
Probabilistic Bilevel Coreset Selection

Xiao Zhou^{*1} Renjie Pi^{*1} Weizhong Zhang^{*1} Yong Lin¹ Tong Zhang^{1,2}

Abstract

The goal of coreset selection in supervised learning is to produce a weighted subset of data, so that training only on the subset achieves similar performance as training on the entire dataset. Existing methods achieved promising results in resource-constrained scenarios such as continual learning and streaming. However, most of the existing algorithms are limited to traditional machine learning models. A few algorithms that can handle large models adopt greedy search approaches due to the difficulty in solving the discrete subset selection problem, which is computationally costly when coreset becomes larger and often produces suboptimal results. In this work, for the first time we propose a continuous probabilistic bilevel formulation of coreset selection by learning a probabilistic weight for each training sample. The overall objective is posed as a bilevel optimization problem, where 1) the inner loop samples coresets and train the model to convergence and 2) the outer loop updates the sample probability progressively according to the model’s performance. Importantly, we develop an efficient solver to the bilevel optimization problem via unbiased policy gradient without trouble of implicit differentiation. We provide the convergence property of our training procedure and demonstrate the superiority of our algorithm against various coreset selection methods in various tasks, especially in more challenging label-noise and class-imbalance scenarios.

1. Introduction

In the last decade, deep neural networks (DNNs) have achieved tremendous successes in multiple areas such as computer vision (Simonyan & Zisserman, 2015; He et al.,

^{*}Equal contribution ¹The Hong Kong University of Science and Technology ²Google Research. Correspondence to: Tong Zhang <tongzhang@tongzhang-ml.org>.

2016) and natural language processing (Vaswani et al., 2017). These superior performances are mostly achieved via learning from huge amounts of data. However, this data-driven paradigm also poses several new challenges: 1) the cumbersome dataset becomes harder to store and transfer; 2) for some real applications, such as continual learning, one can only access a small number of training data at each stage of training; 3) in some more extreme scenarios, where the training data is incorrectly labelled, or they are collected from different domains, more training data may even hurt the model’s performance. To address these issues, a natural idea is to select a small subset (i.e., *coreset*) comprised of most informative training samples, such that training on this subset can achieve comparable or even better performance with that on the full dataset, which is verified in Appendix D. Therefore, how to construct a good coreset for DNNs now becomes a crucial problem.

We notice that, coreset selection has been investigated for the traditional machine learning models, e.g., SVM (Tsang et al., 2005), logistic regression (Huggins et al., 2016) and Gaussian mixture model (Lucic et al., 2017), for a long time to accelerate the training process and lots of effective methods have been developed. The idea of these studies is to find a small weighted training subset, whose objective function is close to the one of the full training set at any point in the parameter space. A typical formulation is uniform function approximation, that is to find a small subset $\hat{\mathcal{D}}$ together with the non-negative weights $\hat{\mathcal{W}} = \{w_i : i \in \hat{\mathcal{D}}\}$ from the training dataset \mathcal{D} satisfying that

$$|\mathcal{L}(\theta) - \hat{\mathcal{L}}(\theta)| \leq \epsilon \mathcal{L}(\theta), \text{ for any } \theta \in \mathbb{R}^p, \quad (1)$$

where $\mathcal{L}(\theta) = \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} \ell(\theta; \mathbf{x}_i, \mathbf{y}_i)$ is the objective function on the full dataset and $\hat{\mathcal{L}}(\theta) = \frac{1}{|\hat{\mathcal{D}}|} \sum_{i \in \hat{\mathcal{D}}} w_i \ell(\theta; \mathbf{x}_i, \mathbf{y}_i)$ is the one on the selected subset $\hat{\mathcal{D}}$, ϵ is a small number to control the approximation error. When ϵ is small enough, the comparable performance of the model learned from $\hat{\mathcal{L}}$ can be guaranteed. Nevertheless, it has been shown that these methods cannot be applied to DNNs directly (Borsos et al., 2020). The reason is that as DNNs are always highly nonconvex and the hypothesis set is significantly larger than traditional models, to obtain small uniform approximation error ϵ , one has to select a very large coreset $\hat{\mathcal{D}}$.

More recently, bilevel optimization (Borsos et al., 2020) has

been introduced to construct the coreset for DNNs. Their motivation is that the only thing we really care about is the performance of the model trained on the coreset, i.e., the optimum of $\hat{\mathcal{L}}$, instead of achieving small approximation error for the loss function in the *whole* parameter space. Therefore, the bilevel optimization with a cardinality constraint on $\hat{\mathcal{D}}$ presented below would be a more reasonable framework for constructing the coreset:

$$\begin{aligned} & \min_{\hat{\mathcal{W}}, \hat{\mathcal{D}} \subset \mathcal{D}, |\hat{\mathcal{D}}| \leq K} \mathcal{L}(\theta^*(\mathcal{W}, \hat{\mathcal{D}})) \\ \text{s.t. } & \theta^*(\hat{\mathcal{W}}, \hat{\mathcal{D}}) = \arg \min_{\theta} \hat{\mathcal{L}}(\theta). \end{aligned} \quad (2)$$

However, despite some promising empirical results are reported, we notice that existing methods adopt greedy strategies to update $\hat{\mathcal{D}}$ as the outer loop involves optimization on the discrete variable set $\hat{\mathcal{D}}$, which is NP hard. For example, (Borsos et al., 2020) starts from several randomly sampled training data and then sequentially adds sample with the largest gradient with respect to the outer objective into $\hat{\mathcal{D}}$ in each iteration. Such greedy methods generally lead to suboptimal performance and the selected coreset would be unnecessarily large. The drawback becomes more prominent if the initially sampled data has bad quality, which may misguide the selection, especially when the dataset contains noisy labels or is class-imbalanced.

In this paper, to address the above issue, we propose a probabilistic coreset selection method based on bilevel optimization for DNNs. In contrast with the greedy methods, our key idea is to continualize the discrete bilevel optimization problem above by probabilistic reparameterization, making the gradient-based optimization possible. In this way, we can explore the entire dataset and progressively improve the quality of the coreset during training. To be precise, we first assign each training data i with a binary mask m_i to indicate whether data i is included in the coreset or not. Then, to continualize the problem, we parameterize m_i to be a Bernoulli variable with the probability s_i to be 1 and $1 - s_i$ to be 0. Thus, the cardinality of the coreset can be roughly controlled by the sum of s_i and coreset selection is transformed into the problem of learning these probabilities s_i . Therefore, we formulate coreset selection as a continuous bilevel optimization problem with a sparsity constraint (Eqn. (4)). In the inner loop, we sample a coreset according to the probability s_i and use it to train the model (see Eqn. (5)). In the outer loop, we minimize the loss of the learned model on the full dataset by adjusting the sample probabilities. We develop an efficient optimization algorithm (Algorithm 1) utilizing unbiased policy gradient estimator, which calculates the probabilities' gradients via only forward instead of backward propagation and avoids the complicate calculation of implicit differentiation.

Notably, our proposed method has the following advantages:

- 1) Our algorithm gradually improves the quality of the selected subset by explores the training set globally unlike greedy methods which make urgent decisions at early stage and cannot remove redundant data once added to the coreset.
- 2) Our algorithm obtains much more competence in the setting with label noise or class imbalance without the bother of making decisions on adding detrimental or redundant data at early stage. In contrast, the misselected data may mislead the future selection for greedy methods.
- 3) Our algorithm develops an efficient policy gradient solver to the bilevel optimization problem without cumbersome implicit gradient calculations.

Moreover, we provide the property of our training algorithm in convergence. We demonstrate the superiority of our method through extensive experimental results on various tasks, including data summarization, continual learning, streaming and feature selection and surpass state-of-the-art methods by a large margin.

Our main contributions can be summarized as follows:

- To the best of our knowledge, our method is the first global coreset selection method for DNNs, where we propose a novel continualized bilevel optimization formulation and develop an efficient policy gradient solver without implicit gradient calculation.
- We provide the convergence property showing that our optimization method can converge similarly with the standard nonconvex projected stochastic gradient descent algorithm.
- We empirically demonstrate the superiority of our method through various experiments including data summarization, continual learning, streaming, feature selection, especially in more challenging scenarios with noisy labels and class imbalance.

2. Related Works

2.1. Coreset Selection

Coreset selection aims to solve the problem of finding the most informative subset from the full set, which can be used to solve the optimization problem and obtain similar performance as the full set. Several coreset selection methods are designed for specific learning algorithms, such as K-means (Feldman & Langberg, 2011; Har-Peled & Kushal, 2007), SVM (Tsang et al., 2005), logistic regression (Huggins et al., 2016) and Gaussian mixture model (Lucic et al., 2017; Feldman et al., 2011). These methods only work for traditional models and can not be directly adopted for DNNs. (Borsos et al., 2020) proposed a bilevel optimization framework for coreset selection for DNN. However, due to

the difficulty in solving the discrete optimization problem, they adopt greedy search algorithm to sequentially select new samples. Since the samples selected in the early stage can never be removed afterwards, their method generally results in suboptimal performance and deteriorates severely in challenging scenarios, where the initial samples can have bad quality and misguide the search. Moreover, the greedy-based algorithm faces expensive computational cost due to its demand in solving a bilevel optimization problem for every sample added to the coreset. Our method searches the coreset globally and does not suffer from those drawbacks.

2.2. Continual Learning and Streaming

Continual learning (CL) (Kirkpatrick et al., 2017; Lopez-Paz & Ranzato, 2017; Rebuffi et al., 2017) aims to tackle the scenario where a series of different tasks are learnt sequentially using the same model. In this work, we mainly focus on replay-based CL methods, which keeps a constant number of data of previous tasks to alleviate the catastrophic forgetting problem. Streaming (Aljundi et al., 2019b; Hayes et al., 2019; Chrysakis & Moens, 2020) is more challenging in the sense that it does not have the concept of tasks and data is sequentially given to the model. In these cases, coresets can be adopted to construct the replay memory to choose the informative data which well represents each task.

2.3. Feature Selection

Feature selection (FS) (Cai et al., 2018; Li et al., 2017; Miao & Niu, 2016) aims to select a subset of important features to represent the original data, which reduces the computation and storage cost. The majority of existing works of feature selection are mainly focused on traditional models (Güneş et al., 2010; Sulaiman & Labadin, 2015; Radovic et al., 2017). However, to the best of our knowledge, the exploration in this direction for deep neural networks is limited. We note that feature selection can also be considered as an instance level coreset selection task which can be naturally addressed by our proposed framework, where the selected features can be viewed as the coreset.

2.4. Dataset Distillation

An alternative approach for dataset compression is dataset distillation (Wang et al., 2018; Nguyen et al., 2020), which is inspired by knowledge distillation (Hinton et al., 2015; Gou et al., 2021; Yao et al., 2021a). Instead of distilling knowledge from the model parameters, dataset distillation learns a few synthetic data points for each class. These methods work well when the model for deployment is the same as the one used for learning the synthetic data. However, as the learnt data points also encode information of the model’s architecture and initialization weights, their performances drop significantly when the synthetic data learned on one model is used to train another model. In contrast, coreset

selection is not sensitive to the model, since we do not alter the data directly.

2.5. Bilevel Optimization

Bilevel optimization (Sinha et al., 2017) has garnered a lot of attention in recent years due to its ability to handle hierarchical decision making processes. Previous works utilize bilevel optimization in multiple areas of research, such as hyper-parameter optimization (Lorraine et al., 2020; Maclaurin et al., 2015; Pedregosa, 2016; MacKay et al., 2019; Franceschi et al., 2017; Vicol et al., 2021), meta learning (Finn et al., 2017; Nichol & Schulman, 2018), neural architecture search (Pham et al., 2018; Liu et al., 2018; Pham et al., 2018; Shi et al., 2020; Yao et al., 2021b; Gao et al., 2022; 2021; Shi et al., 2021) and sample re-weighting (Ren et al., 2018; Shu et al., 2019; Wang et al., 2022a; Zhou et al., 2022). Prior to our work, (Borsos et al., 2020) formulates coreset selection into a bi-level optimization problem and solves it using a greedy algorithm.

3. Probabilistic Bilevel Coreset Selection

In this section, we first present our probabilistic coreset selection framework in Section 3.1 and then develop an efficient training method for this framework in Section 3.2.

3.1. Bilevel Framework for Coreset Selection

Consider a neural network $f(\mathbf{x}; \boldsymbol{\theta})$ with $\boldsymbol{\theta}$ being the trainable parameters and $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ is the training dataset, we first formulate coreset selection into the following discrete bilevel optimization paradigm:

$$\begin{aligned} \min_{\mathbf{m} \in \tilde{\mathcal{C}}} \tilde{\Phi}(\mathbf{m}) &= \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) = \frac{1}{n} \sum_{i=1}^n \ell(f(\mathbf{x}_i; \boldsymbol{\theta}^*(\mathbf{m})), \mathbf{y}_i), \quad (3) \\ \text{s.t. } \boldsymbol{\theta}^*(\mathbf{m}) &\in \arg \min_{\boldsymbol{\theta}} \hat{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{m}) = \frac{1}{K} \sum_{i=1}^n m_i \ell(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i), \end{aligned}$$

where the mask $\mathbf{m} \in \{0, 1\}^n$ is a binary vector with $m_i = 1$ indicating sample i is selected into the coreset and otherwise excluded. K is a positive integer controlling the coreset size and $\tilde{\mathcal{C}} = \{\mathbf{m} : m_i = 0 \text{ or } 1, \|\mathbf{m}\|_0 \leq K\}$ is the feasible region of \mathbf{m} . Intuitively, the inner loop trains the network to converge on the selected coreset to obtain the model $\boldsymbol{\theta}^*(\mathbf{m})$. The outer loop evaluates the loss of $\boldsymbol{\theta}^*(\mathbf{m})$ on the full set and optimizes it to guide the learning of \mathbf{m} .

Remark 3.1. The difference between our discrete bilevel formulation (3) and the existing one in (2) is that our formulation has no weight w_i for each sample in the coreset. We remove these weights for two considerations: 1) our empirical results show that we can achieve good performance without weighting the coreset; 2) it enables us to develop extremely efficient training algorithm (see Section 5.3).

Noticing that the discrete nature of the mask \mathbf{m} makes

directly solving the above bilevel optimization problem intractable, we now turn to continualize it by probabilistic reparameterization, making gradient based optimization method possible. Our main idea is to view each mask m_i as an independent binary random variable and transform the problem (3) from optimizing in the discrete vector space into the probability space, which is continuous. Specifically, we reparameterize m_i as a Bernoulli random variable with probability s_i to be 1 and $1 - s_i$ to be 0, that is $m_i \sim \text{Bern}(s_i)$, where $s_i \in [0, 1]$. Assuming the variables m_i are independent, then we can get the distribution function of \mathbf{m} , i.e., $p(\mathbf{m}|\mathbf{s}) = \prod_{i=1}^n (s_i)^{m_i} (1 - s_i)^{(1-m_i)}$. Thus, we can control the coreset size via the sum of the probabilities s_i , i.e., $\mathbf{1}^\top \mathbf{s}$, since $\mathbb{E}_{\mathbf{m} \sim p(\mathbf{m}|\mathbf{s})} \|\mathbf{m}\|_0 = \sum_{i=1}^n s_i$. Therefore, $\tilde{\mathcal{C}}$ can be approximately transformed into $\mathcal{C} = \{\mathbf{s} : 0 \preceq \mathbf{s} \preceq \mathbf{1}, \|\mathbf{s}\|_1 \leq K\}$. Finally, problem (3) can be naturally relaxed into the following expected loss minimization problem:

$$\min_{\mathbf{s} \in \mathcal{C}} \Phi(\mathbf{s}) = \mathbb{E}_{p(\mathbf{m}|\mathbf{s})} \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})), \quad (4)$$

$$s.t. \boldsymbol{\theta}^*(\mathbf{m}) \in \arg \min_{\boldsymbol{\theta}} \hat{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{m}) \quad (5)$$

where $\mathcal{C} = \{\mathbf{s} : 0 \preceq \mathbf{s} \preceq \mathbf{1}, \|\mathbf{s}\|_1 \leq K\}$ is the domain.

Some appealing features of our Formulation (4) are:

- Our formulation is a tight relaxation (although not equivalent) of Problem (3). The reasons are:
 - It is easy to know that $\min_{\mathbf{s} \in \mathcal{C}} \Phi(\mathbf{s}) \leq \min_{\mathbf{m} \in \tilde{\mathcal{C}}} \tilde{\Phi}(\mathbf{m})$ as any deterministic binary mask \mathbf{m} can be represented as a particular stochastic one by letting s_i be either 0 or 1.
 - Our constraint \mathcal{C} induces sparsity on \mathbf{s} due to the ℓ_1 -norm and the range $[0, 1]$, making most components of the optimal \mathbf{s} either 0 or 1. That is, our finally learned stochastic mask is nearly deterministic, which will be empirically verified in Section 5.2.
- Due to our sparsity constraint, the selected coreset size of the inner loop, i.e., $\|\mathbf{m}\|$ is always small, which makes the optimization of $\boldsymbol{\theta}^*$ very efficient.
- As shown in Eqn.(6), our outer objective is $\Phi(\mathbf{s})$ is differentiable, allowing us to use general gradient based methods for optimization.

3.2. Optimization

Existing bilevel optimization algorithms (Pedregosa, 2016; Grazi et al., 2020; 2021) are often computationally costly due to the expensive implicit differentiation in their chain-rule based gradient estimator. To be precise, if applied to our problem, they generally estimate the gradient in the form of

$$\nabla_{\mathbf{s}} \Phi(\mathbf{s}) \approx \nabla_{\mathbf{s}} \boldsymbol{\theta}^*(\mathbf{m}) \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})).$$

Hence, they need to compute the implicit differentiation of the inner loop optimum, i.e, $\nabla_{\mathbf{s}} \boldsymbol{\theta}^*(\mathbf{m})$, which is expensive since they have to compute the inverse of a huge hessian matrix or unroll the backward propagation for multiple steps.

Even though some efficient bilevel optimization algorithms have been proposed to alleviate the computational burden, for instance, (Lorraine et al., 2020) adopted Neumann series to approximate the hessian inverse, the approximation still requires much time and leads to inefficiency.

Thanks to our probabilistic formulation of the bilevel problem, we are able to avoid these expensive computations by using Policy Gradient Estimator (PGE), which calculates the gradient using forward instead of backward propagation. Our key idea can be illustrated by the following equations:

$$\begin{aligned} \nabla_{\mathbf{s}} \Phi(\mathbf{s}) &= \nabla_{\mathbf{s}} \mathbb{E}_{p(\mathbf{m}|\mathbf{s})} \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) \\ &= \nabla_{\mathbf{s}} \int \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) p(\mathbf{m}|\mathbf{s}) d\mathbf{m} \\ &= \int \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) \frac{\nabla_{\mathbf{s}} p(\mathbf{m}|\mathbf{s})}{p(\mathbf{m}|\mathbf{s})} p(\mathbf{m}|\mathbf{s}) d\mathbf{m} \\ &= \int \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) \nabla_{\mathbf{s}} \ln p(\mathbf{m}|\mathbf{s}) p(\mathbf{m}|\mathbf{s}) d\mathbf{m} \\ &= \mathbb{E}_{p(\mathbf{m}|\mathbf{s})} \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) \nabla_{\mathbf{s}} \ln p(\mathbf{m}|\mathbf{s}). \end{aligned} \quad (6)$$

It shows that $\mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) \nabla_{\mathbf{s}} \ln p(\mathbf{m}|\mathbf{s})$ is an unbiased stochastic gradient of $\nabla_{\mathbf{s}} \Phi(\mathbf{s})$, which is called policy gradient. Therefore, given the inner loop optimum $\boldsymbol{\theta}^*(\mathbf{m})$, we can update \mathbf{s} by projected stochastic gradient descent:

$$\mathbf{s} \leftarrow \mathcal{P}_{\mathcal{C}}(\mathbf{s} - \eta \mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m})) \nabla_{\mathbf{s}} \ln p(\mathbf{m}|\mathbf{s})). \quad (\text{PGE})$$

It is clear that PGE does not involve any implicit differentiation and its component $\mathcal{L}(\boldsymbol{\theta}^*(\mathbf{m}))$ can be computed via forward propagation. Moreover, $\ln p(\mathbf{m}|\mathbf{s})$ has a very simple form and this projection has a closed form solution (given in the appendix B) since the constraint \mathcal{C} is quite simple. Therefore, we can update \mathbf{s} via PGE very efficiently.

Remark 3.2. As we mentioned in Section 3.1, we remove the weights from the original framework (2) because: 1) we empirically find that training using the coreset with binary weights can already achieve competitive performance, and 2) if the weights are not removed, then $\boldsymbol{\theta}(\mathbf{m})$ would be $\boldsymbol{\theta}(\mathbf{w}, \mathbf{m})$ and we have to compute implicit differentiation to get $\nabla_{\mathbf{w}} \boldsymbol{\theta}^*(\mathbf{w}, \mathbf{m})$, since the gradient of deterministic variable cannot be estimated via PGE.

Hence, we can solve our bilevel optimization problem (4) by alternatively: 1) sampling a mask \mathbf{m} , i.e., a coreset, from $p(\mathbf{m}|\mathbf{s})$ for the inner loop and train the model on this coreset to get $\boldsymbol{\theta}^*(\mathbf{m})$; 2) updating the probability \mathbf{s} using PGE. The detailed steps are given in Algorithm 1. Notably, our algorithm has the following advantages:

Algorithm 1 Probabilistic Bilevel Coreset Selection

Require: a network θ , dataset \mathcal{D} and coreset size K .

- 1: Initialize probabilities $\mathbf{s}^1 = \frac{K}{|\mathcal{D}|} \mathbf{1}$.
- 2: **for** training iteration $t = 1, 2 \dots T$ **do**
- 3: Sample mask \mathbf{m} according to the probability \mathbf{s}^t .
- 4: Train the inner loop to converge satisfies:

$$\theta^*(\mathbf{m}) \leftarrow \arg \min_{\theta} \hat{\mathcal{L}}(\theta; \mathbf{m})$$

- 5: Sample a mini-batch of data:

$$\mathcal{B} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_B, \mathbf{y}_B)\}$$

- 6: Update \mathbf{s} using PGE based on $\theta^*(\mathbf{m})$ and \mathcal{B} :

$$\mathbf{s}^{t+1} \leftarrow \mathcal{P}_{\mathcal{C}}(\mathbf{s}^t - \eta \mathcal{L}_{\mathcal{B}}(\theta^*(\mathbf{m})) \nabla_{\mathbf{s}} \ln p(\mathbf{m} | \mathbf{s}^t))$$

- 7: **end for**

output The coreset $\{(\mathbf{x}_i, \mathbf{y}_i) : m_i \neq 0 \text{ and } (\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}\}$ with \mathbf{m} sampled from $p(\mathbf{m} | \mathbf{s}^{T+1})$.

- While the greedy methods can never remove the redundant data once they are added to the coreset, our algorithm behaves like a process of sampling coreset with replacement (Step 4), where the quality of the selected subset is evaluated on the outer objective and the sampling probability \mathbf{s} is adjusted accordingly (Step 6) to progressively improve the coreset quality. This enables us to explore the training set more globally.
- Due to our sparsity constraint \mathcal{C} , most of the probabilities s_i would automatically converge to either 0 or 1 during optimization, thus the uncertainty of the obtained coreset can be finally reduced to nearly 0. This is empirically verified in Section 5.2.
- The superiority of our method is more prominent in more challenging tasks involving data with corrupted labels and class imbalance, which are shown in the Section 4. The reason is that in those scenarios, more explorations are needed to gain a global view of the entire set before deciding which samples to be added to the coreset, whereas greedy algorithms have to start making decisions in the early stage without enough knowledge (demonstrated in Section 5.1).
- As we discussed above, our method is computationally efficient because 1) for the outer loop, PGE enables updating the probability without computing any implicit differentiation; 2) for the inner loop, the selected coreset size is always small, which makes the derivation of θ^* efficient; and 3) as opposed to the greedy algorithm, the running time of our method does not increase rapidly with the coreset size, since the number of outer updates is fixed for all coreset sizes.

The property below shows that if we solve the inner loop problem to convergence, our training algorithm can converge similarly with the standard nonconvex projected stochastic gradient descent algorithms (Ghadimi et al., 2016).

Property 3.3. [Informal] Under the mild assumptions on $\Phi(\mathbf{s})$ and the step size η , then the average of the expectation of the gradient mapping norm, i.e.,

$$\left\| \frac{1}{\eta} (\mathbf{s}^t - \mathcal{P}_{\mathcal{C}}(\mathbf{s}^t - \eta \nabla_{\mathbf{s}} \Phi(\mathbf{s}^t))) \right\|_2,$$

can converge to a small value as $T \rightarrow \infty$.

4. Experiments

We conduct the following experiments in common application scenarios of coreset selection: 1) data summarization, where the selected coreset is directly used to train the model; 2) continual learning (Kirkpatrick et al., 2017; Lopez-Paz & Ranzato, 2017; Rebuffi et al., 2017) and streaming (Aljundi et al., 2019b; Hayes et al., 2019; Chrysakis & Moens, 2020), where coresets are selected from training data to construct the replay memory and resist catastrophic forgetting after sequentially learning a series of tasks; 3) feature selection (Cai et al., 2018; Li et al., 2017; Miao & Niu, 2016), where only a subset of features are selected for training and inference.

In reality, the quality of training data can not be guaranteed. For instance, the data for each task of continual learning may not be balanced or even mislabelled. To test the effectiveness of our method on these challenging scenarios, for each of those applications mentioned above, we create more difficult settings by imposing label noise and class imbalance in the training data. We observe that other approaches including the greedy coreset method (Borsos et al., 2020) fails significantly in such settings, while our method can still discover promising coresets, which is credited to the learnt global information before constructing the final coreset.

4.1. Data Summarization

We examine the quality the coreset constructed by our method via evaluating the performance of the neural network trained on the produced coresets. Specifically, we conduct experiments on two widely used benchmarks, i.e., MNIST (Deng, 2012) and CIFAR10 (Krizhevsky et al., 2009). To make a fair comparison, we follow (Borsos et al., 2020) to use the same model settings (shown in the Appendix C) when conducting the experiments.

We compare with the following competitive baselines: 1) Uniform sampling, 2) K-center clustering using the embedding from last layer (Sener & Savarese, 2017), 3) iCaRL’s selection (Rebuffi et al., 2017), 4) Hardest sampling (Aljundi et al., 2019a) and 5) Greedy coreset (Borsos et al., 2020).




Figure 1. Performance comparison between our method and other baselines on data summarization task with various coreset sizes and different scenarios. For experiments with label noise and class imbalance, the coreset size is set to 1000 for MNIST and 5000 for CIFAR10. Our method consistently surpasses other baselines by a large margin. Notably, the performance of our method is stable even under challenging settings, while other methods begin to fail significantly.

The results in the first column of Figure 1 demonstrate that our method consistently outperforms other baselines across various coreset sizes. Notably, our method’s superiority is more prominent in the small-sized region. It is interesting that hardest sampling fails for data summarization tasks, which can be that the majority of hard samples come from only a few classes and therefore causes the selected coreset to be unbalanced.

4.2. Data Summarization with Label Noise and Class Imbalance

We further conduct experiments in more challenging and practical scenarios, where the dataset contains corrupted labels and class-imbalanced data to showcase the effectiveness of our method. The model setting is the same as stated in Section 4.1 and the outer objective is calculated based on a held-out balanced validation dataset with 100 samples, comprised of 10 uniformly sampled data from each class.

For the label noise experiment, we adopt 2 types of noises: pairwise noise and symmetric noise. For the class imbalance experiment, we adopt similar setting as in Cui et al. (2019). The detailed descriptions are given in the Appendix C due to space limit. The coreset sizes for the MNIST and CIFAR10 experiments are 1000 and 5000, respectively.

In these challenging settings, our algorithm’s advantage over other methods becomes much more prominent, as shown in Figure 1. Notably, our method is much less sensitive to the quality of entire dataset, which can be credited to the global information learnt by repeatedly sampling before

constructing the final coreset. On the other hand, the greedy counterpart begins to fail dramatically as the label noise ratio and imbalance factor grow higher, which is in line with our intuition that greedy methods can be severely affected if the samples in the early phase are redundant or even detrimental. More analysis is given in Section 5.1.

4.3. Continual Learning and Streaming

Two important applications of coreset construction are continual learning (Kirkpatrick et al., 2017; Lopez-Paz & Ranzato, 2017; Rebuffi et al., 2017; Wang et al., 2022b;c) and streaming (Aljundi et al., 2019b; Hayes et al., 2019; Chrysakis & Moens, 2020; Zou et al., 2019). Specifically, continual learning aims to learn a series of tasks sequentially using the same model, and a constant number of data can be reserved for previous tasks to alleviate the catastrophic forgetting of early knowledge. It is thus essential to choose the most informative data for each task when constructing the replay memory. Streaming is similar to continual learning except that the data stream is not divided into tasks, which is more challenging.

Continual Learning. We conduct experiments on the following datasets commonly adopted by the community: 1) PermMNIST (Goodfellow et al., 2013) constructs 10 tasks by performing different random permutation on the image pixels for each task; 2) SplitMNIST (Zenke et al., 2017) splits MNIST into five tasks, each containing two adjacent classes; and 3) SplitCIFAR10 is based on CIFAR10 and splitted in the same way as SplitMNIST. To make a fair

comparison, we use the same settings as adopted in (Borsos et al., 2020), the details are shown in the Appendix C. Moreover, we design more challenging CL tasks with label noise and class imbalance to showcase the superiority of our method, where symmetric noise with 20% noise ratio and class imbalance factor of 50 is applied to each dataset.

We compare our algorithm with data selection methods mentioned in Section 4.1. The results demonstrated in Table 1 verify that our proposed method consistently dominates other baselines across all tasks. Our advantage is more prominent especially under challenging settings, where the performance of other methods drop significantly and our method continues to perform well. Remarkably, our method surpasses the greedy counterpart (Borsos et al., 2020) by around 10% on tasks with label noise.

Streaming. We follow Borsos et al. (2020) to conduct experiment in streaming setting (details are in the Appendix). We compare our method with Reservoir sampling (Vitter, 1985) and greedy coreset (Borsos et al., 2020). We also conduct experiments on label noise scenario, where symmetric label noise with 20% noise ratio is imposed to the data. The results in Table 2 demonstrate that our method dominates other baselines on these streaming tasks.

The success on CL and streaming tasks proves that our method can select informative data that well represent each task to construct the replay memory, which performs well even under challenging conditions.

4.4. Feature Selection




Figure 2. Left: We select a subset of pixel locations, then retain the corresponding pixels of all the training data for training and inference. Right: Pixel selection on dataset with gaussian noise applied to the images. Our method consistently surpasses other baselines by a large margin given different subset size constraints and the advantage is more evident with gaussian noise.

We apply our algorithm for feature selection (FS) tasks on the MNIST dataset to further validate its effectiveness. FS can be naturally viewed as a coreset construction problem, where the selected subset of features is the coreset. The

adaptation of our method in this case is straightforward: the inner loop is to train the model using the selected pixels of the images, while the outer loop is to update the probability of each pixel based on the loss on all the pixels.

Specifically, we select n pixels locations and the corresponding pixels are used during both training and inference. As shown in Figure 2, for various number of selected features, our method consistently outperforms other baselines including F-score (Güneş et al., 2010), mutual information (Sulaiman & Labadin, 2015), MRMR (Radovic et al., 2017) and greedy coreset (Borsos et al., 2020) by a large margin. The selected pixels are demonstrated in Figure 3, which shows that our method is able to retain the most informative features. To further showcase the advantage of our algorithm under more challenging scenario, we impose gaussian noise (with mean set to 0 and stand deviation 2.5) to the image pixels. As demonstrated on the right part of Figure 2, while the performance of most other methods deteriorate significantly, our method still achieves stable performance.




Figure 3. Visualization of selected pixels by our algorithm, where the numbers from 0 to 9 are demonstrated. The white pixels are selected by our algorithm and the black dots are those overlapped with the digits. We can see that the selected pixel locations effectively capture the important information of the images.

5. Ablation Study and Analysis

5.1. Advantage of Global Algorithm Compared with Greedy Counterpart

In this experiment, we analyze how the quality of coreset changes as the size increases using MNIST with 1) 90% corrupted labels and 2) imbalance factor of 200. Specifically, we monitor the change of mislabeling ratio and imbalance criterion in the selected coreset as the size increases. As shown in Figure 4, at the early stage of greedy coreset (blue line), the quality of selected data is much worse due to the lack of global knowledge. Those samples can never be removed and continues to misguide the search, which finally leads to suboptimal coreset. On the other hand, since our method (orange line) samples subsets globally with replacement during the optimization, more information can be learned in this trial-and-error process before arriving at

Table 1. Experiment result on continual learning for PermMNIST, SplitMNIST and CIFAR10 datasets. Normal stands for the standard dataset without modification; Noise represents symmetric label noise with 20% corruption ratio and Imbalance means the data has an class imbalance factor of 50. As demonstrated, using our proposed approach to construct the replay memory consistently surpasses other methods. Remarkably, the advantage of our global search strategy against the greedy counterpart becomes more prominent on challenging tasks: for tasks with label noise, our method surpasses the greedy counterpart by around 10%.

Datasets	PermMNIST			SplitMNIST			SplitCIFAR-10		
	Normal	Noise	Imbalance	Normal	Noise	Imbalance	Normal	Noise	Imbalance
Uniform	78.46	32.12	43.70	93.70	44.32	53.32	36.20	15.80	20.50
k-center embeddings	78.57	37.53	54.53	94.55	51.89	71.50	36.91	16.68	21.44
Hardest samples	76.79	15.38	38.72	91.57	17.23	56.39	28.10	9.63	15.63
iCaRL	79.68	39.36	67.53	95.13	60.99	72.23	34.52	18.47	25.47
Greedy Coreset	79.26	65.84	68.67	96.50	81.75	84.58	37.60	24.23	31.28
Ours	80.60	74.26	75.32	98.15	92.23	94.30	39.10	31.20	35.30

Table 2. Experiment result on Streaming for MNIST with and without label noise. Using our proposed approach to construct the replay memory consistently surpasses other methods significantly.

Datasets	PermMNIST		SplitMNIST	
	Normal	Noise	Normal	Noise
Reservoir sampling	73.21	25.82	90.72	22.03
Greedy Coreset	74.44	61.30	92.59	82.52
Ours	75.5	70.33	94.20	90.69




Figure 4. Analysis of how the quality of coreset changes as the coreset size increases. Left: The noise ratio of samples in the coreset. Right: The class imbalance of samples in the coreset, where the y-axis is defined as the difference between the most class and the least class, then divided by the total coreset size. We can see that the poor data quality at the early phase of greedy-based method misguides the search and results in suboptimal coreset.

the final coreset (more analysis in Section 3.3).

5.2. Evolution of Coreset during Search

In Figure 5, we analyze the search process of our method. In this experiment, we select 100 samples from 1000 training data with noise ratio set to 0.9. The left part visualizes how the distribution of probabilities evolve as the search goes on. The initial sample probabilities are evenly distributed and equal to 0.1. As the search goes on, most of the probabilities




Figure 5. Left: The distribution of probability scores during the search. As the search progresses, most of the score values converge to either 0 or 1, which eventually renders a deterministic coreset with low variance. Right: the noise ratio in the selected coreset as the outer iteration increases. Our method progressively improves the quality of the coreset by learning the global information and updating the probability distribution accordingly.

converge to either 0 or 1, i.e., the uncertainty is gradually reduced to 0, which generates a nearly deterministic sparse mask with low variance. In the right part we show how the noise ratio in the selected coreset evolves. We can observe that the noise ratio continues to decrease, which verifies that our method is able to progressively improve the quality of the coreset by learning the global information.

5.3. Time Complexity Analysis




Figure 6. Comparison of time consumption between our method and the greedy counterpart. As the greedy coreset selection method needs to solve a bilevel problem for every newly added sample, the cost increases rapidly with the coreset size.

The comparison of time complexities between our method and the greedy counterpart is shown in Figure 6. Since the greedy coreset selection method (Borsos et al., 2020) needs to solve a bilevel optimization problem for every newly

added sample, the time complexity increases rapidly with the coreset size. On the other hand, the time required by our method is not sensitive to the coreset size as the number of outer iterations remains fixed. Furthermore, owing to the efficiency of policy gradients, the update of sample probabilities takes much less time.

6. Conclusion

In this paper, we propose a global coreset selection algorithm based on bilevel optimization and adopt probabilistic reparameterization to continualize the discrete optimization problem. Our method is computationally efficient and achieves promising results even on challenging scenarios with label noise or imbalanced classes. We theoretically prove its convergence and conduct extensive experiments on various tasks to demonstrate its superiority.

Acknowledgements

This work is supported by GRF 16201320.

References

- Aljundi, R., Kelchtermans, K., and Tuytelaars, T. Task-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11254–11263, 2019a.
- Aljundi, R., Lin, M., Goujaud, B., and Bengio, Y. Gradient based sample selection for online continual learning. *arXiv preprint arXiv:1903.08671*, 2019b.
- Borsos, Z., Mutn y, M., and Krause, A. Coresets via bilevel optimization for continual learning and streaming. *arXiv preprint arXiv:2006.03875*, 2020.
- Cai, J., Luo, J., Wang, S., and Yang, S. Feature selection in machine learning: A new perspective. *Neurocomputing*, 300:70–79, 2018.
- Chrysakis, A. and Moens, M.-F. Online continual learning from imbalanced data. In *International Conference on Machine Learning*, pp. 1952–1961. PMLR, 2020.
- Cui, Y., Jia, M., Lin, T.-Y., Song, Y., and Belongie, S. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9268–9277, 2019.
- Deng, L. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- