

LPLgrad: Optimizing Active Learning Through Gradient Norm Sample Selection and Auxiliary Model Training

Shreen Gul[†], Mohamed Elmahallawy[‡], Sanjay Madria[†], Ardhendu Tripathy[†]

[†]Computer Science Department, Missouri University of Science and Technology, Rolla, MO 65401, USA

[‡] School of Engineering & Applied Sciences, Washington State University, Richland, WA 99354, USA

Emails: sgchr@mst.edu, mohamed.elmahallawy@wsu.edu, madrias@mst.edu, astripathy@mst.edu

Abstract—Machine learning models are increasingly being utilized across various fields and tasks due to their outstanding performance and strong generalization capabilities. Nonetheless, their success hinges on the availability of large volumes of annotated data, the creation of which is often labor-intensive, time-consuming, and expensive. Many active learning (AL) approaches have been proposed to address these challenges, but they often fail to fully leverage the information from the core phases of AL, such as training on the labeled set and querying new unlabeled samples. To bridge this gap, we propose a novel AL approach, Loss Prediction Loss with Gradient Norm (LPLgrad), designed to quantify model uncertainty effectively and improve the accuracy of image classification tasks. LPLgrad operates in two distinct phases: (i) *Training Phase* aims to predict the loss for input features by jointly training a main model and an auxiliary model. Both models are trained on the labeled data to maximize the efficiency of the learning process—an aspect often overlooked in previous AL methods. This dual-model approach enhances the ability to extract complex input features and learn intrinsic patterns from the data effectively; (ii) *Querying Phase* that quantifies the uncertainty of the main model to guide sample selection. This is achieved by calculating the gradient norm of the entropy values for samples in the unlabeled dataset. Samples with the highest gradient norms are prioritized for labeling and subsequently added to the labeled set, improving the model’s performance with minimal labeling effort. Extensive evaluations on real-world datasets demonstrate that the LPLgrad approach outperforms state-of-the-art methods by order of magnitude in terms of accuracy on a small number of labeled images, yet achieving comparable training and querying times in multiple image classification tasks. Our code is available at Github.

Index Terms—Active learning, image classification, uncertainty quantification, loss prediction

I. INTRODUCTION

Machine learning models are being adopted rapidly across various fields due to their exceptional performance and generalization capabilities. These models rely on both data and ground-truth labels to excel in their tasks. However, obtaining ground-truth labels is often challenging. For instance, in medical imaging, domain experts must be compensated to annotate data points, and in speech recognition, labeling audio data at the word level can take significantly more time than the actual speech duration [1]. The manual annotation process is both time-consuming and labor-intensive [2], [3].

Active learning (AL) offers a solution to these challenges by intelligently selecting the most informative data points for

This work was supported by the National Science Foundation (NSF) under Grant No. 2246187.

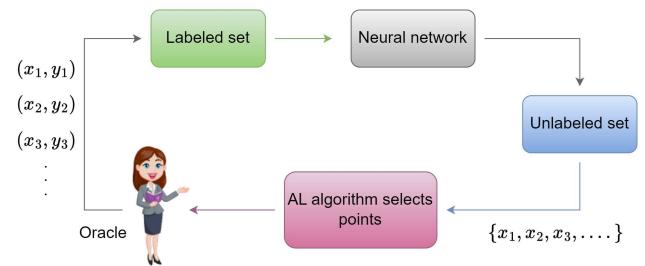


Fig. 1: Typical active learning pipeline.

labeling, thereby reducing the overall annotation effort [4]. In each active learning round, a set of new unlabeled points is selected for annotation and added to the labeled set, and then the target model will be trained on this updated labeled set (Fig. 1 presents an overview of a typical AL process). Various AL methodologies have been proposed in recent years, generally categorized into *uncertainty sampling* and *diversity sampling*. Uncertainty sampling targets data points where the model is most uncertain about their categories, with methods such as entropy sampling [5], margin sampling [6], and least-confidence sampling [5] being popular examples. Diversity sampling, on the other hand, aims to select the most diverse samples that represent the entire dataset, with recent approaches including Coreset [7], variational adversarial AL (VAAL) [8], and Wasserstein adversarial AL (WAAL) [9].

Challenges. Although AL methods reduce traditional manual annotation efforts, they often fail to fully leverage the information from the core phases of AL such as training on labeled set and querying new unlabeled samples. For example, loss prediction loss (LPL), an uncertainty sampling method, integrates a loss prediction module with the target model, yielding two loss values—target loss and prediction loss—through joint learning. The unlabeled points with the highest prediction loss values are then chosen for labeling [10]. However, the hyperparameters in the loss prediction module can be highly sensitive in large-scale datasets like Tiny ImageNet [11] and EMNIST [12], leading to performance degradation [1]. Furthermore, we empirically observe that selecting points based on prediction loss values is *less effective* compared to selecting points based on their entropy values (see section IV).

