PT4AL: Using Self-Supervised Pretext Tasks for Active Learning

John Seon Keun Yi^{1*}, Minseok Seo
2,4*, Jongchan Park³, and Dong-Geol Choi
4**

Georgia Institute of Technology¹, SI Analytics², Lunit Inc³, Hanbat National University⁴

Abstract. Labeling a large set of data is expensive. Active learning aims to tackle this problem by asking to annotate only the most informative data from the unlabeled set. We propose a novel active learning approach that utilizes self-supervised pretext tasks and a unique data sampler to select data that are both difficult and representative. We discover that the loss of a simple self-supervised pretext task, such as rotation prediction, is closely correlated to the downstream task loss. Before the active learning iterations, the pretext task learner is trained on the unlabeled set, and the unlabeled data are sorted and split into batches by their pretext task losses. In each active learning iteration, the main task model is used to sample the most uncertain data in a batch to be annotated. We evaluate our method on various image classification and segmentation benchmarks and achieve compelling performances on CIFAR10, Caltech-101, ImageNet, and Cityscapes. We further show that our method performs well on imbalanced datasets, and can be an effective solution to the cold-start problem where active learning performance is affected by the randomly sampled initial labeled set. Code is available at https://github.com/johnsk95/PT4AL

Keywords: Active Learning, Self-supervised Learning, Pretext Task

1 Introduction

The recent success in deep learning has shown remarkable advancements in computer vision tasks such as classification [19, 12] and semantic segmentation [7, 30]. This has been possible due to the advent of deep convolutional neural networks (CNNs) and large annotated datasets such as ImageNet [12] and COCO [27].

As deep learning models are trained in a data-driven manner, having a large enough training set is crucial to achieve high performance. However, building a large labeled dataset is prohibitively time-consuming and expensive. Labeling costs increase with the size of data and complexity of the tasks. Instead of labeling the entire data, active learning (AL) [39] aims to select informative subsets to label that achieve the highest performance within a fixed labeling budget.

^{*} Equal contribution

^{**} Corresponding author (dgchoi@hanbat.ac.kr)



Fig. 1: The overall framework of the proposed method. Unlabeled data are sorted by pretext task losses, split into batches, and sampled for training

Existing AL approaches can be divided into two main groups: distributionbased and uncertainty-based methods. Distribution-based methods [38, 5] aim to sample data that well covers the distribution of the feature space. The advantage of such methods is that they can sample *representative* points: data points from high density regions that well represent the overall feature distribution. However, distribution-based sampling fails to select data that are placed near the decision boundary (i.e. high uncertainty data points). Uncertainty-based approaches [26] resolve this problem by sampling the most uncertain points. Simple approaches that utilize class posterior probabilities [26, 25], entropy [40, 21], and loss prediction [46, 23] were revealed to perform well on various settings. While these approaches effectively sample *uncertain* or *difficult* data near decision boundaries in the feature space, they do not capture the overall distribution of the data, according to our qualitative analysis in Fig. 4. Our method aims to capture the best of both worlds by sampling both representative and difficult data.

This paper proposes Pretext Tasks for Active Learning (PT4AL), a novel active learning framework that utilizes self-supervised pretext tasks combined with an uncertainty-based sampler. We train a pretext task model [16, 48] with unlabeled data, and the pretext task loss is highly correlated to the main task loss. In order to sample diversely from both representative and difficult data, the unlabeled data are sorted in descending order by their pretext task loss, and split into batches to be used for each AL iteration. Starting from the batch containing data with the highest losses, the most uncertain K data points are sampled from each batch, based on the posterior class probability of the previous main task learner. The uncertainty-based sampler enables PT4AL to sample difficult data, while the batch split allows balanced sampling across the entire data distribution.

PT4AL also resolves the innate problem in active learning: the cold start problem. Existing approaches start from a randomly sampled set of labeled data, rendering the overall performance highly dependent on the distribution of the initial set. Since our method learns the representation of the unlabeled set in advance, we can sample informative data from the first iteration. This approach avoids the issue of high variance and decrease in performance that can stem from randomly sampling the initial labeled set.

We validate our proposed method on various image classification and semantic segmentation datasets and achieve state-of-the-art or compelling results across different datasets and tasks. Additionally, we demonstrate the robustness of PT4AL on a class imbalanced setting by evaluating on an artificially created class-imbalanced CIFAR10 dataset.

2 Related Work

Active Learning Various AL approaches has been proposed, such as information theoretical approaches [32], ensemble approaches [33, 15], uncertainty based methods [43, 21] and Bayesian AL methods [22]. However, these traditional methods have not been verified in large-scale datasets for large-scale models, such as in the field of CNN-based deep learning, which has achieved state-ofthe-art in various computer vision tasks.

