of [11] proposed an enhanced acquisition function by adopting the variance over the data augmentation of the input image.

3 Proposed Active Learning

In this section, we present the details of the ST-CoNAL algorithm.

3.1 Consistency Measure-Based Acquisition

The proposed algorithm follows the setup of a pool-based AL for the K -image class classification task. Suppose that the labeling budget is fixed to b data samples. In the j th sample acquisition step, we are given the dataset \mathcal{D}^j , which consists of a set of labeled data \mathcal{D}_L^j and a set of unlabeled data \mathcal{D}_U^j . $f_{\mathcal{D}^j}$ is the model trained with dataset \mathcal{D}^j . We consider two cases; 1) $f_{\mathcal{D}^j}$ is trained with only \mathcal{D}_L^j via supervised learning and 2) $f_{\mathcal{D}^j}$ is trained with $\mathcal{D}_L^j \cup \mathcal{D}_U^j$ via SSL.

The AL algorithm aims to acquire the set of b data samples $\{x_1^*, \dots, x_b^*\}$ from the pool of unlabeled data \mathcal{D}_U^j . To select the data samples to be labeled, the acquisition function $a_{\text{ST-CoNAL}}(x_i, f_{\mathcal{D}^j})$ is used to score the given data instance x_i with the currently trained model $f_{\mathcal{D}^j}$. The b samples that yield the highest score are selected as

$$\{x_1^*, \dots, x_b^*\} = \underset{\{x_1, \dots, x_b\} \subset \mathcal{D}_U^j}{\operatorname{argmax}} \sum_{i=1}^b a_{\text{ST-CoNAL}}(x_i, f_{\mathcal{D}^j}). \quad (1)$$

The selected samples are then labeled, i.e., $\mathcal{D}_L^{j+1} \leftarrow \mathcal{D}_L^j \cup \{x_1^*, \dots, x_b^*\}$ and $\mathcal{D}_U^{j+1} \leftarrow \mathcal{D}_U^j \setminus \{x_1^*, \dots, x_b^*\}$, and the model is retrained using the newly labeled dataset $\mathcal{D}^{j+1} = \{\mathcal{D}_L^{j+1}, \mathcal{D}_U^{j+1}\}$. This procedure is repeated until a labeled dataset is obtained as large as required within the limited budget.

Suppose that we have Q student models $f_{\mathcal{D}^j}^{(s,1)}, \dots, f_{\mathcal{D}^j}^{(s,Q)}$ and one teacher model $f_{\mathcal{D}^j}^{(t)}$ trained with the dataset \mathcal{D}^j . Each model produces the K dimensional output through the softmax output layer. The method for obtaining these models will be discussed in the next section. We first apply sharpening [3, 4, 43] to the output of the teacher model as

$$\bar{f}_{\mathcal{D}^j}^{(t)}(k) = \frac{\left(f_{\mathcal{D}^j}^{(t)}(k)\right)^{\frac{1}{T}}}{\sum_{k=1}^K \left(f_{\mathcal{D}^j}^{(t)}(k)\right)^{\frac{1}{T}}}, \quad (2)$$

where $f(k)$ denotes the k th element of the output produced by the function $f(\cdot)$ and T denotes the temperature hyper-parameter. Then, the acquisition function of the ST-CoNAL is obtained by accumulating the KL divergence between the predictions produced by each student model and the teacher model, i.e.,

$$a_{\text{ST-CoNAL}}(x_i, f_{\mathcal{D}^j}) = \frac{1}{Q} \sum_{q=1}^Q KL(\bar{f}_{\mathcal{D}^j}^{(t)}(x_i) || f_{\mathcal{D}^j}^{(s,q)}(x_i)), \quad (3)$$

where $KL(p||q) = \sum_i p_i \log(p_i/q_i)$. The KL divergence in this acquisition function measures the inconsistency between the student and teacher models. As the sharpening makes the prediction of teacher model to have low-entropy, the high KL divergence value reflects the uncertainty in the prediction of the student models.