On the other hand, methods that utilize output entropy to

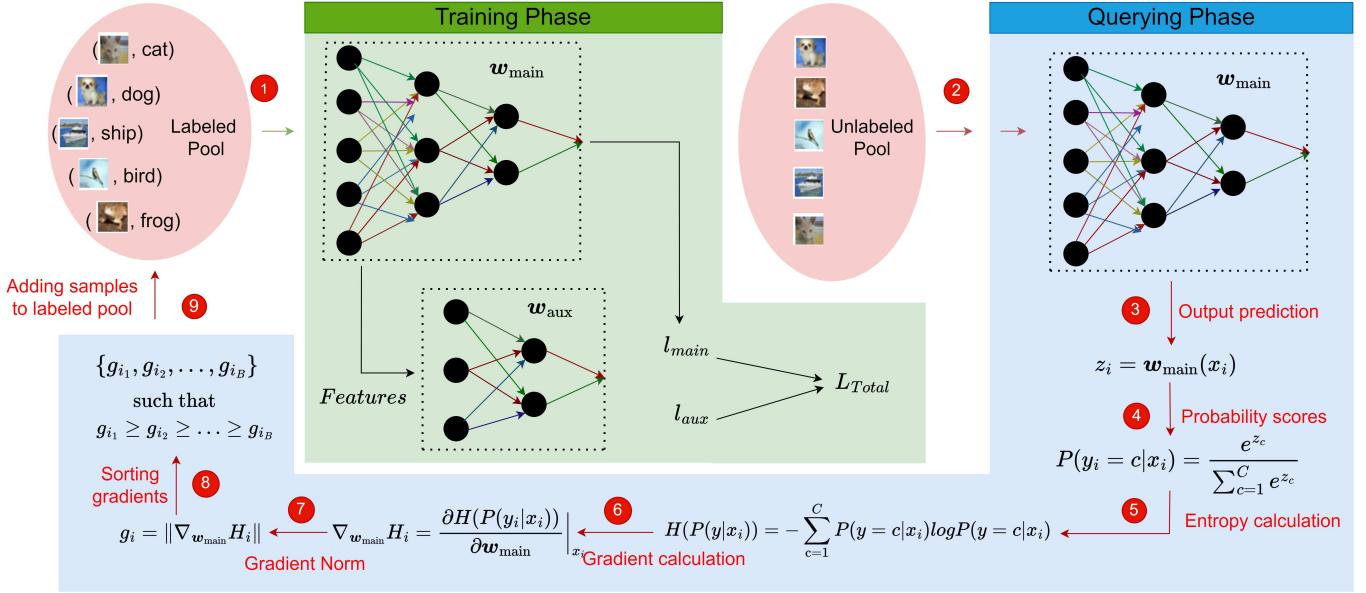


Fig. 2: A visualization of our LPLgrad approach. It highlights the two key components of LPLgrad: training and querying phases. The process involves 9 steps (represented in red circles): In steps 1 and 2, labeled images \mathcal{L} are processed through the main model w_{main} , where feature maps are extracted and subsequently fed into the auxiliary model w_{aux} . This combination of models yields two distinct losses, l_{main} and l_{aux} , which are then aggregated to compute the total loss L_{total} . From steps 3 to 9, samples from the unlabeled set \mathcal{U} are processed, producing scores that are passed through a softmax classifier to obtain entropy values for each sample. These entropy values are back-propagated through w_{main} to compute the gradient. Then, LPLgrad calculates the gradient norm and selects the samples with the highest gradient norms for annotating and adding to \mathcal{L} .

calculate gradient norm, as proposed by Wang et al. [13] emphasize the connection between selected samples and model performance on test data to guide sample selection.

Contributions. In this paper, we propose a novel AL approach named Loss Prediction Loss with Gradient Norm (LPLgrad). LPLgrad incorporates an entropy-based method to quantify model uncertainty, enhancing the accuracy of image classification. Specifically, our proposed AL approach is inspired by [10] but differs in three key ways: (i) We do not use the loss prediction module during the querying phase as done by [10]; instead, we integrate the loss prediction module as an auxiliary model with the main model only during training phase (see the green region in Fig. 2); (ii) During the querying phase we utilize network probability scores instead of input features, as they directly indicate the model’s confidence in its predictions (see the blue region in Fig. 2); (iii) We conduct enhanced training by learning the main model and auxiliary model together, and calculate the output entropy and subsequently the gradient norm of the unlabeled instances as measures of uncertainty, instead of predicted loss values as in [10], to capture the best aspects of both the training and querying phases.

To sum up, our contributions can be summarized as follows:

- We propose a novel AL algorithm called LPLgrad that leverages a loss prediction module to learn input data

features and effectively quantify the network’s uncertainty on unlabeled data based on their gradient norms. To the best of our knowledge, this approach is the first of its kind that addresses a common gap in the literature by utilizing information from *both the training and sampling phases* of AL, enabling more informed and deliberate decisions about sample selection.

- We integrate a main model throughout the training process with the auxiliary model that predicts loss to simultaneously learn the *parametric auxiliary model and main model*. This allows us to effectively extract input features and complex patterns within the data. Subsequently, the model utilizes the gradient norm of unlabeled samples as a metric of uncertainty, aiding in informed decision-making during sample querying.
- Extensive evaluation with different annotation budgets on the visual task such as image classification demonstrates the superior performance of the proposed method against the state-of-the-art AL approaches.

II. RELATED WORK

In this section, we will review recent works in AL, which can be categorized into three main approaches: uncertainty sampling, diversity sampling and hybrid sampling.

A. Uncertainty based Methods

This category of AL methods evaluates the *informativeness* of unlabeled data points by assessing the uncertainty of the target deep network regarding these points. They prioritize selecting those unlabeled points for annotation and adding to the labeled set where the model exhibits significant uncertainty. In this context, Wang et al. [5] introduce a metric for data selection based on model uncertainty, known as entropy sampling. This metric is one of the most widely used for uncertainty quantification and data selection. Some of the recent works that propose active learning approaches include [14]–[16]. For instance, the authors of [14] proposed a technique that incorporates both known and unknown data distributions to measure the model’s uncertainty. Another work, [15] introduced a method that utilizes noise stability in the model’s parameters as an uncertainty metric. A recent approach by [16] estimates the model’s uncertainty by employing a Gaussian process (GP) model as a surrogate for the baseline neural network learner. Another recent work Verified Pseudo-label Selection for Deep Active Learning (VERIPS) proposed by [4] that uses a pseudo-label verification mechanism that consists of a second network only trained on data approved by the oracle and helps to discard questionable pseudo-labels.

While the aforementioned methods effectively reduce labeling effort, they share a common drawback: *they are susceptible to selecting outliers due to their high uncertainty*. Additionally, focusing predominantly on sampling uncertain points can lead to unreliable model predictions and querying redundant data, ultimately decaying model’s performance [2].

B. Diversity based Methods

In this category of AL methods, the learner queries examples that are representative of the entire data distribution such as the work in [7], [8]. In [7], the authors proposed the Coreset approach, which is among the most prominent methods in diversity-based AL. It frames AL as a coresset problem; selecting unlabeled samples based on their geometric properties. Despite its effectiveness, this method is computationally intensive and has delayed sampling times because it requires storing an array of labeled samples for comparison with new samples. The authors of [8] proposed a variational adversarial AL (VAAL) approach that utilizes a variational autoencoder to learn the distribution of labeled data in latent space, coupled with an adversarial network that discriminates between labeled and unlabeled data. However, VAAL necessitates retraining the VAE multiple times rendering it computationally intensive.

While diversity-based methods effectively capture the underlying data distribution, they may fail to fully leverage the information from unlabeled data necessary for training the task learner. Moreover, these techniques might be insensitive to data points near the decision boundary, even though such points could be crucial for the target model to query [2].

In this paper, we propose an uncertainty-based AL algorithm that addresses the limitations above by leveraging both the training and querying phases. Our algorithm effectively learns the features of input data through the joint training of models,

thereby extracting inherent patterns in the input points and reducing the likelihood of selecting redundant data. During the selection phase, it chooses unlabeled points based on their *gradient norm values*, which provably reduces the test loss.

III. PROPOSED METHODOLOGY

In this section, we provide a comprehensive explanation of the components comprising our proposed AL approach, LPLgrad. We start with the problem formulation and an overview of the framework, followed by a detailed description of the two main building blocks of LPLgrad: the Training Phase and the Querying Phase.

A. Problem Formulation and Framework

Given a pool of unlabeled set of data samples denoted as $\mathcal{U} = \{x_i\}_{i=1}^N$ with N is the total number of samples, we aim to solve multi-class classification problem with C categories. To do that, we first construct a labeled set of multi-class $\mathcal{L}_B^t = \{x_i, y_i\}_{i=1}^B$ (y_i represents the label of the data point x_i) by randomly selecting B samples from the unlabeled pool \mathcal{U}_{N-B}^t . Here, the superscript $t = 0, 1, \dots$ signifies the current round of AL, which increases by one as the training progresses. We then utilize a set of models $\mathbf{w} = \{\mathbf{w}_{\text{main}}, \mathbf{w}_{\text{aux}}\}$ and train them on \mathcal{L}_B^t . The training procedure involves an augmented approach where both the main model \mathbf{w}_{main} and the auxiliary model \mathbf{w}_{aux} are jointly learned.

Once the training on the selected set \mathcal{L}_B^t is completed, we compute the output entropy for all the samples in \mathcal{U}_{N-B}^t (which includes only the remaining unselected samples). These entropy values represent the loss incurred by \mathbf{w}_{main} . Subsequently, we update the \mathbf{w}_{main} parameters and store the gradient norm for each sample in the set. These stored values will be then sorted, and a new B' samples with the highest gradient norm will be selected for labeling in the next AL cycle ($t+1$).

In the subsequent cycle, the updated labeled and unlabeled sets are denoted as $\mathcal{L}_{B'}^{t+1}$ and $\mathcal{U}_{N-B'}^{t+1}$, respectively. Then, the main and auxiliary models will be trained on $\mathcal{L}_{B'}^{t+1}$ to update their model weights as $\mathbf{w}_{\text{main}}^{t+1}, \mathbf{w}_{\text{aux}}^{t+1}$, respectively. This process of training and querying new samples continues in subsequent AL cycles until a certain accuracy threshold is achieved, a predefined budget of iterations is exhausted, or any other termination criterion.

Below, we provide a detailed description of our proposed LPLgrad framework (see red numbered circles in Fig. 2):

Step ① In each AL round t , we select a set of labeled images \mathcal{L}_B^t to feed both the main model \mathbf{w}_{main} and the auxiliary model \mathbf{w}_{aux} . The main model \mathbf{w}_{main} is designed to extract features from the selected labeled images and then feed these features to the auxiliary model \mathbf{w}_{aux} . Subsequently, both models \mathbf{w}_{main} and \mathbf{w}_{aux} produce losses, which are jointly learned to generate the total loss $l_{\text{total}} = l_{\text{main}} + l_{\text{aux}}$.

Step ② Next, the remaining unlabeled images \mathcal{U}_{N-B}^t will be fed into the trained \mathbf{w}_{main} to calculate the prediction scores.

Step ③ Then, the main model w_{main} outputs a $1 \times C$ vector of prediction scores, where C is the number of classes in the dataset for our multi-class problem.

Step ④ These output prediction scores are then transformed into a vector of posterior probabilities of the same size using a softmax classifier.

Step ⑤ After that, we use these resulting probabilities to compute the entropy (a scalar value) of the model w_{main} for each sample x_i in the unlabeled set \mathcal{U}_{N-B}^t .

Step ⑥ In this step, the entropy value is treated as a loss that the model w_{main} incurs for that particular sample. Subsequently, the gradient of this loss is calculated with respect to the model's parameters. The size of this gradient corresponds to the number of layers in the model.

Step ⑦ Subsequently, the norm of the gradient value computed in the previous step is determined as $\|\nabla_{w_{\text{main}}} H_i\|$.

Step ⑧ These gradient norms are then sorted, and the unlabeled samples with the highest gradient norms are selected for annotation and inclusion in the labeled set \mathcal{L}_B^t .

Step ⑨ These newly queried samples are annotated by the oracle and included in the labeled pool, which is subsequently used to train the main model.

The above steps will be repeated in each AL round r until a termination criterion is met such as achieving a target accuracy.

B. LPLgrad Training Phase

The training phase of LPLgrad involves two primary models: the main model w_{main} and the auxiliary model w_{aux} . LPLgrad trains the model w_{main} alongside w_{aux} , which is integrated into its architecture to effectively capture intricate patterns and characteristics within input data. Here's how it works:

For each data point x_i , we obtain two values: one is the prediction of the main model $y_{\text{main}} = w_{\text{main}}(x_i)$, and the other is a feature map F , which is processed by auxiliary the model w_{aux} to output the predicted loss $l_{\text{aux}} = w_{\text{aux}}(F)$. The loss of the main model w_{main} is calculated using cross-entropy loss that takes the predicted value y_{main} and ground-truth label y_i of the sample x_i as inputs, which can be expressed as

$$l_{\text{main}} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{\text{CE}}(y_i, y_{\text{main}}) \quad (1)$$

The loss for the auxiliary model w_{aux} is computed based on the predicted loss l_{aux} and its corresponding ground-truth loss value l_{main} , which can be presented as:

$$l_{\text{aux}} = \frac{1}{P} \sum_{i=1}^P \max(0, M - d_i \cdot (l_{\text{aux},i} - l_{\text{main},i})) \quad (2)$$

Here, M is a parameter for margin which explains how much the predicted loss should differ from the ground-truth loss before a penalty is applied. d in the equation is used to

Algorithm 1: LPLgrad Training Phase

Input: Labeled pool \mathcal{L} , main model w_{main} , auxiliary module w_{aux} , number of AL rounds \mathcal{T} , number of epochs in each round \mathcal{E}

Output: w_{main}^T

- 1 Initialize w_{main} , w_{aux}
- 2 **for** $t = 0, 1, 2, \dots, \mathcal{T}$ **do**
- 3 **for** $e = 0, 1, 2, \dots, \mathcal{E}$ **do**
- 4 Feedforward the input data of \mathcal{L} to w_{main} and extract the features of input images as well as the output predictions of the w_{main} .
- 5 Calculate the loss of the main model l_{main} using equation (1).
- 6 Feedforward the extracted features to the w_{aux} to obtain l_{aux} using equation (2).
- 7 Add both of these losses l_{main} and l_{aux} using equation 4.
- 8 Train both the w_{main} and w_{aux} in conjunction using the calculated losses.
- 9 Call Algorithm 2
- 10 **return** The final model w_{main}^T for the round \mathcal{T}

determine the direction of the margin penalty and is computed as

$$d_i = \max(0, l_{\text{main},i}) \quad (3)$$

This ensures that the margin is adjusted correctly, either penalizing or not penalizing the predicted loss, depending on the relative difference between the predicted and true losses. To make the auxiliary model w_{aux} robust to overall scale variations in the loss, we construct a mini-batch of P examples from \mathcal{L}_B . We form $P/2$ data pairs, denoted as $\{x^p = (x_m, x_n)\}$, where x^p represents a pair of examples m and n . The superscript p indicates the loss for a pair of data points, denoted as $l_{\text{aux},i}$ and $l_{\text{main},i}$ for the auxiliary model and the main model, respectively, as shown in equation (2).

Note that $l_{\text{aux},i}$ in equation (2) represents the predicted loss for a specific sample in the pair, which is obtained by processing the extracted features of input through the w_{aux} while the overall loss for auxiliary model is denoted by l_{aux} . The auxiliary model w_{aux} is learned by comparing the differences between the predicted losses l_{main}^p and l_{aux}^p for each data pair.

