

# Semi-Supervised Active Learning with Temporal Output Discrepancy

Siyu Huang<sup>1</sup> Tianyang Wang<sup>2</sup> Haoyi Xiong<sup>1</sup> Jun Huan<sup>3</sup> Dejing Dou<sup>1</sup>  
<sup>1</sup>Baidu Research <sup>2</sup>Austin Peay State University <sup>3</sup>Styling AI

{huangsiyu, xionghaoyi, doudejing}@baidu.com wangt@apsu.edu lukehuan@shenshangtech.com

## Abstract

While deep learning succeeds in a wide range of tasks, it highly depends on the massive collection of annotated data which is expensive and time-consuming. To lower the cost of data annotation, active learning has been proposed to interactively query an oracle to annotate a small proportion of informative samples in an unlabeled dataset. Inspired by the fact that the samples with higher loss are usually more informative to the model than the samples with lower loss, in this paper we present a novel deep active learning approach that queries the oracle for data annotation when the unlabeled sample is believed to incorporate high loss. The core of our approach is a measurement Temporal Output Discrepancy (TOD) that estimates the sample loss by evaluating the discrepancy of outputs given by models at different optimization steps. Our theoretical investigation shows that TOD lower-bounds the accumulated sample loss thus it can be used to select informative unlabeled samples. On basis of TOD, we further develop an effective unlabeled data sampling strategy as well as an unsupervised learning criterion that enhances model performance by incorporating the unlabeled data. Due to the simplicity of TOD, our active learning approach is efficient, flexible, and task-agnostic. Extensive experimental results demonstrate that our approach achieves superior performances than the state-of-the-art active learning methods on image classification and semantic segmentation tasks.

## 1. Introduction

Large-scale annotated datasets are indispensable and critical to the success of modern deep learning models. Since the annotated data is often highly expensive to obtain, learning techniques including unsupervised learning [6], semi-supervised learning [59], and weakly supervised learning [44] have been widely explored to alleviate the dilemma. In this paper<sup>1</sup> we focus on active learning [4] which aims to selectively annotate unlabeled data with lim-

ited budgets while resulting in high performance models.

In existing literature of active learning, two mainstream approaches have been studied, namely the diversity-aware approach and the uncertainty-aware approach. The diversity-aware approach [15] aims to pick out diverse samples to represent the distribution of a dataset. It works well on low-dimensional data and classifier with a small number of classes [40]. The uncertainty-aware approach [52] aims to pick out the most uncertain samples based on the current model. However, the uncertainty heuristics, such as distance to decision boundary [2] and entropy of posterior probabilities [41], are often task-specific and need to be specifically designed for individual tasks such as image classification [19], object detection [50], and semantic segmentation [8].

In this paper, we consider that the samples with higher loss would be more informative than the ones with lower loss. Specifically in supervised learning settings, when samples are correctly labeled, the averaged loss function over all samples should be gradually minimized during the learning procedure. Moreover, in every iteration the training model would backward propagated error according to the loss of every sample [28], while the sample with high loss usually brings informative updates to the parameters of the training model [16]. In this work, we generalize these evidences to active learning problems and propose a simple yet effective loss estimator Temporal Output Discrepancy (TOD), which could measure the potential loss of a sample only relied on the training model, when the ground-truth label of the sample is not available. Specifically, TOD computes the discrepancy of outputs given by models at different optimization steps, and a higher discrepancy corresponds to a higher sample loss. Our theoretical investigation shows that TOD well measures the sample loss.

On basis of TOD, we propose a deep active learning framework that leverages a novel unlabeled data sampling strategy for data annotation in conjunction with a semi-supervised training scheme to boost the task model performance with unlabeled data. Specifically, the active learning procedure can be split into a sequence of training cycles starting with a small number of labeled samples. By the

<sup>1</sup>Code is available at <https://github.com/siyuhuang/TOD>

end of every training cycle, our data sampling strategy estimates Cyclic Output Discrepancy (COD), which is a variant of TOD, for every sample in the unlabeled pool and selects the unlabeled samples with the largest COD for data annotation. The newly-annotated samples are added to the labeled pool for model training in the next cycles. Furthermore, with the aid of the unlabeled samples, we augment the task learning objective with a regularization term derived from TOD, so as to improve the performance of active learning in a semi-supervised manner.

Compared with the existing deep active learning algorithms, our approach is more efficient, more flexible, and easier to implement, since it does not introduce extra learnable models such as the loss prediction module [54] or the adversarial network [43, 57] for uncertainty estimation. In the experiments, our active learning approach shows superior performances in comparison with the state-of-the-art baselines on various image classification and semantic segmentation datasets. Further ablation studies demonstrate that our proposed TOD can well estimate the sample loss and benefit both the active data sampling and the task model learning.

The contributions of this paper are summarized as follows.

1. This paper proposes a simple yet effective loss measure TOD. Both theoretical and empirical studies validate the efficacy of TOD.
2. This paper presents a novel deep active learning framework by incorporating TOD into an active sampling strategy and a semi-supervised learning scheme.
3. Extensive active learning experiments on image classification and semantic segmentation tasks evaluate the effectiveness of the proposed methods.

## 2. Related Work

**Active Learning.** Active learning aims to incrementally annotate samples that result in high model performance and low annotation cost [4]. Active learning has been studied for decades of years and the existing methods can be generally grouped into two categories: the query-synthesizing approach and the query-acquiring approach. The query-synthesizing approach [60, 32] employs generative models to synthesize new informative samples. For instance, ASAL [33] uses generative adversarial networks (GANs) [14] to generate high-entropy samples. In this paper, we focus on the query-acquiring active learning which develops effective data sampling strategies to pick out the most informative samples from the unlabeled data pool.

The query-acquiring methods can be categorized as diversity-aware and uncertainty-aware methods. The

diversity-aware methods [36, 15] select a set of diverse samples that best represents the dataset distribution. A typical diversity-aware method is the core-set selection [40] based on the core-set distance of intermediate features. It is theoretically and empirically proven to work well with a small scale of classes and data dimensions.

The uncertainty-aware methods [21, 52, 12, 13, 9] actively select the most uncertain samples in the context of the training model. A wide variety of related methods has been proposed, such as Monte Carlo estimation of expected error reduction [39], distance to the decision boundary [47, 2], margin between posterior probabilities [38], and entropy of posterior probabilities [41, 19, 31].

The diversity-aware and uncertainty-aware approaches are complementary to each other thus many hybrid methods [29, 51, 53, 58, 30, 26] have been proposed for specific tasks. In more recent literature, adversarial active learning [7, 43, 57] is introduced to learn an adversarial discriminator to distinguish the labeled and unlabeled data.

Compared to the existing works in active learning, our method falls into the category of uncertainty-aware active learning by directly utilizing the task model for uncertainty estimation. The relevant works include the ones which utilize the expected gradient length [42] or output changes on input perturbation [11, 20] for uncertainty estimation. In the realm of loss estimation, Yoo *et al.* [54] propose to learn a loss prediction module to estimate the loss of unlabeled samples. Different from existing methods which require extra deep models such as loss prediction network [54] or adversarial network [43, 57] for uncertainty estimation, we propose a learning-free principle for efficient active learning by evaluating the discrepancy of model outputs at different active learning cycles. Except its efficiency and task-agnostic property, we demonstrate that it is a lower bound of the accumulated sample loss, ensuring that the data samples of loss of higher lower bound can be picked out.

**Semi-Supervised Learning.** This work is also related to semi-supervised learning which seeks to learn from both labeled and unlabeled data, since we also develop the proposed loss estimation method to improve the learning of task model using unlabeled data. There has been a wide variety of semi-supervised learning approaches such as transductive model [18], graph-based method [56], and generative model [23]. We refer to [48] for an up-to-date overview.

More recently, several semi-supervised methods including  $\Pi$ -model [27] and Virtual Adversarial Training [34] apply consistency regularization to the posterior distributions of perturbed inputs. Further improvements including Mean Teacher [46] and Temporal Ensembling [27] apply the consistency regularization on models at different time steps. However, the consistency regularization has been seldom exploited for active learning.

Compared to the existing efforts in semi-supervised

learning for neural networks, our proposed loss measure TOD could be considered as an alternative solution of consistency regularization. TOD can be well adapted to active learning by developing a novel active sampling method COD. COD only relies on the models learned after every active learning cycle. In contrast, the existing temporal consistency-based uncertainty measurements often require access to a number of previous model states. For instance, the computing of Mean Teacher [46] and Temporal Ensembling [27] require the historical model parameters and the historical model outputs, respectively.

On the other hand, there have not been sufficient theoretical interpretations for the success of consistency regularization. Athiwaratkun *et al.* [1] reveals that the consistency regularization on perturbed inputs is an unbiased estimator for the norm of the Jacobian of the network. However, there is still a lack of interpretations on the temporal consistency regularization. In this paper, we show that the temporal consistency regularization can be connected to the lower bound of the accumulated sample loss. Thus, the temporal consistency regularization is an theoretically effective solution to loss estimation as well as semi-supervised learning.

### 3. Temporal Output Discrepancy

Measuring the sample loss on a given neural network  $f$ , when the label of the sample is unavailable, is a key challenge for many learning problems, including active learning [38, 31, 13], continual learning [9], and self-supervised learning [46, 27]. In this work, we present Temporal Output Discrepancy (TOD), which estimates the sample loss based on the discrepancy of outputs of a neural network at different learning iterations. Given a sample  $x \in \mathbb{R}^d$ , we have  $\text{TOD } D_t^{\{T\}} : \mathbb{R}^d \rightarrow \mathbb{R}$

$$D_t^{\{T\}}(x) \stackrel{\text{def}}{=} \|f(x; w_{t+T}) - f(x; w_t)\|. \quad (1)$$

$D_t^{\{T\}}(x)$  characterizes the distance<sup>2</sup> between outputs of model  $f$  with parameters  $w_{t+T}$  and  $w_t$  obtained in the  $(t+T)$ -th and  $t$ -th gradient descend step during learning (e.g.,  $T > 0$ ), respectively.

In the following, we show that a larger  $D_t^{\{T\}}(x)$  indicates a larger sample loss<sup>3</sup>  $L_t(x) = \frac{1}{2}(y - f(x; w_t))^2$ , where  $y \in \mathcal{R}$  is the label corresponding to sample  $x$ . We first give the upper bound of one-step output discrepancy  $D_t^{\{1\}}(x)$ .

**Theorem 1** *With an appropriate setting of learning rate  $\eta$ ,*

$$D_t^{\{1\}}(x) \leq \eta \sqrt{2L_t(x)} \|\nabla_w f(x; w_t)\|^2. \quad (2)$$

<sup>2</sup>For brevity,  $\|\cdot\|$  denotes the  $L_2$  norm  $\|\cdot\|_2$  in this paper.

<sup>3</sup>Here we take Euclidean loss as an example. The cross-entropy loss has similar results.

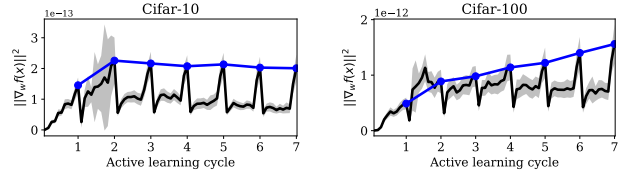


Figure 1.  $\|\nabla_w f\|^2$ , vs., the active learning cycle. The dark lines denote  $\|\nabla_w f\|^2$  averaged over the training (i.e., labeled and unlabeled) pool. The blue lines denote average  $\|\nabla_w f\|^2$  after every active learning cycle.

The proofs of Theorem 1 and the following corollaries can be found in the supplementary material. From Theorem 1, the upper bound of  $T$ -step output discrepancy  $D_t^{\{T\}}(x)$  can be easily deduced.