3.2 Generation of Student and Teacher Models via Optimization Path

The ST-CoNAL computes the student and teacher models using the temporal self-ensemble networks. The parameters of these self-ensemble networks are obtained by capturing the network weights at the intermediate check points of the weight trajectory. The study in [1] revealed that the average of the weights obtained through the SGD iterates can yield the weights corresponding to the

Algorithm 1: ST-CoNAL

Input:
Initial Dataset: $\mathcal{D}^1 = (\mathcal{D}_L^1, \mathcal{D}_U^1)$
Total number of acquisition steps: J
Labeling budget: b
Temperature parameter: T
for $j = 1$ **to** J **do**
 Training:
 Generate the student models $f_{\mathcal{D}^j}^{(s,1)}, \dots, f_{\mathcal{D}^j}^{(s,Q)}$ using SGD with the modified LR schedule in (4).
 Sample acquisition:
 Compute the teacher model $f_{\mathcal{D}^j}^{(t)}$ via weight averaging.
 $\bar{f}_{\mathcal{D}^j}^{(t)} = \text{Sharpen}(f_{\mathcal{D}^j}^{(t)}, T)$.
 Evaluate $a_{\text{ST-CoNAL}}(x_i, f_{\mathcal{D}^j})$ for all $x_i \in \mathcal{D}_U^j$.
 Acquire b samples according to (1).
 Dataset update:
 $\{(x_1^*, y_1^*), \dots, (x_b^*, y_b^*)\} = \text{Annotate}(\{x_1^*, \dots, x_b^*\})$
 $\mathcal{D}_L^{j+1} \leftarrow \mathcal{D}_L^j \cup \{x_1^*, \dots, x_b^*\}$
 $\mathcal{D}_U^{j+1} \leftarrow \mathcal{D}_U^j \setminus \{x_1^*, \dots, x_b^*\}$
 $\mathcal{D}^{j+1} \leftarrow (\mathcal{D}_L^{j+1}, \mathcal{D}_U^{j+1})$
end
Train the model using \mathcal{D}^{J+1} for SSL or \mathcal{D}_L^{J+1} for supervised learning.
return Trained model

flat minima of the loss surface, which are known to generalize well [29, 44]. In ST-CoNAL, the average of these weights forms a teacher model. To obtain the diverse student models, we adopt the learning rate (LR) schedule used for SWA [17, 29, 44]. The learning rate l used for our method is given by

$$l = \begin{cases} l_0, & \text{if } t < T_0 \\ \gamma l_0, & \text{otherwise} \end{cases} \quad (4)$$

where l_0 denotes the initial learning rate, $\gamma < 1$ denotes the parameter for learning rate decay, t is the current training epoch, and T_0 denotes the training epoch at which the learning rate will be switched to a smaller value. After running T_0 epochs (i.e., $t \geq T_0$), the Q intermediate network weights are stored every c epoch, which constitute the weights of the student networks. We obtain Q self-ensemble networks $f_{\mathcal{D}^j}^{(s,1)}, \dots, f_{\mathcal{D}^j}^{(s,Q)}$ and these models are treated as the student models. The teacher model $f_{\mathcal{D}^j}^{(t)}$ is then obtained by taking the equally weighted average of the parameters of $f_{\mathcal{D}^j}^{(s,1)}, \dots, f_{\mathcal{D}^j}^{(s,Q)}$. Note that this learning rate schedule does not require any additional training cost for obtaining the temporal self-ensemble.

3.3 Summary of Proposed ST-CoNAL

Algorithm 1 presents the summary of the ST-CoNAL algorithm.

4 Experiments

In this section, we evaluate the performance of ST-CoNAL via experiments conducted on four public datasets, CIFAR-10 [21], CIFAR-100 [21], Caltech-256 [13] and Tiny ImageNet [23].

4.1 Experiment Setup

Datasets. We evaluated ST-CoNAL on four benchmarks: CIFAR-10 [21], CIFAR-100 [21], Caltech-256 [13] and Tiny ImageNet [23]. CIFAR-10 contains 50k training examples with 10 categories and CIFAR-100 contains 50k training examples with 100 categories. Caltech-256 has 24,660 training examples with 256 categories and Tiny ImageNet has 100k training examples with 200 categories.

To see how ST-CoNAL performed on class-imbalanced scenarios, we additionally used the synthetically imbalanced CIFAR-10 datasets; *the step imbalanced CIFAR-10* [7, 20] and the *long-tailed CIFAR-10* [5, 7]. The step imbalance CIFAR-10 has 50 samples for the first five smallest classes and 5,000 samples for the last five largest classes. Long-tailed CIFAR-10 followed the configuration used in [5], where 12,406 training images were generated with 5,000 samples from the largest class and 50 samples from the smallest class. The imbalance ratio of 100 is the widely used setting adopted in the literature [5, 20].

Implementation Details. In all experiments, an 18-layer residual network (ResNet-18) was used as the backbone network. The detailed structure of this network was provided in [14]. For CIFAR-10 and CIFAR-100 datasets, models were trained with standard data augmentation operations including random translation and random horizontal flipping [14]. For Caltech-256 and Tiny-ImageNet datasets, we applied the data augmentation operations including resizing to 256×256 , random cropping to 224×224 , and horizontal flipping. In the inference phase, we applied the image resizing method followed by the center cropping to generate 224×224 input images [14]. We also applied our ST-CoNAL method to SSL setup. We used the mean teacher method using the detailed configurations described in [39]. We conducted all experiments with a TITAN XP GPU for training.

Configurations for AL. According to [34, 35], selection of b samples from a large pool of unlabeled samples requires computation of the acquisition function for all unlabeled samples in the pool, yielding considerable computational overhead. To mitigate this issue, the subset of unlabeled samples \mathcal{S} was randomly drawn from the unlabeled pool \mathcal{D}_U [2, 20, 45]. Then, acquisition functions were used to acquire the best b samples from \mathcal{S} .

In all experiments, the temperature parameter T was set to 0.7, which was determined through empirical optimization. For CIFAR-10, we set $\gamma = 1.0$, $c =$

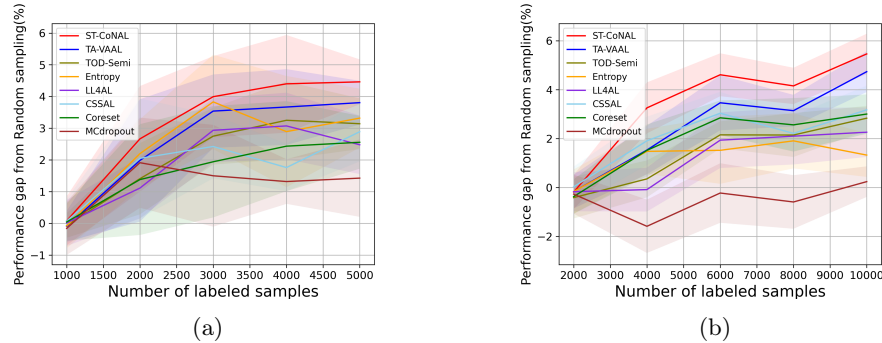


Fig. 2: Average accuracy improvement from random sampling versus the number of labeled samples on (a) CIFAR-10 and (b) CIFAR-100 datasets.

10, $b = 1k$ and $|\mathcal{S}| = 10k$, and for CIFAR-100, we set $\gamma = 0.5$, $c = 10$, $b = 2k$ and $|\mathcal{S}| = 10k$. For Tiny ImageNet, we set $\gamma = 0.5$, $c = 10$, $b = 2k$ and $|\mathcal{S}| = 20k$. For Caltech-256, we set $\gamma = 0.3$, $c = 10$, $b = 1k$ and $|\mathcal{S}| = 10k$. These configurations followed the convention adopted in the existing methods. The parameter T_0 used for the learning rate scheduling was set to 160 epochs for all datasets except Tiny ImageNet. We used T_0 of 60 epochs for Tiny ImageNet.

Baselines We compared our method with the existing AL methods including Core-set [34], MC-Dropout [9], LL4AL [45], Entropy [36], CSSAL [11], TOD-Semi [16], TA-VAAL [20], and random sampling. We adopted the average classification accuracy as a performance metric. We report the mean and standard deviation of the classification accuracy measured 5 times with different random seeds. In our performance comparison, we kept the network optimization and inference methods the same and only changed the acquisition criterion. TOD-Semi, TA-VAAL and LL4AL are exceptions because we used the optimization and inference methods proposed in their original paper. In Fig. 2 to Fig. 5, the performance difference of the target method from the random sampling baseline was used as a performance indicator.

4.2 Performance Comparison

CIFAR-10 and CIFAR-100: Fig. 2 (a) and (b) show the performance of ST-CoNAL with respect to the number of labeled samples evaluated on CIFAR-10 and CIFAR-100 datasets, respectively. As mentioned earlier, we use the performance difference obtained compared to the classification accuracy of random sampling as a performance indicator. We observe that the proposed ST-CoNAL outperforms other AL methods on both CIFAR-10 and CIFAR-100. After the last acquisition step, our method achieves a performance improvement of 4.68% and 5.47% compared to the random sampling baseline on CIFAR-10 and CIFAR-100 datasets, respectively. It should be noted that ST-CoNAL achieves performance comparable to TA-VAAL, a state-of-the-art method that requires the addition of

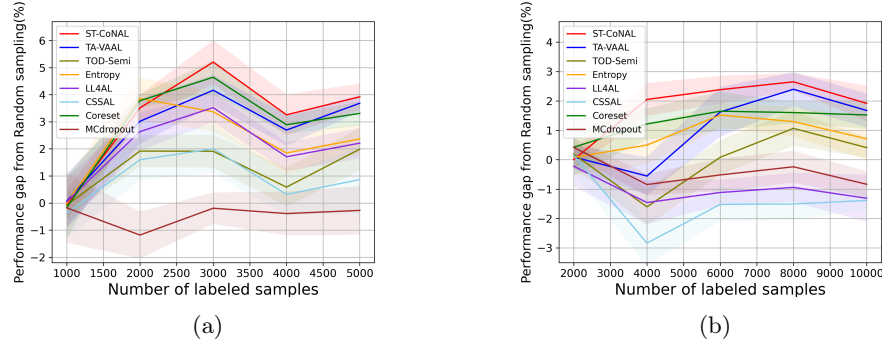


Fig. 3: Average accuracy improvement from random sampling versus the number of labeled samples on (a) Caltech-256 and (b) Tiny ImageNet datasets.

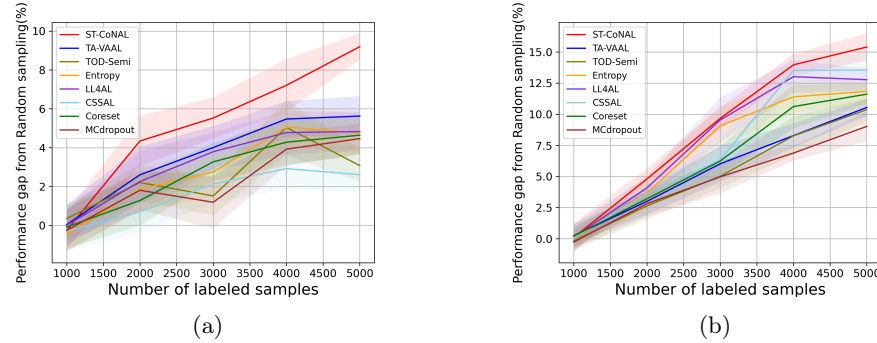


Fig. 4: Average accuracy improvement from random sampling versus the number of labeled samples on (a) the step imbalanced CIFAR-10 and (b) the long-tailed CIFAR-10. The imbalance ratio was set to 100 for both experiments.

subnetworks and complex training optimization. Entropy shows a good performance at the beginning acquisition steps but the performance improves slowly as compared to other methods. After the last acquisition, ST-CoNAL achieves 1.37% and 4.15% higher accuracy than Entropy on CIFAR-10 and CIFAR-100, respectively. Our method outperforms CSSAL, another consistency-based AL method that only considers data uncertainty, not model uncertainty.

Caltech-256 and Tiny ImageNet: Fig. 3 (a) and (b) show the performance of several AL methods on Caltech-256 and Tiny ImageNet, respectively. The ST-CoNAL also achieves significant performance gain on these datasets. After the last acquisition step, ST-CoNAL achieves 3.92% higher accuracy than the random sampling baseline on Caltech-256 while it achieves 1.92% gain on Tiny-Imagenet dataset. Note that the proposed method far outperforms the CSSAL and performs comparable to more sophisticated representation-based methods, Coreset and TA-VAAL.




Fig. 5: Average accuracy improvement from random sampling versus the number of labeled samples on (a) CIFAR-10 and (b) CIFAR-100 dataset. The model was trained by the mean teacher [39], one of SSL methods.

4.3 Performance Comparison on Other Setups

Class-Imbalanced Datasets. We also evaluate ST-CoNAL on class imbalanced datasets. Fig. 4 (a) and (b) present the performance of ST-CoNAL on (a) the step imbalanced CIFAR-10 [20] and (b) the long-tailed CIFAR-10 [5]. We see that ST-CoNAL achieves remarkable performance improvements compared to random sampling on the class imbalanced datasets. Specifically, after the last acquisition step, ST-CoNAL achieves a 9.2% performance improvement on the step-imbalanced CIFAR-10 and a 15.41% improvement on the long-tailed CIFAR-10. Furthermore, ST-CoNAL achieves larger performance gain over the current state-of-the-art, TA-VAAL. It achieves 3.58% and 4.86% performance gains over TA-VAAL on the step imbalanced CIFAR-10 and the long-tailed CIFAR-10, respectively. Note that the consistency-based method CSSAL does not perform well under this class-imbalanced setup.

SSL Setups. Fig. 5 (a) and (b) show the performance evaluated on CIFAR-10 and CIFAR-100 datasets when ST-CoNAL is applied to SSL. ST-CoNAL maintains performance gains even in SSL settings and outperforms other AL methods by a larger margin than in supervised learning settings. After the last acquisition step, ST-CoNAL achieves the 4.87% and 6.66% performance gains over the random sampling baseline on CIFAR-10 and CIFAR-100, respectively. In particular, Fig. 5 (b) shows that ST-CoNAL achieves substantial performance gains over the existing methods in CIFAR-100.

4.4 Performance Analysis

In this subsection, we investigate the benefit of using self-ensemble for an acquisition criterion through some experiments.

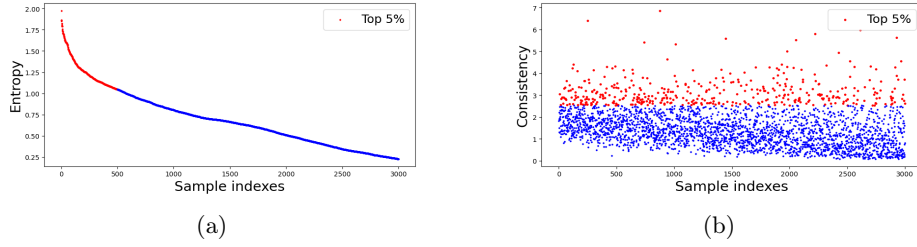


Fig. 6: The entropy and consistency measure are evaluated based on 10k unlabeled samples. The model is trained after 4k samples are acquired on CIFAR-10 dataset. (a) Entropy evaluated for the samples sorted in descending order by entropy and (b) consistency measure evaluated for the samples sorted in the same order. Top 5% of samples are marked in red.

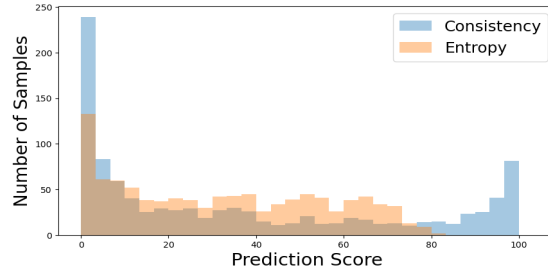


Fig. 7: Histogram of prediction scores produced by the samples selected by the proposed acquisition criterion and entropy-based one.

Comparison Between Consistency Measure versus Uncertainty Measure. To understand the behavior of the proposed consistency measure, we compare it with the popularly used uncertainty measure, Entropy [36]. The sample acquisition process was conducted with the setup described in *Experiment Setup* section. We stopped the acquisition when 4k samples are labeled. Then, we evaluated both the entropy measure and the consistency measure of ST-CoNAL with 10k samples randomly selected from the remaining unlabeled samples. The 10k samples were sorted in descending order according to entropy and sample indexes were assigned in order. Fig. 6 (a) shows the plot of entropy value versus the sample index. Obviously, due to sorting operation, the entropy value decreases with the sample index. Fig. 6 (b) presents the plot of the consistency measure over the samples sorted in the same order. The samples ranked in the top 5% by each measure are marked in red in each figure. We observe that quite different samples were identified in the top 5% according to two criteria. From this, we can conclude that the consistency measure offers the quite different standard for acquisition from the entropy measure.

Fig. 7 shows the distribution of the prediction scores produced by the samples selected by the proposed acquisition criterion and the conventional entropy-

CIFAR-10	1k samples labeled	2k samples labeled	3k samples labeled	4k samples labeled	5k samples labeled
EMA	46.36	60.73	69.10	77.49	81.51
Our EWA	47.13	61.48	71.85	79.61	83.05
CIFAR-100	2k samples labeled	4k samples labeled	6k samples labeled	8k samples labeled	10k samples labeled
EMA	20.34	32.96	42.75	51.77	55.49
Our EWA	20.35	34.70	44.87	52.47	57.49

Table 1: Ablation study on the equally-weighted average (EWA) strategy compared to exponential moving average (EMA) of MT [39].

based one [36]. We use a prediction score that corresponds to the true class index for the given sample. We used the same setup in Fig. 6. Obviously, the entropy-based criterion tends to select the samples with low confidence. It does not select the samples whose prediction score is above 80%. On the contrary, the proposed acquisition criterion selects samples with widely spread prediction scores. Note that it even selects samples for which predictions were made with near 100% confidence because they yielded the low consistency measure. The fact that the proposed method outperforms entropy-based acquisitions indicates that selecting samples with the lowest confidence is not necessarily the best strategy, and consistency-based acquisitions could promise a better solution by leveraging the advantage of various temporal self-ensemble models.

Comparison Between Equally-Weighted Average versus Exponential Moving Average. Table 1 compares two acquisition functions that construct a teacher model using the equally-weighted average (EWA) of the proposed method versus the exponential moving average (EMA) of the MT method. For fair comparison, only the teacher models were differently constructed while all other configurations were set equally for both training and inference. Table 1 shows that EWA achieves higher performance gains over EWA as the acquisition step proceeds on both CIFAR-10 and CIFAR-100 datasets. This shows that the proposed method can construct more effective teacher model for sample acquisition than EMA.

T-SNE Visualization. We also visualize the samples acquired according to the entropy and the proposed consistency measure using T-SNE [28]. T-SNE projects the feature vectors onto the embedding space of lower dimension for visualization. The setup used in Fig. 6 was also used to perform T-SNE. In Fig. 8, the colored dots correspond to the unlabeled sample points projected in the embedding space. The samples selected by the entropy and by the consistency measure are marked by black cross in Fig. 8 (a) and (b), respectively. We observe from Fig. 8 that the samples selected by two acquisition criteria have different distribution, which is consistent with our finding in previous section. Samples selected by entropy tend to gather near the boundary, while samples selected by consistency measure tend to spread further. This shows that more diverse samples are selected according to consistency measure than entropy.

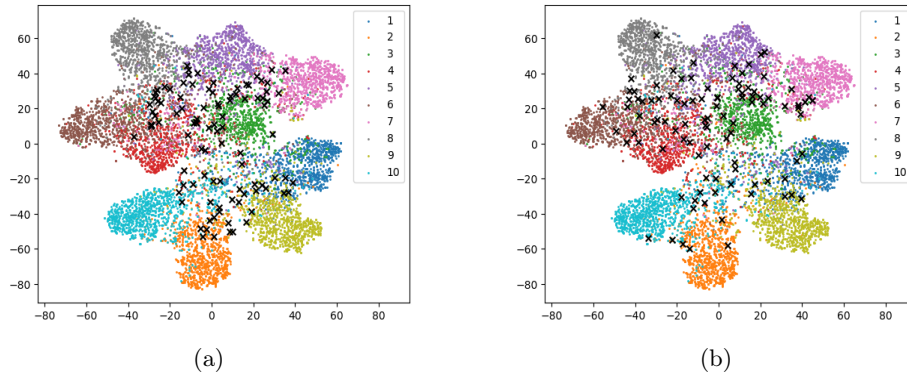


Fig. 8: T-SNE [28] visualizes the feature at the penultimate layer of the teacher model. The model is trained by $4k$ samples on CIFAR-10 dataset. The colored dots correspond to the unlabeled samples assigned to each true class and the black cross marks in (a) and (b) correspond to the top 1% sample selected by the entropy and the consistency measure, respectively.

5 Conclusions

In this study, we proposed a new acquisition function for pool-based AL, which used the temporal self-ensemble generated during the SGD optimization. The proposed sample acquisition criterion measures the discrepancy between the predictions from the student model and the teacher model, where the Q student models are provided as the self-ensemble models and the teacher model is obtained through a weighted average of the student models. Using the consistency measure based on KL divergence, the proposed ST-CoNAL can acquire better samples for AL. We also showed that the sharpening operation applied to the logits of the teacher model can further improve the AL performance. The experiments conducted on CIFAR-10, CIFAR-100, Caltech-256 and Tiny ImageNet datasets demonstrated that ST-CoNAL achieved the significant performance gains over the existing AL methods. Furthermore, our numerical analysis indicated that our ST-CoNAL performs better than other AL baselines in class-imbalanced and semi-supervised learning scenarios.

Acknowledgments This work has supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT)(No.2020R1A2C2012146) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT)(No.2020-0-01373, Artificial Intelligence Graduate School Program (Hanyang University)).

References

1. Athiwaratkun, B., Finzi, M., Izmailov, P., Wilson, A.G.: There are many consistent explanations of unlabeled data: Why you should average. In: International Conference on Learning Representations (ICLR) (2019)

2. Beluch, W.H., Genewein, T., Nürnberger, A., Köhler, J.M.: The power of ensembles for active learning in image classification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9368–9377 (2018)
3. Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. In: *International Conference on Learning Representations (ICLR)* (2019)
4. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 5049–5059 (2019)
5. Cao, K., Wei, C., Gaidon, A., Arechiga, N., Ma, T.: Learning imbalanced datasets with label-distribution-aware margin loss. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 32 (2019)
6. Chen, C., Dong, S., Tian, Y., Cao, K., Liu, L., Guo, Y.: Temporal self-ensembling teacher for semi-supervised object detection. *IEEE Transactions on Multimedia* (2021)
7. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9268–9277 (2019)
8. Drugman, T., Pytköinen, J., Kneser, R.: Active and semi-supervised learning in ASR: Benefits on the acoustic and language models. In: *Interspeech*. pp. 2318–2322 (2016)
9. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In: *Proceedings of the International Conference on Machine Learning (ICML)*. pp. 1050–1059 (2016)
10. Gal, Y., Islam, R., Ghahramani, Z.: Deep Bayesian active learning with image data. In: *Proceedings of the International Conference on Machine Learning (ICML)*. vol. 70, pp. 1183–1192 (2017)
11. Gao, M., Zhang, Z., Yu, G., Arık, S.Ö., Davis, L.S., Pfister, T.: Consistency-based semi-supervised active learning: Towards minimizing labeling cost. In: *European Conference on Computer Vision (ECCV)*. pp. 510–526 (2020)
12. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 529–536 (2005)
13. Griffin, G., Holub, A., Perona, P.: Caltech-256 object category dataset. In: *California Institute of Technology* (2007)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
15. Houlshby, N., Ferenc, H., Zoubin, G., Lengyel, M.: Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745* (2011)
16. Huang, S., Wang, T., Xiong, H., Wen, B., Huan, J., Dou, D.: Temporal output discrepancy for loss estimation-based active learning. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
17. Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., Wilson, A.G.: Averaging weights leads to wider optima and better generalization. In: *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI)*. pp. 876–885 (2018)
18. Johnson, E.H.: Elementary applied statistics: For students in behavioral science. *Social Forces* **44**(3), 455–456 (1966)

19. Joshi, A.J., Porikli, F., Papanikolopoulos, N.: Multi-class active learning for image classification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2372–2379 (2009)
20. Kim, K., Park, D., Kim, K.I., Chun, S.Y.: Task-aware variational adversarial active learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8166–8175 (2021)
21. Krizhevsky, A.: Learning multiple layers of features from tiny images. Technical report, Department of Computer Science, University of Toronto (2009)
22. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: *International Conference on Learning Representations (ICLR)* (2017)
23. Le, Y., Yang, X.: Tiny imagenet visual recognition challenge. *CS 231N* **7**(7), 3 (2015)
24. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: *Proceedings of the International Conference on Machine Learning Workshop on Challenges in Representation Learning*. vol. 3, p. 2 (2013)
25. Lee, J., Chung, S.Y.: Robust training with ensemble consensus (2020)
26. Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: *SIGIR '94*. pp. 3–12. Springer (1994)
27. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. In: *International Conference on Learning Representations (ICLR)* (2017)
28. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* **9**(Nov), 2579–2605 (2008)
29. Maddox, W.J., Izmailov, P., Garipov, T., Vetrov, D.P., Wilson, A.G.: A simple baseline for Bayesian uncertainty in deep learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 7026–7037 (2019)
30. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**, 1979–1993 (2018)
31. Oliver, A., Odena, A., Raffel, C.A., Cubuk, E.D., Goodfellow, I.: Realistic evaluation of deep semi-supervised learning algorithms. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 3235–3246 (2018)
32. Rhee, P., Erdenee, E., Shin, D.K., Ahmed, M., Jin, S.: Active and semi-supervised learning for object detection with imperfect data. *Cognitive Systems Research* **45**, 109–123 (2017)
33. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 1163–1171 (2016)
34. Sener, O., Savarese, S.: Active learning for convolutional neural networks: A core-set approach. In: *International Conference on Learning Representations (ICLR)* (2018)
35. Settles, B.: Active learning literature survey. Tech. rep., University of Wisconsin-Madison Department of Computer Sciences (2009)
36. Shannon, C.E.: A mathematical theory of communication. *Bell system technical journal* **27**(3), 379–423 (1948)
37. Sinha, S., Ebrahimi, S., Darrell, T.: Variational adversarial active learning. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 5972–5981 (2019)
38. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with

- consistency and confidence. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020)
39. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 1195–1204 (2017)
 40. Tomanek, K., Hahn, U.: Semi-supervised active learning for sequence labeling. In: *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*. vol. 2, pp. 1039–1047. Association for Computational Linguistics (2009)
 41. Tong, S., Koller, D.: Support vector machine active learning with applications to text classification. *The Journal of Machine Learning Research* **2**, 45–66 (2001)
 42. Verma, V., Lamb, A., Kannala, J., Bengio, Y., Lopez-Paz, D.: Interpolation consistency training for semi-supervised learning. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. pp. 3635–3641. IJCAI.org (2019)
 43. Xie, Q., Dai, Z., Hovy, E., Luong, T., Le, Q.: Unsupervised data augmentation for consistency training. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 33, pp. 6256–6268 (2020)
 44. Yang, G., Tianyi, C., Kirichenko, P., Wilson, A.G., De Sa, C.: Swalp: Stochastic weight averaging in low precision training. In: *Proceedings of the International Conference on Machine Learning (ICML)* (2019)
 45. Yoo, D., Kweon, I.S.: Learning loss for active learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 93–102 (2019)
 46. Zhang, B., Li, L., Yang, S., Wang, S., Zha, Z.J., Huang, Q.: State-relabeling adversarial active learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8756–8765 (2020)
 47. Zhu, X., Lafferty, J.D., Ghahramani, Z.: Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In: *Proceedings of the International Conference on Machine Learning Workshop on the Continuum from Labeled to Unlabeled Data*. pp. 58–65 (2003)