The total loss during the training phase is then calculated as follows:

$$L_{\text{total}} = l_{\text{aux}} + l_{\text{main}} \quad (4)$$

LPLgrad leverages the extraction of multi-level input features obtained from various layers of the main model w_{main} , which are subsequently fed into the auxiliary model w_{aux} . Specifically, the model w_{aux} comprises a series of blocks corresponding to the layers within the model w_{main} . Each block consists of two distinct layers: a global average pooling layer and a fully-connected layer. These blocks process the feature

Algorithm 2: LPLgrad Querying Phase

Input: Unlabeled data pool \mathcal{U} , initial labeled pool \mathcal{L} , main model \mathbf{w}_{main} , number of AL rounds \mathcal{T} , annotation budget in each round B

Output: $\mathbf{w}_{\text{main}}^T$

- 1 **for** $t = 0, 1, 2, \dots, \mathcal{T}$ **do**
- 2 Calculate the entropy of each sample in \mathcal{U} using equation 6.
- 3 Calculate the gradient of this calculated entropy with respect to main model \mathbf{w}_{main} parameters using equation 7 and subsequently its norm.
- 4 Select B samples with the highest values for gradient norm from \mathcal{U} .
- 5 Annotate these selected data points and add them to \mathcal{L} .
- 6 Call algorithm 1
- 7 **return** main model parameters $\mathbf{w}_{\text{main}}^T$ for the round \mathcal{T}

maps F derived from the layers of \mathbf{w}_{main} model, producing scalar values representing predicted losses l_{aux} . Our objective is to jointly minimize both the predicted loss l_{aux} generated by the model \mathbf{w}_{aux} and the actual loss l_{main} determined by the model \mathbf{w}_{main} . This optimization strategy enables the model to not only discern relevant input features but also effectively integrate rich, multi-level input space information.

Algorithm 1 summarizes the LPLgrad's training phase.

C. LPLgrad Querying Phase

After training the main model \mathbf{w}_{main} alongside the auxiliary model \mathbf{w}_{aux} , LPLgrad transitions to its second phase, which focuses on querying new samples for labeling. While the method proposed by Yoo et al. [10] utilizes the \mathbf{w}_{aux} model to identify the most informative points, our empirical results reveal a more effective strategy. Specifically, selecting samples based on their entropy values presents a robust alternative to the aforementioned approach. The hyperparameters in the loss prediction module can be highly sensitive in large-scale datasets like [11] and EMNIST [12] leading to performance degradation. Moreover, the model's uncertainty is better estimated using its entropy because it incorporates the model's confidence scores directly rather than with an attached loss prediction module. Specifically, selecting samples based on their gradient norm values presents a robust alternative to the aforementioned approach.

To implement this, we begin by extracting the embeddings of each sample in the unlabeled set $x_i \in \mathcal{U}_B$ such that $z_i = \mathbf{w}_{\text{main}}(x_i)$. We then use a softmax classifier to obtain the posterior probabilities $P(y_i|x_i)$, which can be expressed as

$$P(y_i = c|x_i) = \frac{e^{z_c}}{\sum_{c=1}^C e^{z_c}} \quad (5)$$

These posterior probabilities are then used to calculate the output entropy of each sample, which can be given as

$$H(P(y_i|x_i)) = - \sum_{c=1}^C P(y_i = c|x_i) \log P(y_i = c|x_i) \quad (6)$$

where $P(y_i = c|x_i)$ is the predicted probability for class c given the sample x_i , and C is the total number of classes in the dataset.

We treat this entropy as a loss and perform a backward pass on the loss function to compute the gradient of the \mathbf{w}_{main} model parameters for each sample x_i :

$$\nabla_{\mathbf{w}_{\text{main}}} H_i = \frac{\partial H(P(y_i|x_i))}{\partial \mathbf{w}_{\text{main}}} \Big|_{x_i} \quad (7)$$

The Frobenius norm of these gradients is then calculated to assess the network's sensitivity to the input as

$$g_i = \|\nabla_{\mathbf{w}_{\text{main}}} H_i\|_F \quad (8)$$

These gradient norms of all the inputs are subsequently stored and sorted to identify the highest values, which can be represented as

$$\{g_{i_1}, g_{i_2}, \dots, g_{i_B}\} \quad \text{s.t.} \quad g_{i_1} \geq g_{i_2} \geq \dots \geq g_{i_B} \quad (9)$$

In our querying phase, we follow the theoretical insights presented in [13], which suggest that selecting samples with higher gradient norms from the unlabeled set can lead to a reduction in the upper bound of the total loss. Thus, we guide our selection process by prioritizing samples with the largest gradient norms for annotation. Subsequently, these selected samples are added to the labeled set \mathcal{L} . Importantly, the parameters of the \mathbf{w}_{main} model will not be updated during the selection process of the new samples.

Algorithm 2 summarizes the LPLgrad's querying phase.

IV. EXPERIMENTS

In this section, we evaluate the performance of LPLgrad across various datasets and compare its accuracy with several baseline methods. Additionally, we assess LPLgrad's computational performance, including both querying and training time.

A. Experimental Setup

To ensure accurate results, we average all the experiments over 5 trials and report the mean outcomes. Below, we describe our setup in detail:

Datasets. We evaluate our proposed approach, LPLgrad, using four publicly available benchmark datasets: CIFAR-10 [17], CIFAR-100 [17], SVHN [6], Caltech-101 [18], which are commonly used in state-of-the-art (SOTA) comparisons. In addition, We validate LPLgrad on another real-world dataset, *comprehensive disaster dataset (CDD)* [19], to assess its robustness. Below, we provide a brief description of each dataset.

TABLE I: Datasets, Models, and Training Parameters.

Dataset	Size	Classes	Model	Parameters	Epochs	mini-batch
CIFAR 10	70,000	10	ResNet-18	1,100,000	200	128
CIFAR-100	70,000	10	ResNet-18	1,100,000	200	128
SVHN	73257	10	ResNet-18	1,100,000	200	128
Caltech101	6110	101	ResNet-18	1,100,000	50	80
CDD	8677	6	ResNet-18	1,100,000	50	128
$\eta = 0.001 - 0.1$		SGD Momentum= 0.9				

- **CIFAR-10 & CIFAR-100 [17]:** are two datasets, with the former containing 10 different classes and the latter containing 100 classes. Each dataset comprises 60,000 color images, evenly distributed across the classes. Specifically, in CIFAR-10, each class has 5000 samples, whereas in CIFAR-100, each class has 600 samples. The images are all 32×32 pixels in size and feature various objectives such as animals, vehicles, etc.
- **SVHN [6]:** The SVHN (Street View House Numbers) dataset comprises a total of 630,420 colored images of house numbers, categorized into 10 classes for each digit. The images are of size 32×32 , similar to CIFAR-10 and CIFAR-100. The dataset is divided into three sets: 73,257 images for training, 26,032 images for testing, and an extra set containing 530,420 images. For a fair comparison with other AL methods, we do not use the extra set provided in the SVHN dataset.
- **Caltech-101 [18]:** It contains images of objects belonging to 101 categories, making it a highly imbalanced dataset with over 100 classes. For instance, there are about 40 to 800 images per category, with most categories having around 50 images. It contains 5800 training while 2877 test images with each image being roughly 300×200 pixels in size.
- **Comprehensive Disaster Dataset (CDD) [19]:** This dataset consists of a total of 10,733 images, categorized into fire disaster, human damage, land disaster, water disaster, and etc. It is divided into two main sets: 8,591 training images and 2,142 test images, with a total of 6 classes. It is an extremely imbalanced dataset, with the number of samples per class varying from 29 to 1,668. These images have a resolution of 32×32 pixels.

Baselines. We compare our proposed LPLgrad algorithm with state-of-the-art AL approaches discussed in the literature [7], [10], [13], [20]. Specifically, we consider the two main categories: i) uncertainty-based methods such as LearningLoss [10], Ent-GradNorm [13], and entropy [5], and ii) diversity-based methods such as coresnet [7], which have been reviewed in Sections I and II. In addition, we compare LPLgrad with non-AL strategies, where unlabeled samples are selected randomly to augment the labeled set. We refer to this method as random sampling (Rand).

Training Models. To ensure a fair comparison with our baselines, we train an 18-layer residual network (ResNet-18) as our main model w_{main} to perform image classification across all the experiments. We adhere to the settings and

TABLE II: Values of B , R and A for all datasets.

	CIFAR10	CIFAR100	SVHN	Caltech101	CDD
B	1,000	2,500	1,000	500	500
R	25,000	25,000	25,000	2,000	1,000
A	10,000	20,000	10,000	3,500	2,500
B	500	1,000	500	200	200
R	5,000	15,000	10,000	1,000	1,000
A	5,000	8,000	5,000	1,400	1,000

hyperparameters recommended by the baseline studies for reproducing their results. Additionally, we employ stochastic gradient descent (SGD) as the optimizer, with a momentum of 0.9 for all datasets. The rest of the training settings for each dataset are as follows:

- For CIFAR10 and CIFAR100, we use a learning rate (lr) of 0.1 and train both w_{main} and w_{aux} models for 200 epochs, reducing the lr by a factor of 0.1 after 160 epochs.
- For SVHN, we use a lr of 0.01 and train both w_{main} and w_{aux} models for 200 epochs, reducing the lr by 0.1 after 40 epochs.
- For Caltech-101, the w_{main} and w_{aux} models are trained for 50 epochs with a lr of 0.01, which is decayed by 0.1 after 40 epochs.
- For CDD, the w_{main} and w_{aux} models are trained for 200 epochs with a lr of 0.01, which is decayed by 0.1 after 40 epochs.

For training the ResNet-18 model, we begin with a randomly selected initial pool of samples B from the large unlabeled pool to be annotated and transferred to the labeled set, which is then added to the training dataset. However, we note that this random selection might result in overlapping images and similar selections each time. To mitigate this, we employ the technique used in [21] to obtain a random subset $S_R \subset \mathcal{U}_N$ from the pool of unlabeled samples. This simple technique proves to be efficient as it reduces redundancy. The value of R is adjusted for each dataset, depending on the total number of samples and the size of B , considering differences in dataset size and class imbalance, as each class within the dataset has a different number of data points.

In Table II, we detail the query batch size B , the size of the randomly chosen set R , and the annotation budget (A) for all experiments across all datasets. The values for A and R have been chosen following the standard practice in AL literature [22]–[25]. The upper part of the table provides details for a high A (with plots shown in Fig. 3), while the lower part corresponds to a low A (with plots shown in Fig. 4).

B. Results

Comparing LPLgrad with baselines. We first evaluate the performance of our LPLgrad approach against the baselines in terms of image classification with a high A . In Table III, we present detailed results showing that LPLgrad consistently outperforms its counterparts across almost all datasets. Notably, LPLgrad achieves a significant margin of superiority, especially with challenging datasets such as CIFAR-

TABLE III: Comparison of LPLgrad vs. baselines on all datasets: CIFAR10, CIFAR100, SVHN, and Caltech-101 for high A.

Actively chosen data	CIFAR10			CIFAR100			SVHN			Caltech-101		
	3000	6000	9000	7500	12500	20000	3000	6000	9000	1500	2500	3500
LearningLoss	71.3 \pm 1.1	86.0 \pm 0.2	89.7 \pm 0.3	43.7 \pm 1.1	60.6 \pm 0.4	66 \pm 0.5	90 \pm 0.3	93.9 \pm 0.1	95.1 \pm 0.1	36.0 \pm 0.6	46.8 \pm 0.5	53.5 \pm 1.0
Entropy	65.2 \pm 2.9	84.1 \pm 0.8	88.9 \pm 0.4	44.4 \pm 1.0	60.4 \pm 0.3	66.5 \pm 0.2	89.6 \pm 0.3	93.7 \pm 0.1	94.9 \pm 0.1	36.9 \pm 0.4	45.6 \pm 0.8	54.1 \pm 1.5
Ent-GradNorm	68.0 \pm 1.3	84.8 \pm 0.4	89.0 \pm 0.2	46.6 \pm 0.8	61.7 \pm 0.5	66.7 \pm 0.2	89.7 \pm 0.3	93.4 \pm 0.1	94.8 \pm 0.1	38.3 \pm 1.4	47.5 \pm 1.6	54.4 \pm 1.4
Coreset	64.0 \pm 1.9	81.7 \pm 0.6	86.0 \pm 0.3	44.3 \pm 1.0	59.8 \pm 0.8	64.5 \pm 0.6	82.3 \pm 0.3	87.6 \pm 0.1	89.7 \pm 0.1	37.8 \pm 1.3	45.5 \pm 2.1	51.5 \pm 2.0
Rand	71.7 \pm 0.8	83.8 \pm 0.7	86.2 \pm 0.5	44.4 \pm 1.8	59.3 \pm 0.6	64.3 \pm 0.1	81.9 \pm 0.5	87.5 \pm 0.1	89.6 \pm 0.0	36.0 \pm 1.4	44.1 \pm 1.6	50.8 \pm 1.7
LPLgrad (Ours)	72.4 \pm 1.1	86.2 \pm 0.2	90.1 \pm 0.3	46.8 \pm 1.1	62.8 \pm 0.4	67 \pm 1.4	90.2 \pm 0.1	93.7 \pm 0.1	95.1 \pm 0.0	38.6 \pm 1.9	49.1 \pm 2.0	56.6 \pm 1.2

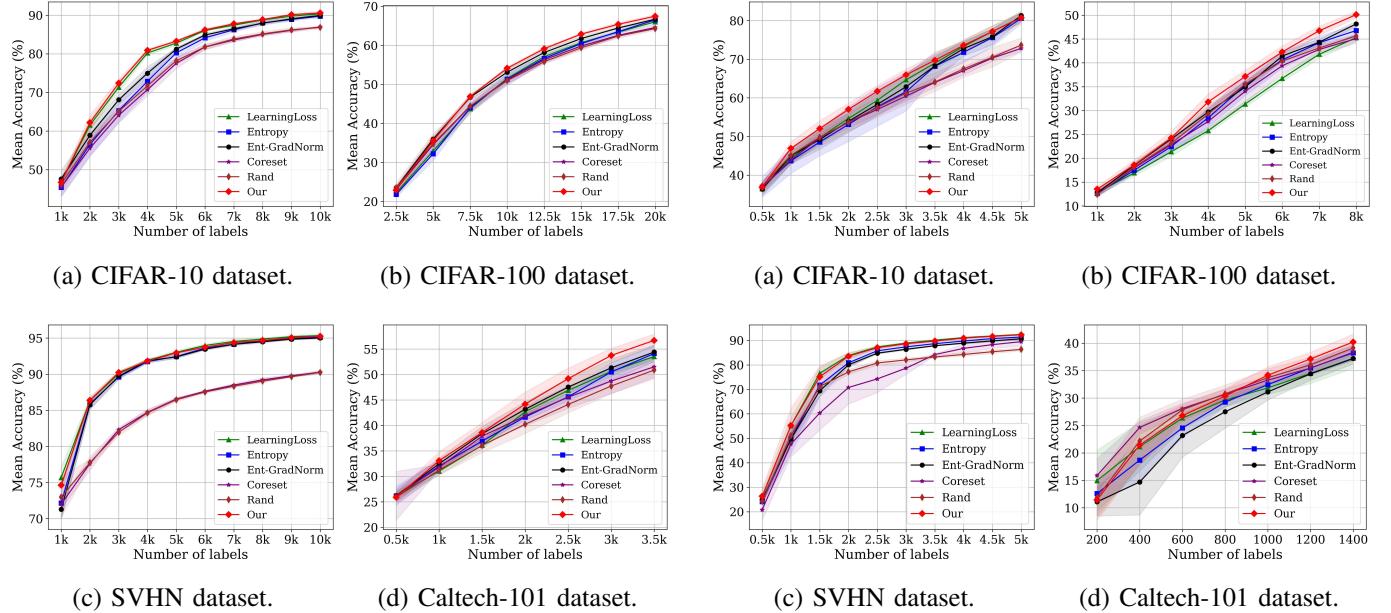


Fig. 3: Classification performance of LPLgrad compared to baseline methods on CIFAR-10, CIFAR-100, SVHN, and Caltech-101 datasets with high A.

100 and Caltech101. For example, on the Caltech101 dataset, which is highly imbalanced with over 100 classes, LPLgrad demonstrates a clear advantage over all baselines, improving accuracy by approximately 5%. LearningLoss [10] emerges as the second-best performing method, sometimes surpassing Ent-GradNorm [13]. Entropy often ranks third, followed by Coreset and Rand. Interestingly, Coreset performs similarly to Rand in almost all scenarios.

Evaluating the effect of A. In Fig. 3 and 4, we present the classification accuracy of LPLgrad across various datasets under two different annotation budget (A) settings: high and low. These results are compared against our baseline approaches. For instance, in Fig. 3(a), we show the accuracy curves of our baselines when classifying the CIFAR-10 dataset. Our approach, LPLgrad, outperforms the baselines across all As, achieving the highest accuracy of 91% by the final round. LearningLoss [10] is the second-best performing method, followed by Ent-GradNorm [13] and entropy [5]. Entropy lags behind Ent-GradNorm for the first seven rounds but catches up in the last three rounds. Coreset [7], the only diversity-based method, performs worse than all other methods in the initial

half of the AL rounds, eventually aligning with Rand in the latter half.

For the CIFAR-100 dataset, which is the most challenging due to its 100 fine-grained classes, Fig. 3(b) demonstrates that LPLgrad has superior performance compared to all baselines, reaching a high accuracy of 68%. Here, Ent-GradNorm is the second-best performing algorithm, followed by entropy sampling, indicating their effectiveness in a large-scale dataset setting. LearningLoss performs comparably to entropy until the end. Coreset and Rand are the worst-performing methods. As a purely diversity-based method, Coreset does not outperform the uncertainty-based methods like entropy, Ent-GradNorm, and LPLgrad.

Fig. 3(c) represents the accuracy curves of LPLgrad on the SVHN dataset. All algorithms perform well on this dataset, including our LPLgrad approach, achieving $\geq 95\%$ accuracy, except for Coreset and Rand, which achieve an accuracy of 90%. LearningLoss performs worse than LPLgrad in the first half of the rounds but slightly leads in the later rounds. Entropy and Ent-GradNorm dip around the fifth AL round but stabilize and perform on par with other algorithms. Finally, Fig. 3(d)

TABLE IV: Comparison of computational overheads (in seconds) of our LPLgrad against baselines across different datasets.

	CIFAR10		CIFAR100		SVHN		Caltech101		CDD	
	Train Time	Querying Time	Train Time	Querying Time	Train Time	Querying Time	Train Time	Querying Time	Train Time	Querying Time
LearningLoss	63.7	15.7	99.7	95	124.9	210	280.9	75	63.9	40
Entropy	61.9	14.29	97.5	93	124.4	200.4	252.5	50	59.0	35
Ent-GradNorm	63.1	40	97.9	115	135	202.4	182.5	56.6	55.7	37.4
Coreset	68.42	2319	97.1	1123	130	2124	197.4	676	63.1	54.5
LPLgrad (Ours)	60.5	14	92	92.3	123	215	252	49.1	60	34.8

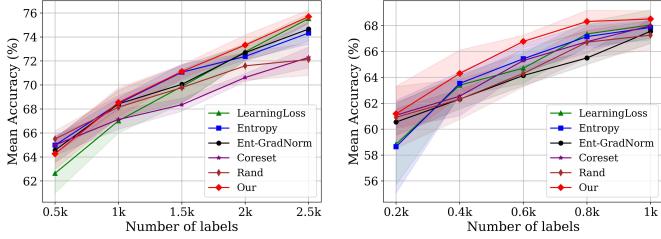


Fig. 5: Classification performance comparison on Comprehensive Disaster Dataset.

compares LPLgrad against the baselines on the Caltech-101 dataset. LPLgrad clearly outperforms the others, achieving 57.4% accuracy. Ent-GradNorm follows as the second-best performer. Entropy and LearningLoss initially perform poorly but improve and end up close to Ent-GradNorm. Coreset is the second-worst AL method, followed by Rand. Uncertainty-based methods consistently lead over diversity-based methods. However, all approaches, including LPLgrad, show the lowest performance on the Caltech-101 dataset compared to the other datasets, suggesting that increasing the A could improve results, albeit with more training time needed.

Examining the results with a low A in Fig. 4, LPLgrad consistently achieves the highest accuracy, outperforming all baselines across all datasets. An exception is observed in the Caltech-101 dataset, where the Coreset method initially achieves higher accuracy than LPLgrad. However, as the learning curves converge, LPLgrad again surpasses the baseline methods in accuracy.

Evaluating LPLgrad’s accuracy on a challenging dataset. We evaluate LPLgrad on a more challenging dataset, CDD, in Fig. 5 to demonstrate its effectiveness even with difficult datasets compared to our baselines. In Fig. 5(a), which depicts results with a high budget A , LPLgrad starts with an accuracy of around 64%. By the first AL round, it surpasses all other baselines and continues to achieve the best performance across all rounds, reaching close to 75%. LearningLoss also shows strong performance, ending as the second-best performing algorithm. Ent-GradNorm attempts to catch up for almost half the rounds but then its performance degrades significantly. A similar trend is seen with entropy. Coreset and Rand both start well, but their performance declines over the rounds, with Coreset ending up performing worse than Rand.

Fig. 5(b) presents the results of CDD classification in a

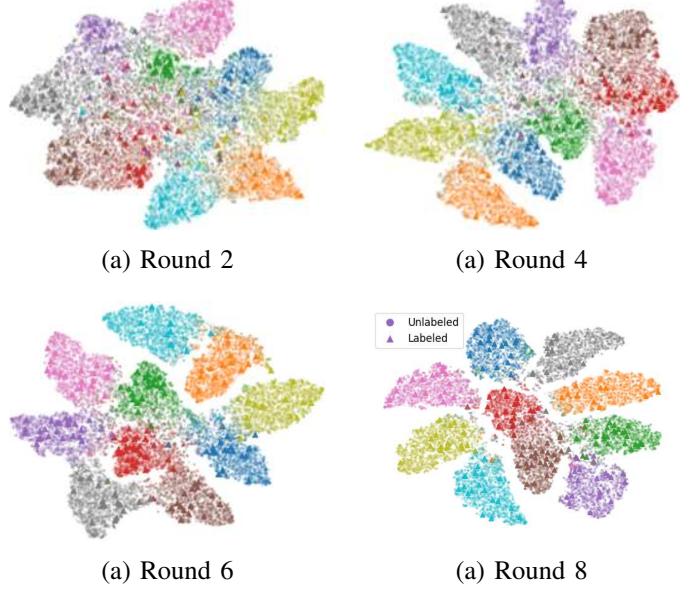
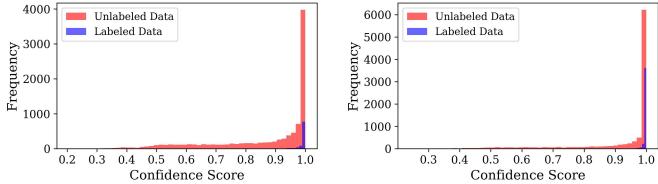


Fig. 6: t-SNE visualization of feature embeddings generated by LPLgrad on the CIFAR-10 dataset.

low-budget regime. Here, LPLgrad consistently outperforms other methods throughout all AL rounds, achieving 69% accuracy. Due to the low-budget regime, the overall scale of accuracy is lower compared to the high annotation budget scenario. Entropy performs fairly well, securing the second-best position. Interestingly, Ent-GradNorm ends up last in this case. Coreset, random, and entropy continue to alternate their positions until the final round.

Comparing LPLgrad’s Computing time against baselines. In tabel IV, we show the computation time required during the training and querying phases of LPLgrad against our baselines (using Nvidia A100 80GB RAM GPUs). Our observations indicate that Coreset is the most time-consuming algorithm due to its need to store all the data and compare new points with this stored data. On the other hand, LPLgrad performs comparably or even better than Entropy, Ent-GradNorm, and LearningLoss. While in some cases, LPLgrad may require more querying time due to backpropagation and parameter updates, this trade-off leads to improved accuracy, rendering the additional time investment worthwhile.

Evaluating LPLgrad’s feature embeddings. We first present in Fig. 6 the t-distributed stochastic neighbor embedding (t-SNE) visualization of LPLgrad to illustrate the distribution of input features across four distinct AL rounds (2, 4, 6, 8). The visualization notably demonstrates that, *after only*



(a) Round 2

(b) Round 4

(c) Round 6

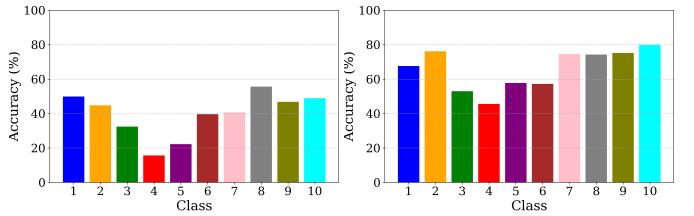
(d) Round 8

Fig. 7: Confidence scores of w_{main} across various AL rounds on the CIFAR-10 dataset.

8 rounds, LPLgrad effectively labels CIFAR-10 images into their respective 10 classes. LPLgrad exhibits not only a strong discriminative ability in differentiating between instances of various classes but also shows enhanced proficiency in accurately assigning labels to these instances as the AL rounds progress.

Evaluating LPLgrad’s Confidence score. Next, in Fig. 7, we present the frequency distribution of the model’s confidence scores for labeled and unlabeled samples of CIFAR-10 dataset, particularly those close to 0.99, across the same AL rounds (2, 4, 6, 8). We plot the confidence scores for 10,000 unlabeled images in all the subfigures, while the number of labeled images increases by 1,000 in each subsequent graph, starting from 1,000. Specifically, Fig. 7(a) illustrates that the confidence scores for the unlabeled samples are more widely distributed across the confidence intervals. There are approximately 4,000 images for which the model is confident in their predictions, but there are 1,000 images with a confidence frequency of less than 0.99. In contrast, for the labeled images, almost 700 out of 1,000 have a confidence score between 0.99 and 1. Fig. 7(b) shows the frequency of scores for the fourth AL round. Evidently, 6,000 unlabeled images now have confidence scores between 0.99 and 1. Similarly, about 3,500 out of 4,000 are predicted correctly for the labeled samples, hence decreasing the spread of both the labeled and unlabeled samples’ confidence scores. Moving to Fig. 7(c), we observe a further decrease in the spread of bins, indicating that the model is more confident with both the labeled and unlabeled images. It gives scores between 0.9 and 1 for about 7,500 unlabeled and 6,000 labeled images. In Fig. 7(d), we show that the labeled images have surpassed the unlabeled ones, with about 8,500 and 8,200 out of 10,000, respectively, having scores between 0.99 and 1.

Evaluating LPLgrad’s class accuracy. In Fig. 8, we analyze the accuracy of LPLgrad across all 10 classes in the CIFAR-10 dataset across the four AL classes. Specifically,



(a) Round 2

(b) Round 4

(c) Round 6

(d) Round 8

Fig. 8: Per class accuracy values of LPLgrad on CIFAR10.

Fig. 8 shows the results of the second AL round, where the model exhibits a lack of confidence, with the highest accuracy reaching only 55.7%. Moving to Fig. 8(b), which depicts the fourth AL round, we observe significant improvements across all classes, with the model achieving a peak accuracy of 80%. Fig. 8(c) illustrates the sixth AL round, where the model’s proficiency continues to improve, achieving the highest accuracy of 88.6%. Finally, Fig. 8(d) shows the results of the eighth AL round, where the model demonstrates a thorough understanding of the dataset, achieving consistently high accuracies across all classes, with the highest accuracy reaching 93.3%.

C. Ablation Study

In this subsection, we delve into the significance of incorporating an auxiliary model to enhance the accuracy of our training pipeline. We explore the impact of integrating this auxiliary model by comparing results obtained from standalone training (only using the main model w_{main}) with those achieved when the auxiliary model is used. The primary role of the auxiliary model is to bolster the main model’s capacity to learn and process input features more efficiently. It achieves this by introducing an additional set of layers that are trained concurrently with the main model. This setup allows the auxiliary model to capture supplementary information and nuances from the data, which the main model alone might overlook. By doing so, we aim to improve the overall performance, robustness, and generalization of the main model. The following sections will present detailed results and analyses of these approaches to highlight the tangible benefits of including the auxiliary model in our training regimen.

In Table V and VI, we present a comparison under two different annotation budget settings, A , across four distinct datasets. These results illustrate that the auxiliary model consistently supports the main model in achieving higher performance throughout the entire training process. Specifically,

TABLE V: Accuracy values across different epochs with and without the auxiliary model model for **low** A.

	with w_{aux}			w/o w_{aux}		
	1/3 of A	1/2 of A	Entire A	1/3 of A	1/2 of A	Entire A
CIFAR-10	52.1	65.7	77.4	49.8	63.3	76.2
CIFAR-100	24.6	43.7	50.2	24.1	42.3	47.7
SVHN	75.2	88.7	92.4	69.7	86.3	91.2
Caltech-101	22.5	31.4	37.6	14.9	27.6	34.8

TABLE VI: Accuracy values across different epochs with and without the auxiliary model for **high** A.

	with w_{aux}			w/o w_{aux}		
	1/3 of A	1/2 of A	Entire A	1/3 of A	1/2 of A	Entire A
CIFAR-10	72.2	85.5	90.8	69.8	83.7	89.9
CIFAR-100	47.7	63.4	68.7	47.6	62.1	67.9
SVHN	91	93.5	95.1	89.8	93.3	95
Caltech-101	33.5	44.2	53.8	33.4	43.8	51.2

the inclusion of the auxiliary model leads to an approximate improvement in accuracy of around 8%. This enhancement is observed across all training epochs, demonstrating the auxiliary model's effectiveness in improving the main model's performance and stability. The tables provide detailed evidence of the auxiliary model's positive impact on both accuracy and result consistency, regardless of the training phase.

V. CONCLUSION AND FUTURE WORK

In this paper, we proposed LPLgrad, a novel AL approach aimed at addressing a common gap in the literature: the underutilization of the core phases of AL, specifically the training and querying phases. To fully exploit the labeled data, we adopted an augmented approach wherein two models—a main model and an auxiliary model—are trained together to optimally learn the features of the input data. Additionally, to effectively query the most informative samples, we compute the entropy values of the unlabeled set and backpropagate these values to obtain loss values. This loss is minimized, and its gradients with respect to the main model's parameters are calculated. The Frobenius norm of gradients is then computed, sorted, and used to identify the samples with the highest gradient norm values, which are selected for labeling and added to the labeled set. We extensively evaluate LPLgrad on diverse image classification datasets and a real-world dataset to validate its efficacy. Our findings demonstrate that LPLgrad surpasses state-of-the-art approaches by achieving higher accuracy with fewer labels and less computing time.

As future work, we plan to explore the application of LPLgrad to more complex and larger-scale datasets to further validate its robustness and scalability. We also aim to investigate the potential of LPLgrad in domains beyond image classification, such as natural language processing or time-series data, which will be a key focus. Finally, we will explore using LPLgrad in real-time systems where rapid decision-making with limited labeled data is critical.

REFERENCES

- X. Zhan *et al.*, “A comparative survey of deep active learning,” *arXiv preprint arXiv:2203.13450*, 2022.
- D. Li, Z. Wang, Y. Chen, R. Jiang, W. Ding, and M. Okumura, “A survey on deep active learning: Recent advances and new frontiers,” *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
- S. Gul, M. Elmahallawy, S. Madria, and A. Tripathy, “Fishermask: Enhancing neural network labeling efficiency in image classification using fisher information,” *arXiv preprint arXiv:2411.05752*, 2024. [Online]. Available: <https://arxiv.org/abs/2411.05752>
- S. Gilhuber, P. Jahn, Y. Ma, and T. Seidl, “Verips: Verified pseudo-label selection for deep active learning,” in *2022 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2022, pp. 951–956.
- D. Wang and Y. Shang, “A new active labeling method for deep learning,” in *2014 International joint conference on neural networks (IJCNN)*. IEEE, 2014, pp. 112–119.
- Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, A. Y. Ng *et al.*, “Reading digits in natural images with unsupervised feature learning,” in *NIPS workshop on deep learning and unsupervised feature learning*, vol. 2011, no. 5. Granada, Spain, 2011, p. 7.
- O. Sener and S. Savarese, “Active learning for convolutional neural networks: A core-set approach,” *arXiv preprint arXiv:1708.00489*, 2017.
- S. Sinha, S. Ebrahimi, and T. Darrell, “Variational adversarial active learning,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 5972–5981.
- C. Shui, F. Zhou, C. Gagné, and B. Wang, “Deep active learning: Unified and principled method for query and training,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1308–1318.
- D. Yoo and I. S. Kweon, “Learning loss for active learning,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 93–102.
- Y. Le and X. Yang, “Tiny imagenet visual recognition challenge,” *CS 231N*, vol. 7, no. 7, p. 3, 2015.
- G. Cohen, S. Afshar, J. Tapson, and A. van Schaik, “Emnist: an extension of mnist to handwritten letters,” *arXiv preprint arXiv:1702.05373*, 2017.
- T. Wang, X. Li, P. Yang, G. Hu, X. Zeng, S. Huang, C.-Z. Xu, and M. Xu, “Boosting active learning via improving test performance,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 8, 2022, pp. 8566–8574.
- B. Safaei, V. Vibashan, C. M. de Melo, and V. M. Patel, “Entropic open-set active learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 5, 2024, pp. 4686–4694.
- X. Li, P. Yang, Y. Gu, X. Zhan, T. Wang, M. Xu, and C. Xu, “Deep active learning with noise stability,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 12, 2024, pp. 13 655–13 663.
- Y. An, S. Park, and K. I. Kim, “Active learning guided by efficient surrogate learners,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 10, 2024, pp. 10 874–10 881.
- A. Krizhevsky, V. Nair, and G. Hinton, “Cifar-10 (canadian institute for advanced research),” URL <http://www.cs.toronto.edu/kriz/cifar.html>, vol. 5, no. 4, p. 1, 2010.
- L. Fei-Fei *et al.*, “Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories,” *Computer Vision and Pattern Recognition Workshop*, 2004.
- A. Verma, “Disaster images dataset, version 1,” May 2021, retrieved June 2, 2024 from [kaggle.com/datasets/varpit94/disaster-images-dataset/data](https://www.kaggle.com/datasets/varpit94/disaster-images-dataset/data).
- W. Luo, A. Schwing, and R. Urtasun, “Latent structured active learning,” *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- W. H. Beluch, T. Genewein, A. Nürnberger, and J. M. Köhler, “The power of ensembles for active learning in image classification,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9368–9377.
- S. Kothawade, N. Beck, K. Killamsetty, and R. Iyer, “Similar: Submodular information measures based active learning in realistic scenarios,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 18 685–18 697, 2021.
- G. Hacohen, A. Dekel, and D. Weinshall, “Active learning on a budget: Opposite strategies suit high and low budgets,” *arXiv preprint arXiv:2202.02794*, 2022.
- E. Zhao, A. Liu, A. Anandkumar, and Y. Yue, “Active learning under label shift,” in *International conference on artificial intelligence and statistics*. PMLR, 2021, pp. 3412–3420.
- F. Wan, T. Yuan, M. Fu, X. Ji, Q. Huang, and Q. Ye, “Nearest neighbor classifier embedded network for active learning,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 11, 2021, pp. 10 041–10 048.