**Corollary 1** *With an appropriate setting of learning rate  $\eta$ ,*

$$D_t^{\{T\}}(x) \leq \sqrt{2}\eta \sum_{\tau=t}^{t+T-1} \left( \sqrt{L_\tau(x)} \|\nabla_w f(x; w_\tau)\|^2 \right). \quad (3)$$

Corollary 1 preliminarily connects  $T$ -step output discrepancy  $D_t^{\{T\}}(x)$  to sample loss  $L(x)$ . However, it is almost infeasible to compute  $\|\nabla_w f(x; w_\tau)\|$  on all the  $\tau$ . Fortunately,  $\|\nabla_w f\|$  is approximately a constant under the context of neural networks, as discussed in [45, 49].

**Remark 1** *For a linear layer  $\phi(x; W)$  with ReLU activation, the Lipschitz constant  $\mathcal{L}(W) \leq \|x\|$ .*

Since sample  $x$  is drawn from a distribution  $X$ , we assume  $\|x\|$  is upper-bounded by a constant so that  $f$  is Lipschitz-continuous over  $w$ . Thus, we let  $\|\nabla_w f\|^2$  be upper-bounded by a constant  $C$ . Empirical results on image classification benchmarks including Cifar-10 and Cifar-100 also support this assumption. As shown in Fig. 1, the dark lines are the averaged  $\|\nabla_w f\|^2$  over the training set. The blue lines denote the averaged  $\|\nabla_w f\|^2$  after every active learning cycle.  $\|\nabla_w f\|^2$  has a small variance over samples and it is nearly constant across every active learning cycle.

With  $\|\nabla_w f\|^2 \leq C$ , we rewrite Corollary 1 to connect  $D_t^{\{T\}}(x)$  with the accumulated loss of sample  $x$ .

**Corollary 2** *With appropriate settings of a learning rate  $\eta$  and a constant  $C$ ,*

$$D_t^{\{T\}}(x) \leq \sqrt{2T}\eta C \sqrt{\sum_{\tau=t}^{t+T-1} L_\tau(x)}. \quad (4)$$

Corollary 2 shows that  $\|f(x; w_{t+T}) - f(x; w_t)\|$  is a lower bound of the square root of accumulated loss  $L$  during  $T$  gradient descend steps. Thus, when  $T$  is fixed, e.g., a certain number of iterations of neural network training, TOD can effectively estimate the loss of sample  $x$ . Note

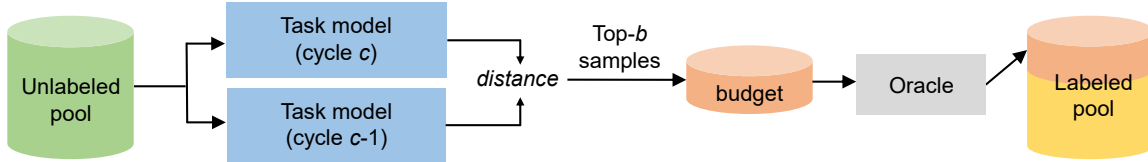


Figure 2. The COD-based unlabeled data sampling strategy for active learning. Data samples with the largest COD are collected from the unlabeled pool. The collected samples are annotated by an oracle and added to the labeled pool.

that the pre-assumptions of Theorem 1 and its corollaries limit the learning rate  $\eta$  not to be too large to dissatisfy the Taylor expansion used in our proofs. In empirical study we find that the commonly used learning rates, *e.g.*,  $\eta=0.1$  or smaller, work well.

## 4. Semi-Supervised Active Learning

### 4.1. Problem Formulation

We first formulate the standard active learning task as follows. Let  $(x_S, y_S)$  denote a sample pair drawn from a set of labeled data  $(X_S, Y_S)$ , where  $X_S$  is the data points and  $Y_S$  is the labels. Let  $x_U$  denote an unlabeled sample drawn from a larger unlabeled data pool  $X_U$ , *i.e.*, the labels  $Y_U$  corresponding to  $X_U$  cannot be observed. In an active learning cycle  $c$ , the active learning algorithm selects a fixed budget of samples from the unlabeled pool  $X_U$  and the selected samples will be annotated by an oracle. The budget size  $b$  is usually much smaller than  $|X_U|$ , the size of the unlabeled pool. The goal of active learning is to select the most informative unlabeled samples for annotation, so as to minimize the expected loss of a task model  $f : X \rightarrow Y$ .

We next present the use of TOD in a *semi-supervised* active learning framework. An active learning algorithm generally consists of two components: (a) an unlabeled data sampling strategy and (b) the learning of a task model. We adapt TOD to these two components, respectively. For component (a), we propose Cyclic Output Discrepancy (COD), a new criterion to select unlabeled samples with the largest *estimated loss* for annotation. For component (b), we develop a TOD-based unsupervised loss term to improve the performance of task model. In the following, we formulate the active learning problem and discuss the details of the two components.

### 4.2. Cyclic Output Discrepancy

In Eq. 4, our proposed TOD characterizes a lower bound of the loss function for supervised learning. Here we introduce a variant of TOD, *i.e.*, Cyclic Output Discrepancy (COD), for active selection of unlabeled samples. COD estimates the sample uncertainty by measuring the difference of model outputs between two consecutive active learning cycles,

$$D_{\text{cyclic}}(x|w_c, w_{c-1}) = \|f(x; w_c) - f(x; w_{c-1})\|, \quad (5)$$

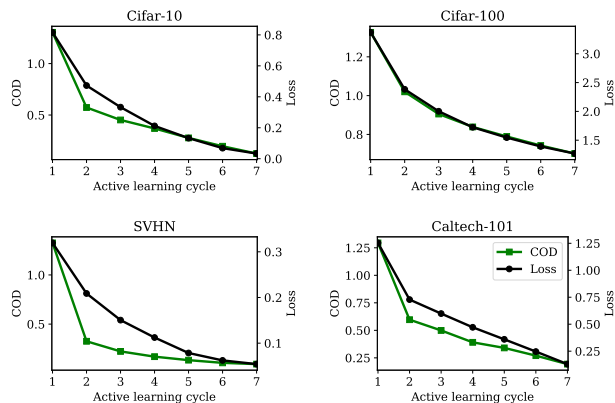


Figure 3. The consistency of Cyclic Output Discrepancy (COD) and the real task loss. We show COD and real loss averaged over the unlabeled samples, *vs.*, percentage of labeled images, under active learning setting.

where model parameters  $w_c$  and  $w_{c-1}$  are obtained after the  $c$ -th and  $(c-1)$ -th active learning cycle, respectively.

Fig. 2 illustrates the procedure of COD-based unlabeled data sampling. Given COD for every sample in unlabeled pool  $X_U$ , our strategy selects  $b$  samples with the largest COD from  $X_U$ . Then, the strategy queries human oracles for annotating the selected samples. The newly-annotated data is added to labeled pool for the next active learning cycle. In the first cycle (*i.e.*,  $c=1$ ), the model  $f$  is trained with a random subset of labeled data, and COD is computed based on the initial model and the model learned after the first cycle. For  $c \geq 2$ , we compute COD  $D_{\text{cyclic}}(x|w_c, w_{c-1})$  for active sample selection systematically.

**Minimax optimization of COD.** As discussed in Corollary 2, COD-based data sampling strategy can find samples of large loss in unlabeled pool, so as to minimize the expected loss of model  $f$  through further training the task model in the next cycle. Fig. 3 preliminarily verifies the consistency between COD and the real loss, where COD shows a similar trend with the real loss and they are both decreasing along with the active learning progress. Instead of minimizing the TOD directly (*i.e.*, min-min optimization which may not be good), COD develops TOD as the criterion of sample selection in active learning, where samples with the maximal TOD are picked up (*e.g.*, max-min strategies). When

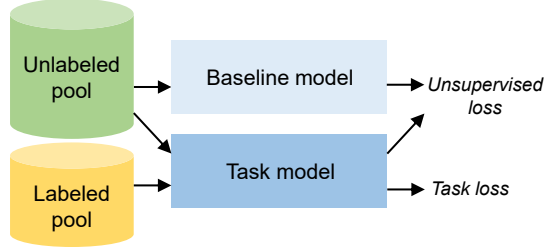


Figure 4. Semi-supervised task learning scheme. For labeled data, the task model is trained with the task loss. For unlabeled data, the task model is trained to minimize the distance between outputs of the task model and the baseline model.

considering labels of samples with potential losses as information gain, our strategy actually maximizes the minimum gain in active learning.

### 4.3. Semi-Supervised Task Learning

**Unsupervised loss.** As suggested by Corollary 2, TOD measures the accumulated sample loss, and thus it is natural to employ TOD as an unsupervised criterion to improve the learning of model  $f$  using the unlabeled data. However, directly applying TOD to unsupervised training with the baseline model obtained at the last cycle  $c - 1$  may lead to an unstable training, due to the following aspects: 1) The iteration interval between current model and baseline model (*i.e.*,  $T$  in Corollary 2) is no longer fixed during model training, thus the loss measurement would be inaccurate; 2) The baseline model only depends on a single historical model state so that it may suffer from a large variance in loss measurement. To address the above issues, we are inspired by Mean Teacher [46] to construct a baseline model by applying an exponential moving average (EMA) to the historical parameters, as

$$\tilde{w} \leftarrow \alpha \cdot \tilde{w} + (1 - \alpha) \cdot w. \quad (6)$$

where  $\tilde{w}$  and  $w$  are parameters of baseline model and current model, respectively, and  $\alpha$  is the EMA decay rate.

Our unsupervised loss minimizes the distance between the current model and the baseline model. In the  $c$ -th cycle, with the unlabeled pool  $X_U^c$ , the unsupervised loss is

$$L_U^c(w) = \frac{1}{|X_U^c|} \sum_{x_U \in X_U^c} \|f(x_U; w) - f(x_U; \tilde{w})\|^2. \quad (7)$$

**Task loss.** For the labeled data, we optimize a supervised task objective. Here we take the cross-entropy (CE) loss for image classification as an example. In the  $c$ -th cycle, given the labeled set  $(X_S^c, Y_S^c)$  in the cycle, the supervised loss is

$$L_S^c(w) = \frac{1}{|X_S^c|} \sum_{(x_S, y_S) \in (X_S^c, Y_S^c)} \text{CE}[f(x_S; w), y_S]. \quad (8)$$

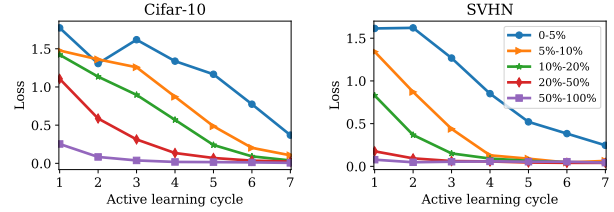


Figure 5. The average real losses of unlabeled samples in a descending order of COD values. For instance, “0-5%” denotes the 5% unlabeled samples which have the largest COD values, and so on.

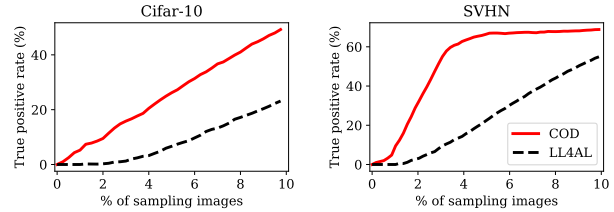


Figure 6. The performance of loss estimation using a learned loss prediction model (LL4AL) [54] and the proposed COD method. We show the proportion of sampled images which have the highest real losses, *vs.*, the proportion of sampling images.

Note that the labeled pool  $(X_S^c, Y_S^c)$  will be enlarged per active learning cycle. Within an active learning cycle, the labeled pool remains unchanged.

**Overall objective.** Our semi-supervised task learning scheme is illustrated in Fig. 4. By integrating the task and unsupervised losses, we minimize an overall learning objective that evolves with the cycle  $c$ , as

$$L_{\text{overall}}^c(w) = L_S^c(w) + \lambda \cdot L_U^c(w), \quad (9)$$

where  $\lambda$  is a trade-off weight to balance the task and unsupervised loss terms. In our experiments,  $\lambda$  is set to 0.05 and the EMA decay rate  $\alpha$  is set to 0.999. See supplementary material for more details.

## 5. Experiments

We conduct extensive experimental studies to evaluate the proposed active learning approach on two computer vision tasks, image classification and semantic segmentation, with five benchmark datasets. The results are reported over 3 runs with different initial network weights and labeled pools. We implement the methods using PyTorch framework [37]. See supplementary material for more details.

### 5.1. Efficacy of TOD as Loss Measure

This work proposes TOD to estimate the loss of an unlabeled sample. Fig. 3 has evaluated the relations between TOD and sample loss as discussed in Theorem 1 and Corollary 2, suggesting that the average COD and average loss have a consistent trend along with active learning cycles. To

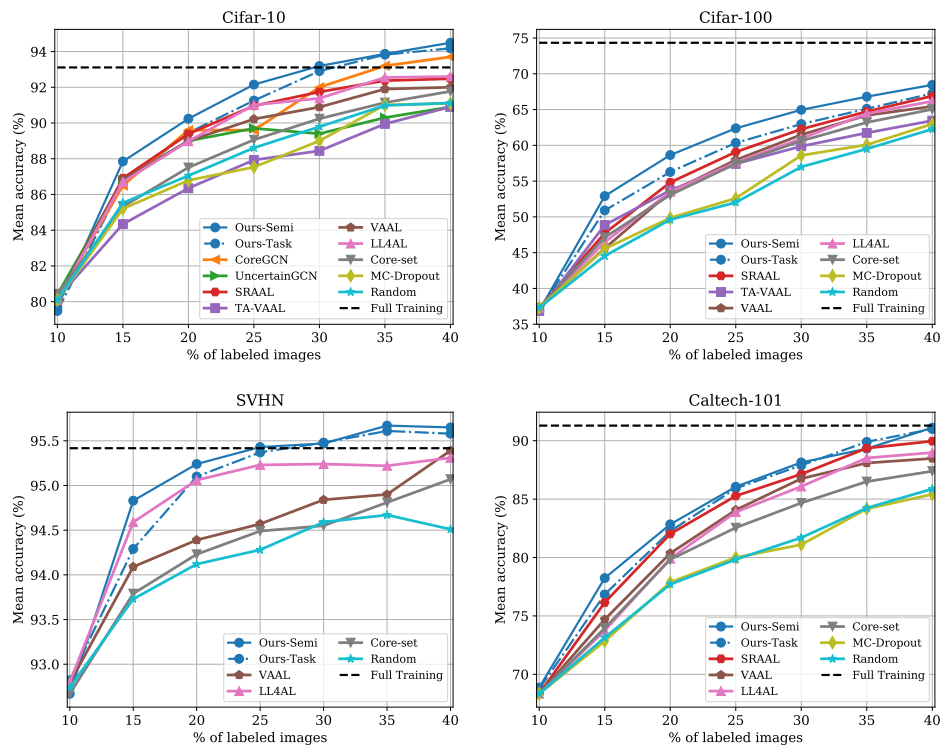


Figure 7. Active learning results of image classification on four benchmark datasets.

further verify the effectiveness of TOD for loss estimation, we study the average loss of unlabeled samples by sorting their COD values. Fig. 5 shows that the larger COD values of samples indicate the higher losses of samples, and this observation is consistent across all the active learning cycles.

In Fig. 6, we compare the loss estimation performance of a learned loss prediction model (LL4AL) [54] and COD. We investigate how many samples of the highest losses can be picked out by using different methods. Fig. 6 shows that COD performs significantly better than LL4AL, as COD is able to pick out more high-loss samples under all the sampling settings. Figs. 3, 5 and 6 demonstrate that COD is an effective loss measure as well as a feasible criterion for active data sampling.

## 5.2. Active Learning for Image Classification

**Experimental setup.** We evaluate active learning methods on four benchmark image classification datasets including Cifar-10 [25], Cifar-100 [25], SVHN [35], and Caltech-101 [10]. Following the conventional practices in deep active learning [54, 43], we employ ResNet-18 [17] as the image classification model. We compare our active learning approach against the state-of-the-art methods including CoreGCN [3], UncertainGCN [3], SRAAL [57], TA-VAAL [22], VAAL [43], LL4AL [54], Core-set [40], and MC-Dropout [12]. In addition, the random selection of unlabeled data (“Random”) and the model trained on the full

training set (“Full Training”) are also included as baselines. “Ours-Semi” indicates our approach trained with the semi-supervised loss and “Ours-Task” is our approach trained with only the task loss.

**Results.** Fig. 7 shows image classification performances of different active learning methods. Our method outperforms all the other methods on the benchmark datasets. Additionally, we have the following observations. (i) Our method consistently performs better than the other methods with respect to the cycles. This is a desired property for a successful active learning method, since the labeling budget may vary for different tasks in real-world applications. For instance, one may only be able to annotate 20% instead of 40% of all the data. (ii) Our method shows robust performances on difficult datasets such as Cifar-100 and Caltech-101. Both datasets include much more classes than Cifar-10, and, Caltech-101 includes images of much higher resolution (*i.e.*,  $300 \times 200$ ). These difficult datasets bring more challenges to active learning, and the superior performances on these datasets demonstrate the robustness of our method. (iii) The performance curves of our method are relatively smooth compared with the other methods. A smooth curve means there are consistent performance improvements from cycle to cycle, indicating that our sampling strategy can take informative data from the unlabeled pool. (iv) Ours-Semi performs better than Ours-Task, demonstrating that our semi-supervised training successfully utilizes the unlabeled data. (v) Our method uses 40% training samples to

Table 1. The active learning performances on 40% labeled data. ‘Base’: standard task model training without active data selection. ‘Semi’: the proposed semi-supervised task learning. ‘Active’: the proposed active data selection strategy.

Dataset	Core-set	LL4AL	VAAL	SRAAL	Base	Base+Semi	Base+Active	Base+Semi+Active
Cifar10	91.8	94.1	92.0	92.5	91.8	92.2 (+0.4)	94.2 (+2.4)	94.5 (+2.7)
Cifar100	65.0	65.2	65.4	66.2	62.3	66.1 (+3.8)	67.3 (+5.0)	68.5 (+6.2)

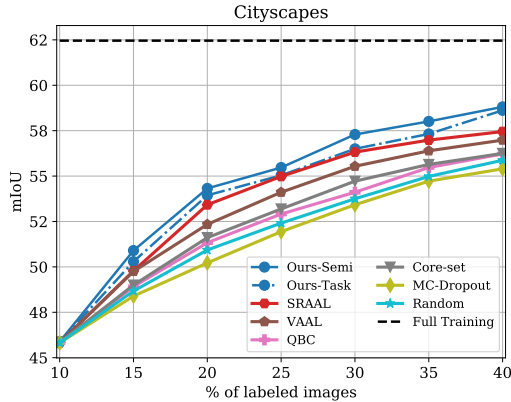


Figure 8. Active learning results of semantic segmentation on Cityscapes dataset.

outperform the full training on Cifar-10 and SVHN, *e.g.*, 94.5% vs. 93.1% on Cifar-10. This interesting finding is in accord with the observations discussed in previous literature [24] that some data in the original dataset might be unnecessary or harmful to model training.

Table 1 compares different active learning methods, *i.e.*, the state-of-the-art algorithms and the proposed methods, for image classification on 40% training labeled data. Both semi-supervised task learning and active data selection strategy contributes to performance improvement, while, active data selection results in a more significant improvement than semi-supervised task learning. We also note that the proposed method can outperform existing algorithms without semi-supervised task learning (see ‘Base+Active’ in Table 1).

### 5.3. Active Learning for Semantic Segmentation

**Experimental setup.** To validate the active learning performance on more complex and large-scale scenarios, we study the semantic segmentation task with the Cityscapes dataset [5] which is a large-scale driving video dataset collected from urban street scenes. Semantic segmentation addresses the pixel-level classification task and its annotation cost is much higher. Following the settings in [43, 57], we employ the 22-layer dilated residual network (DRN-D-22) [55] as the semantic segmentation model. We report the mean Intersection over Union (mIoU) on the validation set of Cityscapes. We compare our method against SRAAL [57], VAAL [43], QBC [26], Core-set [40], MC-Dropout [12], and the random selection.

**Results.** Fig. 8 shows the semantic segmentation performances of different active learning methods on Cityscapes.

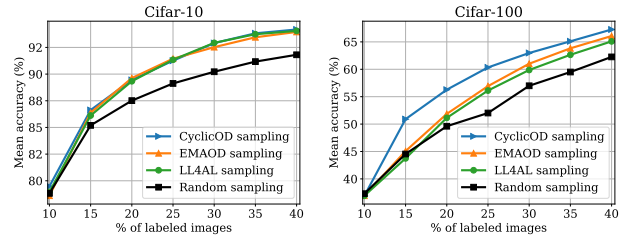


Figure 9. Ablation on active data sampling.

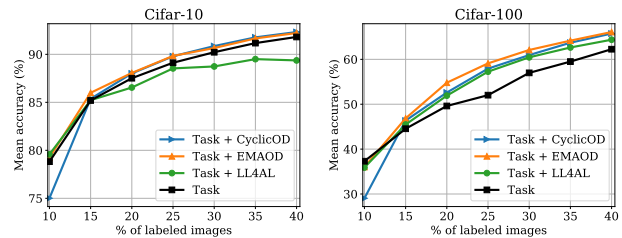


Figure 10. Ablation on semi-supervised task learning.

Both Ours-Semi and Ours-Task outperform the other baselines in terms of mIoU. The results demonstrate the competence of our approach on the challenging semantic segmentation task. Note that in our approach, neither task model training nor data sampling needs to exploit extra domain knowledge. Therefore, our approach is independent of tasks. Moreover, the image size of Cityscapes (*i.e.*,  $2048 \times 1024$ ) is much larger than that of the classification benchmarks, indicating that our method is not sensitive to the data complexity. These advantages make our approach a competitive candidate for complex real-world applications.

### 5.4. Ablation Study

**Active data sampling strategy.** Fig. 9 compares different active data sampling strategies on Cifar-10 and Cifar-100. CyclicOD and EMAOD are two variants of TOD, where CyclicOD employs the model at the end of last cycle as the baseline model while EMAOD employs an exponential moving average of the previous models as the baseline model. LL4AL [54] uses a learned loss prediction module to sample the unlabeled data. Fig. 9 shows that the proposed sampling strategies, *i.e.*, EMAOD and CyclicOD, outperform random sampling and LL4AL sampling on both datasets, validating the effectiveness of TOD-based sampling strategy. CyclicOD performs better than EMAOD on Cifar-100, thus we employ COD as our sampling strategy in the rest of the experiments.

**Semi-supervised task learning.** To evaluate the necessity of semi-supervised task learning in active learning, Fig. 10

Table 2. The class-wise performances on Cityscapes, where 40% labeled data is used for training. ‘Proportion’ denotes the proportions of classes at pixel level. ‘T’ is the model trained using only the task loss. ‘T + U’ is the model trained using both the task loss and the proposed TOD-based unsupervised loss.

Class ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	Ave
Proportion (%)	37.4	5.4	22.3	0.8	0.8	1.5	0.2	0.7	17.3	0.8	3.3	1.3	0.2	6.5	0.3	0.4	0.1	0.1	0.7	-
T (mIoU)	92	67	82	16	27	53	53	63	87	42	84	71	43	86	19	34	20	31	69	54.7
T + U (mIoU)	95	72	87	24	28	56	59	72	90	49	84	77	52	90	24	42	8	39	73	58.9

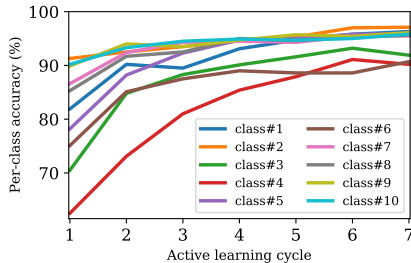


Figure 11. Per-class accuracy on Cifar-10 using the proposed active learning method.

compares different loss functions on Cifar-10 and Cifar-100. CyclicOD loss and EMAOD loss are two TOD-based unsupervised learning criteria. They are minimized on the unlabeled data, and, the settings of their baseline models are identical to those in the study of sampling strategy as discussed above. LL4AL loss [54] minimizes the distance between the predicted loss and the real task loss, and it needs the data labels. All the auxiliary losses are used in a combination with the task loss. The full pipeline and the training with only the task loss are also included in comparison. We observe that either the EMAOD loss or the CyclicOD loss can help to improve the performance, and either of them shows a larger performance improvement than the LL4AL loss. The EMAOD loss demonstrates a more stable performance than the CyclicOD loss, indicating that directly applying COD to unsupervised training may lead to an unstable model training. A moving average of previous model states enables a more stable unsupervised training. We employ EMAOD as our unsupervised loss in the rest of the experiments.

Table 2 shows the per-class performance of standard task model training on Cityscapes, where 40% labeled data is observable. The row of ‘Proportion’ in Table 2 shows the pixel-level proportion of every class, indicating a severe class imbalance problem of Cityscapes. We compare the models trained without (*i.e.*, ‘T’) and with (*i.e.*, ‘T + U’) the unsupervised loss. The semi-supervised learning yields better results on 18 out of 19 classes. More importantly, the semi-supervised learning shows more significant performance improvements on the minority classes than the majority classes, demonstrating that the unsupervised loss imparts robustness to the task model to handle the class imbalance issue.

**Per-class performance.** Fig. 11 shows the per-class accu-

Table 3. Time (seconds) taken for one iteration of active sampling using an NVIDIA GTX 1080Ti GPU. *Number of sampling images and size of images* are shown for each dataset.

Method	Cifar-10 2.5K, 32 <sup>2</sup>	SVHN 3.6K, 32 <sup>2</sup>	Caltech-101 0.4K, 224 <sup>2</sup>	Extra model?
Coreset [40]	91.4	168.7	48.2	×
VAAL [43]	13.0	17.2	32.6	✓
LL4AL [54]	7.7	10.8	39.6	✓
COD (ours)	<b>7.2</b>	<b>10.1</b>	<b>26.9</b>	×

racy on Cifar-10 with the proposed active learning method. The accuracies of the classes are improved along with the increasing of active learning cycles in most cases, such that the performance improvement is not biased towards certain classes. The accuracies of class#3 and class#4 decrease from the 6-th cycle to the 7-th cycle, mainly due to the overfitting.

## 5.5. Time Efficiency

This paper proposes to use COD for active data sampling. Table 3 evaluates the time taken for one active sampling iteration using different active learning methods. On all the three image classification datasets with different number and size of sampling images, COD is faster than the existing active learning methods. COD is task-agnostic and more efficient, since it only relies on the task model itself and it does not introduce extra learnable models such as the adversarial network (VAAL) [43] or the loss prediction module (LL4AL) [54].

## 6. Conclusion

In this paper we have presented a simple yet effective deep active learning approach. The core of our approach is a measurement Temporal Output Discrepancy (TOD) which estimates the loss of unlabeled samples by evaluating the discrepancy of outputs given by models at different gradient descend steps. We have theoretically shown that TOD lower-bounds the accumulated sample loss. On basis of TOD, we have developed an unlabeled data sampling strategy and a semi-supervised training scheme for active learning. Due to the simplicity of TOD, our active learning approach is efficient, flexible, and easy to implement. Extensive experiments have demonstrated the effectiveness of our approach on image classification and semantic segmentation tasks. In future work, we plan to apply TOD to other machine learning tasks and scenarios, as it is an effective loss measure.



## References

- [1] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. There are many consistent explanations of unlabeled data: Why you should average. In *ICLR*, 2019. 3
- [2] Klaus Brinker. Incorporating diversity in active learning with support vector machines. In *ICML*, pages 59–66, 2003. 1, 2
- [3] Razvan Caramalau, Binod Bhattarai, and Tae-Kyun Kim. Sequential graph convolutional network for active learning. In *CVPR*, pages 9583–9592, 2021. 6
- [4] David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. *JAIR*, 4:129–145, 1996. 1, 2
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 7
- [6] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, pages 1422–1430, 2015. 1
- [7] Melanie Ducoffe and Frederic Precioso. Adversarial active learning for deep networks: a margin based approach. *arXiv preprint arXiv:1802.09841*, 2018. 2
- [8] Suyog Dutt Jain and Kristen Grauman. Active image segmentation propagation. In *CVPR*, pages 2864–2873, 2016. 1
- [9] Sayna Ebrahimi, Mohamed Elhoseiny, Trevor Darrell, and Marcus Rohrbach. Uncertainty-guided continual learning with bayesian neural networks. In *ICLR*, 2020. 2, 3
- [10] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE TPAMI*, 28(4):594–611, 2006. 6
- [11] Alexander Freytag, Erik Rodner, and Joachim Denzler. Selecting influential examples: Active learning with expected model output changes. In *ECCV*, pages 562–577, 2014. 2
- [12] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, pages 1050–1059, 2016. 2, 6, 7
- [13] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. *arXiv preprint arXiv:1703.02910*, 2017. 2, 3
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014. 2
- [15] Yuhong Guo. Active instance sampling via matrix partition. In *NIPS*, pages 802–810, 2010. 1, 2
- [16] Hangfeng He and Weijie Su. The local elasticity of neural networks. In *ICLR*, 2020. 1
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [18] Thorsten Joachims. Transductive learning via spectral graph partitioning. In *ICML*, pages 290–297, 2003. 2
- [19] Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *CVPR*, pages 2372–2379, 2009. 1, 2
- [20] Christoph Käding, Erik Rodner, Alexander Freytag, and Joachim Denzler. Active and continuous exploration with deep neural networks and expected model output changes. *arXiv preprint arXiv:1612.06129*, 2016. 2
- [21] Ashish Kapoor, Kristen Grauman, Raquel Urtasun, and Trevor Darrell. Active learning with gaussian processes for object categorization. In *ICCV*, pages 1–8, 2007. 2
- [22] Kwanyoung Kim, Dongwon Park, Kwang In Kim, and Se Young Chun. Task-aware variational adversarial active learning. In *CVPR*, pages 8166–8175, 2021. 6
- [23] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *NIPS*, pages 3581–3589, 2014. 2
- [24] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *ICML*, 2017. 7
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical Report*, 2009. 6
- [26] Weicheng Kuo, Christian Häne, Esther Yuh, Pratik Mukherjee, and Jitendra Malik. Cost-sensitive active learning for intracranial hemorrhage detection. In *MICCAI*, pages 715–723, 2018. 2, 7
- [27] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2, 3
- [28] Yann LeCun, D Touresky, G Hinton, and T Sejnowski. A theoretical framework for back-propagation. In *Proceedings of the 1988 connectionist models summer school*, volume 1, pages 21–28, 1988. 1
- [29] Xin Li and Yuhong Guo. Adaptive active learning for image classification. In *CVPR*, pages 859–866, 2013. 2
- [30] Buyu Liu and Vittorio Ferrari. Active learning for human pose estimation. In *ICCV*, pages 4363–4372, 2017. 2
- [31] Wenjie Luo, Alex Schwing, and Raquel Urtasun. Latent structured active learning. In *NIPS*, pages 728–736, 2013. 2, 3
- [32] Dwarikanath Mahapatra, Behzad Bozorgtabar, Jean-Philippe Thiran, and Mauricio Reyes. Efficient active learning for image classification and segmentation using a sample selection and conditional generative adversarial network. In *MICCAI*, pages 580–588, 2018. 2
- [33] Christoph Mayer and Radu Timofte. Adversarial sampling for active learning. In *WACV*, pages 3071–3079, 2020. 2
- [34] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE TPAMI*, 41(8):1979–1993, 2018. 2
- [35] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 6
- [36] Hieu T Nguyen and Arnold Smeulders. Active learning using pre-clustering. In *ICML*, page 79, 2004. 2
- [37] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 5

- [38] Dan Roth and Kevin Small. Margin-based active learning for structured output spaces. In *ECML*, pages 413–424, 2006. [2](#), [3](#)
- [39] Nicholas Roy and Andrew McCallum. Toward optimal active learning through monte carlo estimation of error reduction. pages 441–448, 2001. [2](#)
- [40] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *ICLR*, 2018. [1](#), [2](#), [6](#), [7](#), [8](#)
- [41] Burr Settles and Mark Craven. An analysis of active learning strategies for sequence labeling tasks. In *EMNLP*, pages 1070–1079, 2008. [1](#), [2](#)
- [42] Burr Settles, Mark Craven, and Soumya Ray. Multiple-instance active learning. In *NIPS*, pages 1289–1296, 2008. [2](#)
- [43] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *ICCV*, pages 5972–5981, 2019. [2](#), [6](#), [7](#), [8](#)
- [44] Yu-Yin Sun, Yin Zhang, and Zhi-Hua Zhou. Multi-label learning with weak label. In *AAAI*, pages 593–598, 2010. [1](#)
- [45] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014. [3](#)
- [46] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, pages 1195–1204, 2017. [2](#), [3](#), [5](#)
- [47] Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. *JMLR*, 2(Nov):45–66, 2001. [2](#)
- [48] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440, 2020. [2](#)
- [49] Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In *NIPS*, pages 3835–3844, 2018. [3](#)
- [50] Keze Wang, Xiaopeng Yan, Dongyu Zhang, Lei Zhang, and Liang Lin. Towards human-machine cooperation: Self-supervised sample mining for object detection. In *CVPR*, pages 1605–1613, 2018. [1](#)
- [51] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-effective active learning for deep image classification. *IEEE TCSVT*, 27(12):2591–2600, 2016. [2](#)
- [52] Zheng Wang and Jieping Ye. Querying discriminative and representative samples for batch mode active learning. *TKDD*, 9(3):1–23, 2015. [1](#), [2](#)
- [53] Lin Yang, Yizhe Zhang, Jianxu Chen, Siyuan Zhang, and Danny Z Chen. Suggestive annotation: A deep active learning framework for biomedical image segmentation. In *MICCAI*, pages 399–407, 2017. [2](#)
- [54] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *CVPR*, pages 93–102, 2019. [2](#), [5](#), [6](#), [7](#), [8](#)
- [55] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *CVPR*, pages 472–480, 2017. [7](#)
- [56] Zheng-Jun Zha, Tao Mei, Jingdong Wang, Zengfu Wang, and Xian-Sheng Hua. Graph-based semi-supervised learning with multiple labels. *Journal of Visual Communication and Image Representation*, 20(2):97–103, 2009. [2](#)
- [57] Beichen Zhang, Liang Li, Shijie Yang, Shuhui Wang, Zheng-Jun Zha, and Qingming Huang. State-relabeling adversarial active learning. In *CVPR*, pages 8756–8765, 2020. [2](#), [6](#), [7](#)
- [58] Zongwei Zhou, Jae Shin, Lei Zhang, Suryakanth Gurudu, Michael Gotway, and Jianming Liang. Fine-tuning convolutional neural networks for biomedical image analysis: actively and incrementally. In *CVPR*, pages 7340–7351, 2017. [2](#)
- [59] Zhi-Hua Zhou, De-Chuan Zhan, and Qiang Yang. Semi-supervised learning with very few labeled training examples. In *AAAI*, pages 675–680, 2007. [1](#)
- [60] Jia-Jie Zhu and José Bento. Generative adversarial active learning. *arXiv preprint arXiv:1702.07956*, 2017. [2](#)