Recent AL methods have been centered on large-scale settings for CNNbased deep learning models. Sener & Savarese [38] proposed a core-set selection method, which chooses data points that cover all data with high diversity based on the feature distribution. This method targets two problems of the previous uncertainty-based methods. First, uncertainty-based methods select only hard samples, resulting in redundant, overlapping data points. Second, the existing methods are not suitable for batch processing on CNNs. The core-set algorithm aims to sample diverse data points in a batch manner. Yoo & Kweon [46] proposed a sub-task module to predict the main task loss of unlabeled data, and sample the high-loss samples from the unlabeled pool. This method samples from a subset of the unlabeled pool to avoid selecting redundant data points when sampling consecutively from the most uncertain data [4]. However, in our qualitative analysis in Fig. 4, uncertainty-based methods like Yoo & Kweon sample data points from decision boundaries with less diversity in distribution. Recently, using a variational autoencoder architecture [41], the discriminator adversarially trains the input data to be unlabeled or labeled. In the data sampling phase, a method that first labels the sample predicted as unlabeled with the lowest confidence was proposed.

Our active learning method uses a self-supervised pretext task to supplement the flaws of the data distribution-based method and the uncertainty-based method. As described above, AL is largely divided into data distribution-based methods [38, 5, 29] and uncertainty-based methods [41, 23, 46, 10]. The data distribution-based method has the disadvantage that it cannot extract hard samples, and the uncertainty-based method has the possibility to sample overlapping data points and it is difficult to extract the representation of the entire data distribution. Other works [20, 45, 1] sample from both representative and difficult data by utilizing variance maximization between labeled and unlabeled data or using separate sampling criteria for data in each category. Our method uses pretext task-based batch split which allows us to select representative samples across the semantic distribution, and an uncertainty-based in-batch sampler which allows us to select difficult samples.

Representation Learning with Pretext Tasks Representation learning aims to learn good pre-trained weights by learning self-supervised pretext tasks with unlabeled data. The pre-trained weights are fine-tuned with a small amount of labeled data to achieve high performance on downstream tasks. The key assumption and the findings in representation learning is that pretext tasks provide enough learning signals without any labels (i.e. direct supervision) provided. Using these assumptions, Liu et al. [28] proposed unsupervised neural architecture search (NAS) using self-supervised pretext tasks [48, 34, 16] and achieved similar performance to supervised NAS baselines. Zhang et al. [48] proposed a pretext task to restore the color of the original image through a network after transforming the input image to gray scale. Noroozi & Favaro [34] improved the performance of representation learning in image classification through the task of dividing input images into grids, mixing them with each other, and inputting each grid into the network. Gidaris et al. [16] proposed a pretext task that rotates the input image by 0° , 90° , 180° , and 270° and training the network to match the rotated angle of the transformed input image. This method achieved the highest performance among representation learning methods utilizing data structures. Recently proposed representation learning methods use contrastive learning [35, 8, 9, 6, 18] to minimize the distance between different pairwise augmentations of the same image, and repel from augmentations of different images. Contrastive learning is proved to be robust on different downstream tasks and provide state-of-the-art results by far.

There have been several efforts to use self-supervised pretext tasks in active learning. Zhu *et al.* [49] uses graph contrastive learning [47] for active learning on graph neural networks. [2, 37, 17] utilizes self-supervised learning to pre-train the main task model, which is then fine-tuned on labeled data. Bhatnagar *et al.* [3] presents a multi-task active learner trained for both pretext task and main task, while being robust to mislabeled samples. Although these methods help justify the use of pretext tasks in active learning, they are limited to specific domains [49, 37, 17], fail to sample both difficult and representative data [2], and does not solve the cold start problem [49, 2, 3].

As pretext tasks provide good initializations for downstream tasks, we assume that the information learned through these tasks is highly correlated to the semantic data distribution. We analyze and identify the correlation between the pretext task loss and the supervised loss in downstream tasks in Section 3. Finally, we propose an active learning method using pretext tasks in Section 4.

3 Using Pretext Tasks for Active Learning

The success of representation learning with self-supervised pretext tasks [8, 18, 9, 28], leads us to believe that there is a high correlation between self-supervised pretext tasks and downstream tasks, and thus pretext tasks can be utilized for active learning. Rather than utilizing the feature distribution after the pretext task training, we resort to a simpler metric for active learning - the pretext task loss. In this section, we propose and validate a hypothesis, and use these evidences to formulate our AL algorithm. Our hypothesis is that:

H1: Pretext task loss is correlated with the main task loss.



Fig. 2: (From left to right) The loss rank correlation plots for the main task loss and the pretext task loss in CIFAR10, Caltech-101 and ImageNet. The x and y axes represent the normalized rank of the two losses, respectively

We think that if a pretext task is correlated or representative of the main task, images that are *hard* (i.e. having high loss values) for the pretext task will also be *hard* for the main task.

Fig. 2 presents scatter plots of the pretext task loss and the main task loss in three benchmark datasets. The x-axis is the normalized rank of the main task loss, and the y-axis is the normalized rank of the pretext task loss. The pretext task and the main task are independently trained with the training set, and the losses are computed in the test set. For ease of interpretation, we visualized 1,000 random samples on the plots. Spearman's rank correlation [42] denoted as ρ is calculated on the full test set.

As illustrated in Fig. 2, the pretext and main task losses have a strong positive correlation. That is, if a data sample has high loss for a pretext task, it is likely for it to have high loss for the main task, and vice versa. We observe high ρ values for all three datasets: CIFAR10 ($\rho = 0.79$), Caltech-101 ($\rho = 0.78$), and ImageNet $(\rho = 0.88)$. Note that these datasets vary in image size, number of classes, and class balance. The strong correlation between the pretext task loss and the main task loss across diverse datasets validates our hypothesis, and thus is a strong evidence for using pretext task losses for active learning. However, there is one caveat to the hypotheses: methods that use contrastive loss as the pretext task [18, 8, 9] do not have a strong loss correlation. ($\rho = -0.001$ for SimSiam) We contribute this result to two main reasons: class bias of the contrastive loss and strong reliance to augmentations. Details are explained in the supplementary material. Even if we could find a way to achieve close correspondence with the main loss, we decide not to use contrastive methods since the large batch size and long training time generally required for these methods beat our purpose of a simple and quick AL model. Details are explained in the supplementary material.

Throughout this work, we validate the efficacy of PT4AL with 4 different pretext tasks: Rotation prediction [16], colorization [48], solving jigsaw puzzles [34], and SimSiam [9]. We compare and analyze the efficacy of different pretext tasks on classification and semantic segmentation in Section 6.2. Since rotation prediction [16] performs best in CIFAR10 and colorization [48] performs the best in Cityscapes, we use rotation prediction for image classification main tasks, and colorization for semantic segmentation.

4 Method

In this section, we introduce the specifics of PT4AL. First, we provide a brief overview of our active learning algorithm. Then we provide details of the pretext task learning for batch split and in-batch sampling in the following sections.

4.1 Overview

In a typical active learning scenario, we are initially provided with a pool of unlabeled data $x_U \in X_U$. The objective of AL is to achieve the best performance in the main task model $F_m(\cdot)$ with a limited amount of labeled data. In specific, we follow the *batch mode* active learning scheme: in the *i*-th AL iteration, we select K samples from X_U^i , add them into labeled pool (X_L^i, Y_L^i) with oracle, train and evaluate $F_m^i(\cdot)$ with (X_L^i, Y_L^i) . The iterations are repeated until the specified labeling budget is reached.

The overall framework of PT4AL is illustrated in Fig. 1. PT4AL is split into two parts: pretext task learning for batch split and in-batch sampling. Pretext task learning is done prior to the AL iterations. We train a pretext task learner with X_U . The unlabeled samples are sorted in descending order of their pretext task losses, and split into batches. The in-batch sampling is done at each AL iteration. At the *i*-th iteration, the sampling module selects K samples from the *i*-th batch, according to the uncertainty of the main task learner in these samples. The main task learner $F_m^i(\cdot)$ is trained with (X_L^i, Y_L^i) and evaluated on the test set.

4.2 Pretext Task Learning for Batch Split

In this section, we explain how a pretext task is used for active learning batch split. The term *batch* refers to a pool of unlabeled data to be sampled in an AL iteration. While any pretext task can be used in our method, we use the widely used rotation prediction task [16] for the explanation. For the rotation prediction task, the backbone neural network [19] is trained on all four orientations (0°, 90°, 180°, 270° degrees) of the input image. The loss function is defined as the average of the losses for each orientation:

$$loss(x_i, \theta_p) = \frac{1}{k} \sum_{y=1}^k \mathcal{L}_{CE}(F_p(g(x_i \mid y) \mid \theta_p), y)$$
(1)

Where \mathcal{L}_{CE} is the cross-entropy loss. The rotation operator $g(\cdot \mid y)$ yields the rotated input image according to the orientation label y. We define k = 4 since we predict four different rotations. F_p represents the probability distribution of

the input image rotated by label y. Note that the rotation label y is unknown to F_p . In inference, four orientations of each image is fed into the trained network F_p and the extracted loss is the same averaged loss $loss(x_i, \theta_p)$ used in training. F_p is trained and tested on the same unlabeled set X_U . The model weights θ_p with the best test accuracy is used for loss extraction.

After training the pretext task learner, we extract pretext task loss values from X_U and split them into batches. Given the pretext task loss values of the unlabeled data $loss_{X_U}$ in the pretext task learning phase, we first sort the losses in descending order. The sorted data \mathcal{X}_U is then divided into I batches of equal size. The number of I is equal to the number of AL iterations: if there are ten iterations(I = 10), there will be ten batches $\mathcal{B} = \{b_i\}_{i=1}^{I=10}$.

4.3 In-batch Sampling

The in-batch sampler selects K samples at each AL iteration. At the *i*-th iteration, the in-batch sampler $\phi(\cdot)$ selects K samples from the *i*-th batch to be annotated by the oracle. The sampler computes the top-1 posterior probability in the given batch using the previous main task learner F_m^{i-1} , and K data points with the lowest confidence scores are selected. In the first iteration, K points are sampled from the first batch b_0 at even intervals. Equation 2 summarizes the sampler $\phi(\cdot)$. The sampling makes use of the main task model from the previous iteration, F_m^{i-1} .

$$\phi(b_i, F_m^{i-1}) = \min_K \{ \max(F_m^{i-1}(b_i \mid \theta_m)) \}$$
(2)

Algorithm 1 illustrates our overall sampling algorithm including batch splitting and in-batch sampling. In the first iteration when we do not have F_m^0 , we uniformly select samples in the first batch, based on our empirical observation that visually similar samples have similar pretext task loss values. Sampled data have

Algorithm 1 Sampling Strategy	
Input: Unlabeled X_U , labeled X_L , pretext	t task losses $loss_{X_U}$, main task model F_m
$\mathcal{X}_U = sort(loss_{X_U})$	\triangleright Sort losses in descending order
Split \mathcal{X}_U into batches \mathcal{B}	
for b_i in \mathcal{B} do	
if $i == 1, X_K = uniform(b_i, loss^i_{X_U})$	\triangleright For the first batch, uniformly sample
else, $X_K = \phi(b_i, F_m^{i-1})$ \triangleright For oth	ner batches, sample top-K uncertain data
$X_U \leftarrow X_U - X_k$	\triangleright Remove from unlabeled pool
$X_L \leftarrow X_L \cup X_k$	\triangleright Add to labeled pool
train F_m^i with X_L	
end for	

two main traits: difficult and representative. Difficult or uncertain data refers to data that the main task model cannot easily distinguish because it is near a decision boundary. Conversely, representative data well defines the distribution 8 JSK. Yi, M. Seo et al.

in the feature space. Our intuition is that if we can sample data from both categories, we can form a labeled pool with the most informative data. This is empirically verified through query analysis in Section 5. Our batch split method combined with the sampler samples both representative and difficult data. Our method is much simpler and well performing compared to previous works that sample data from both traits [20, 45, 1].

5 Experiments

We evaluate the efficacy of our method on two commonly used visual recognition tasks: image classification and semantic segmentation. We choose CI-FAR10 [24], Caltech-101 [14], ImageNet [12] benchmarks for image classification, and Cityscapes [11] for semantic segmentation. To further demonstrate our method's efficacy in a more challenging class-imbalanced setting, we additionally use a class-imbalanced version of CIFAR10. Finally, we show the use of PT4AL as an effective solution to the cold start problem. Unless otherwise specified, all the experiment results are reproduced by ourselves, averaged over multiple runs with different random seeds.

5.1 Image Classification

Dataset We perform experiments on three image classification datasets with varying size and number of classes. CIFAR10 contains 50,000 training and 10,000 testing images of size 32×32 with 10 object categories. We start with 1,000 labeled images, and 1,000 images are added for each iteration. Caltech-101 has 9,144 images of size around 300×200 distributed around 101 classes. We divide the data into 8,046 for training and 1,098 for testing. Similar to CIFAR10 we also start with 1,000 labeled images of 1,000 classes. 1,279,867 and 49,950 images are used for the training and testing set. For ease of experimentation and to avoid noise from similar class labels, ImageNet classes are reduced to 67 based on the WordNet [36] superclasses. ImageNet starts with $K \approx 128,000$ labeled samples, and the same K samples are selected for each iteration. Due to heavy computation, each ImageNet performance is the average of 3 runs.

Baselines and implementation details We compare PT4AL with random sampling, Core-Set [38], Variational Adversarial Active Learning (VAAL) [41], Learning Loss [46], CoreGCN [5], and PAL [3]. For CIFAR10 we add "Learning loss(detached)", where the loss prediction task is detached during supervised learning to avoid influences from multi-task learning. ResNet-18 [19] is used as the backbone network for the pretext task and the main task learner. The final linear layer of the pretext task learner is converted to (512,4) to account for the four orientations of the rotation task. For Caltech-101 and ImageNet, input images are resized into 224×224 . No data augmentation is applied in the pretext task learning phase. Random resized crop and horizontal flip is applied in



Fig. 3: Comparison of image classification performance on CIFAR10, Caltech-101, ImageNet-67. Best viewed in color

the main task phase. The main task is trained for 200 epochs in CIFAR10 and Caltech-101, and 100 epochs in ImageNet. SGD with a multi-stage learning rate is applied. Detailed hyper-parameters are described in the supplement material.

Results Fig. 3a demonstrates the results for CIFAR10. PT4AL clearly outperforms other methods across all AL iterations by a noticeable margin. The accuracy of PT4AL in the final iteration of 10,000 labeled points is 95.13% (\uparrow 8.91%), while the second-best performing learning loss scores 89.93% (\uparrow 3.71%). Note that detached learning loss [46] performs significantly worse than the original multi-task learning approach, where the main task model is simultaneously trained with auxiliary tasks. The significant drop in performance due to the detachment indicates that the multi-task approaches [46, 41, 5] may benefit from multi-task learning. To strictly measure the benefit of AL to select informative samples, we need to compare the detached setting across all methods. Our method also has a significant advantage from the first iteration, achieving an accuracy of 55.83% (\uparrow 9.81%) compared to the other methods' 46.02%. This emphasizes the advantage of PT4AL sampling informative points in the first iteration, instead of random sampling in other AL frameworks. Further details of PT4AL solving the cold-start problem is described in Section 5.4.

Similar results can also be observed in Caltech-101 and ImageNet, in Fig. 3b and Fig. 3c. Our method outperforms other methods across most of the iterations with a considerable advantage from the start.

Query Analysis Fig. 4 illustrates t-SNE [31] embeddings of the CIFAR10 data points sampled by random, learning loss [46] and ours. For a fair comparison, we use embeddings extracted from a ResNet-18 model trained with fully labeled CIFAR10. To visualize the sampled data in different methods across the AL iterations, each of the 1,000 samples from the first iteration are marked in circle, fifth iteration in triangle, and tenth (last) iteration as square. Fig. 4a shows that random sampling queries evenly from the embedding space, but fails to sample difficult data points along the decision boundaries. As shown in Fig. 4b, learning loss [46] has most of its queries concentrated on the border regions. While this



Fig. 4: t-SNE visualization of the CIFAR10 dataset for random, learning loss [46] and PT4AL. Vivid points are sampled for labeling. Best viewed in color



Fig. 5: Comparison of cityscapes semantic segmentation and imbalanced cifar10 classification. Best viewed in color

may be effective for the labeled classifier to learn difficult points, it does not query points that represent the classes well. Fig. 4c shows that PT4AL queries from both difficult and representative regions. The sampled points are either concentrated on the class boundaries or evenly located in the class distributions. Since PT4AL initially samples from batches with higher pretext loss values, selected points from the first iteration are concentrated on the decision boundaries of the embedding space. As the sampler progresses to batches with lower loss values, we can see that the sampled points propagate to the remaining regions of the class clusters. Such sampling behavior is a mix of both distribution and uncertainty-based methods, mitigating their flaws while sampling both difficult and representative data points.

5.2 Semantic Segmentation

Dataset We choose Cityscapes [11], a public benchmark dataset widely used in semantic segmentation. The dataset consists of 2,975 training and 500 validation

images. At each AL iteration, 100 images are sampled for the labeled training set. The original training set is set as the unlabeled set.

Baselines and Implementation Details We choose the state-of-the-art active learning methods for this experiment: Core-Set [38], Learning loss [46], VAAL [41], TA-VAAL [23], and PAL [3]. We choose a widely-used semantic segmentation architecture, DeepLab [7] with a ResNet-101 [19] backbone. The model is initialized with ImageNet [13] pre-trained weights. The input images are resized to (1024,512) and no data augmentation is applied. The training batch size is 1, and all hyper-parameters follow the original paper [7], unless otherwise specified.

Results Fig. 5a demonstrates that PT4AL outperforms all other methods by a noticeable margin across all iterations. The performance improvement at the first iteration is also significant, showing that PT4AL is an effective solution for the cold start problem. Note that learning loss [46], VAAL [41] and Core-Set [38] are not as effective as in CIFAR10, sometimes being worse or on-par with the random selection baseline. One possible reason for the universal effectiveness of PT4AL across tasks is that the nature of PT4AL can dynamically change with respect to the pretext task being used. A more detailed analysis among pretext tasks is described in Section 6.2.

5.3 Image Classification on an Imbalanced Dataset

Dataset & Experiment details To evaluate the efficacy of PT4AL on a more challenging class-imbalanced setting, we recompose the CIFAR10 dataset. The number of images for each class is as follows: airplane-500, automobile-1,000, bird-1,500, cat-2,000, deer-2,500, dog-3,000, frog-3,500, horse-4,000, ship-4,500 and truck-5,000. All implementation details are identical to the balanced CI-FAR10 described in 5.1, except for dataset composition.

Results Fig. 5b demonstrates the performances of PT4AL and other baselines on imbalanced CIFAR10. PT4AL outperforms other baselines across all iterations by a large margin. In the pretext task, data from classes with little training data generally have high loss, and classes that have abundant training data generally have low loss values. Since PT4AL samples from data batches with high to low loss, it can sample in a class-balanced way even in imbalanced settings. Also, unlike other methods using only the main task model related metrics, PT4AL utilizes the pretext task loss which is completely independent from the main task model. Interestingly, unlike the experiment results in balanced CIFAR10, data distribution-based AL methods (Core-Set, CoreGCN) obtains higher performance than the uncertainty-based methods (VAAL, learning loss). These results empirically show that uncertainty-based methods are more negatively affected by the class-imbalanced setting than the distribution-based methods. PT4AL outperforms other methods by a margin, showing robustness on a more challenging class-imbalanced setting. Furthermore, we observe that PT4AL samples data in a more class-balanced way. Details on the class distribution of the sampled data are in the supplementary material.

12 JSK. Yi, M. Seo et al.

5.4 Cold Start Problem in Active Learning

Since most AL approaches require a trained main task model, the first AL iteration starts with randomly selected labeled data. This is what we call *the cold start problem in active learning*. To thoroughly validate the efficacy of our method as a solution to the cold start problem, we take a closer look into the first AL iteration in the CIFAR10 benchmark. Note that all other methods use random selection for the first iteration. For PT4AL, after training the pretext task learner, the unlabeled data are sorted by pretext task loss in a descending order, split into 10 batches, and 1,000 data points are uniformly selected from the first batch. The experiment is repeated 20 times with different random seeds. All implementation details are identical to Section 5.1.

Table 1 summarizes the experiment results. PT4AL displays more stable performance compared with random sampling in the first iteration, as the standard deviation is smaller, and the gap between the max/min accuracy is smaller than that of the random baseline. PT4AL significantly outperforms random in the average accuracy, indicating that more informative data points are sampled for the main task model. These results indicate that PT4AL is a good solution to the cold start problem, and can be used as a good starting point for existing AL methods [46, 41, 23, 29, 5]. More details are in the supplement material.

Table 1: Results of the first active learning iteration in CIFAR10

Method	Mean accuracy	Min / Max
Random	$47.49 \pm 3.15\%$	43.06 / 53.74 %
PT4AL(Rotation)	$55.20 \pm 1.95~\%$	$52.00 \ / \ 57.71 \ \%$

5.5 Computational Overheads

As described in Section 4, the extra computation for PT4AL, apart from the main task model training, is the pretext task learning for batch split, and the unlabeled data inference for uncertainty measurement in in-batch sampling. To fairly compare the computational overheads of different approaches, we measure the wall-clock time of the methods compared in CIFAR10 experiment under the same environment. In Fig. 3a and Table 2, we can observe that PT4AL achieves the best performance while having on-par computational overheads with others. Core-Set [38] has similar computations as the random selection baseline.

Table 2: The wall-clock time of each algorithm in the CIFAR10 experiment

Method	Random Selection	PT4AL	Learning Loss [46]	VAAL $[41]$	CoreGCN [5]
Time	2hr 16min	3hr 36min	2hr 37min	14hr 9min	3hr 48min



Fig. 6: (a) Ablation on two components of PT4AL. (b),(c) PT4AL with different pretext tasks on CIFAR10 and Cityscapes

6 Ablation Study

6.1 Ablation on sampling strategy and pretext task loss

Fig. 6.a shows the ablations results of the two core components of PT4AL. Instead of the pretext task loss, "Sampling only" uses the main task model's entropy. Batches are made by randomly segmenting the unlabeled data. Data in the first iteration are randomly sampled as there is no main task model to begin with. "Pretext only" replaces our proposed sampling method with a naive sampling of high-loss samples or low-loss samples. Compared with PT4AL which uses both heuristics, the two variations display inferior performance. Using both components is imperative to a well-performing model.

6.2 Pretext Tasks

Fig. 6.b and Fig. 6.c presents active learning performance of PT4AL using different pretext tasks. Rotation prediction [16], colorization [48], solving jigsaw puzzles [34], and SimSiam [9] are compared in CIFAR10 and Cityscapes benchmarks. The experiment settings are identical to Section 5.1 and Section 5.2. The inferior performance of SimSiam is analyzed in detail in the supplement material. The rotation prediction task shows the best performance in CIFAR10, and the colorization task performs best in Cityscapes. The best performing pretext task differs by the main task. Since rotation prediction is an image-level task and colorization is a pixel-level task, it is intuitive to match rotation prediction with image classification and colorization with segmentation.

6.3 Sampling Strategy

Sampling in the first iteration In the first iteration, we do not have access to the main task model for uncertainty measurement. Thus, sampling within the first batch resorts to sampling with the pretext task losses. We compare three

14 JSK. Yi, M. Seo et al.

simple sampling methods in CIFAR10: top-K, random, and uniform. The performances for top-K loss, random, and uniform sampling are 44.65%, 51.88%, and 55.20%, respectively. As the uniform sampling outperforms the other two sampling methods, we choose it as our in-batch sampling method for the first iteration. We observe that the samples with similar loss values are visually similar, indicating overlapping semantic information in the top-K sampling. This observation is also coherent with the best performance of uniform sampling, as it avoids selecting data points with visually too similar data points. More details are in the supplementary material.

High loss first vs Low loss first batch split We examine two different strategies for batch split: high loss batch first or low loss batch first. High loss batch first method starts the first iteration with the batch containing the highest pretext task losses, then moves to batches with lower losses for consecutive iterations. Low loss batch first is in reverse. On the CIFAR10 experiment, the high loss first strategy displays slightly better results with 55.20% accuracy in the first iteration and 95.13% from the last iteration. Low loss first strategy scores 53.47% and 94.59% in the first and last iterations. We attribute the small performance difference between the two batch split methods to a finding in curriculum learning [44]. Low loss batch first and high loss batch first are analogous to curriculum learning and anti-curriculum learning, respectively. Wu et al. [44] concludes that curriculum or anti-curriculum is not effective in standard settings, which explains the small performance gap. As the high loss first method performs better across all iterations, we choose it as our batch split method.

7 Conclusion

In this paper we introduce PT4AL, a novel active learning method based on pretext tasks. We demonstrate the correlation between pretext tasks and semantic recognition tasks, and utilize the pretext task losses to split unlabeled samples into batches. In the query analysis in Section 5, we show that the batches are scattered across the whole semantic distribution. Combined with the uncertaintybased in-batch sampler, PT4AL samples both difficult and representative data from the unlabeled pool. We thoroughly examine our method on two widely used vision tasks across various datasets. Our method demonstrates compelling results on datasets with varying resolution, scale and class distribution. We also show that PT4AL is an effective solution for the cold start problem. Although our proposed method performs well on different tasks and datasets, performance varies by the pretext task being used and some tasks such as SimSiam [9] perform poorly. Future research directions may include designing a pretext task that is universal across various recognition tasks.

Acknowledgments: This research was supported by the the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. 2022R1F1A1075019, No. 2021M3E8A2100446) and partially supported by the IITP grant funded by the MSIT (No. 2014-3-00123).

Bibliography

- Behpour, S., Liu, A., Ziebart, B.: Active learning for probabilistic structured prediction of cuts and matchings. In: International Conference on Machine Learning. pp. 563–572. PMLR (2019)
- [2] Bengar, J.Z., van de Weijer, J., Twardowski, B., Raducanu, B.: Reducing label effort: Self-supervised meets active learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1631–1639 (2021)
- [3] Bhatnagar, S., Goyal, S., Tank, D., Sethi, A.: Pal: Pretext-based active learning. arXiv preprint arXiv:2010.15947 (2020)
- Birodkar, V., Mobahi, H., Bengio, S.: Semantic redundancies in image-classification datasets: The 10% you don't need. arXiv preprint arXiv:1901.11409 (2019)
- [5] Caramalau, R., Bhattarai, B., Kim, T.K.: Sequential graph convolutional network for active learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9583–9592 (2021)
- [6] Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. arXiv preprint arXiv:2006.09882 (2020)
- [7] Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE transactions on pattern analysis and machine intelligence 40(4), 834–848 (2017)
- [8] Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
- Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021)
- [10] Cho, J.W., Kim, D.J., Jung, Y., Kweon, I.S.: Mcdal: Maximum classifier discrepancy for active learning. IEEE transactions on neural networks and learning systems (2022)
- [11] Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
- [12] Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- [13] Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255 (2009). https://doi.org/10.1109/CVPR.2009.5206848

- 16 JSK. Yi, M. Seo et al.
- [14] Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In: 2004 conference on computer vision and pattern recognition workshop. pp. 178–178. IEEE (2004)
- [15] Freund, Y., Seung, H.S., Shamir, E., Tishby, N.: Selective sampling using the query by committee algorithm. Machine learning 28(2), 133–168 (1997)
- [16] Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. arXiv preprint arXiv:1803.07728 (2018)
- [17] Hacohen, G., Dekel, A., Weinshall, D.: Active learning on a budget: Opposite strategies suit high and low budgets. arXiv preprint arXiv:2202.02794 (2022)
- [18] He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2020)
- [19] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- [20] Huang, S.J., Jin, R., Zhou, Z.H.: Active learning by querying informative and representative examples. Advances in neural information processing systems 23 (2010)
- [21] Joshi, A.J., Porikli, F., Papanikolopoulos, N.: Multi-class active learning for image classification. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 2372–2379. IEEE (2009)
- [22] Kapoor, A., Grauman, K., Urtasun, R., Darrell, T.: Active learning with gaussian processes for object categorization. In: 2007 IEEE 11th International Conference on Computer Vision. pp. 1–8. IEEE (2007)
- [23] 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. pp. 8166–8175 (2021)
- [24] Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
- [25] Lewis, D.D., Catlett, J.: Heterogeneous uncertainty sampling for supervised learning. In: Machine learning proceedings 1994, pp. 148–156. Elsevier (1994)
- [26] Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: SIGIR'94. pp. 3–12. Springer (1994)
- [27] Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
- [28] Liu, C., Dollár, P., He, K., Girshick, R., Yuille, A., Xie, S.: Are labels necessary for neural architecture search? In: European Conference on Computer Vision. pp. 798–813. Springer (2020)
- [29] Liu, Z., Ding, H., Zhong, H., Li, W., Dai, J., He, C.: Influence selection for active learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9274–9283 (2021)

- [30] Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015)
- [31] Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(11) (2008)
- [32] MacKay, D.J.: Information-based objective functions for active data selection. Neural computation 4(4), 590–604 (1992)
- [33] McCallumzy, A.K., Nigamy, K.: Employing em and pool-based active learning for text classification. In: Proc. International Conference on Machine Learning (ICML). pp. 359–367. Citeseer (1998)
- [34] Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European conference on computer vision. pp. 69–84. Springer (2016)
- [35] Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
- [36] Pedersen, T., Patwardhan, S., Michelizzi, J., et al.: Wordnet:: Similaritymeasuring the relatedness of concepts. In: AAAI. vol. 4, pp. 25–29 (2004)
- [37] Pourahmadi, K., Nooralinejad, P., Pirsiavash, H.: A simple baseline for lowbudget active learning. arXiv preprint arXiv:2110.12033 (2021)
- [38] Sener, O., Savarese, S.: Active learning for convolutional neural networks: A core-set approach. arXiv preprint arXiv:1708.00489 (2017)
- [39] Settles, B.: Active learning literature survey (2009)
- [40] Shannon, C.E.: A mathematical theory of communication. The Bell system technical journal 27(3), 379–423 (1948)
- [41] Sinha, S., Ebrahimi, S., Darrell, T.: Variational adversarial active learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5972–5981 (2019)
- [42] Spearman, C.: The proof and measurement of association between two things. (1961)
- [43] Tong, S., Koller, D.: Support vector machine active learning with applications to text classification. Journal of machine learning research 2(Nov), 45–66 (2001)
- [44] Wu, X., Dyer, E., Neyshabur, B.: When do curricula work? arXiv preprint arXiv:2012.03107 (2020)
- [45] Yang, Y., Loog, M.: A variance maximization criterion for active learning. Pattern Recognition 78, 358–370 (2018)
- [46] Yoo, D., Kweon, I.S.: Learning loss for active learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 93–102 (2019)
- [47] You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., Shen, Y.: Graph contrastive learning with augmentations. Advances in Neural Information Processing Systems 33, 5812–5823 (2020)
- [48] Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: European conference on computer vision. pp. 649–666. Springer (2016)
- [49] Zhu, Y., Xu, W., Liu, Q., Wu, S.: When contrastive learning meets active learning: A novel graph active learning paradigm with self-supervision. arXiv preprint arXiv:2010.16091 (2020)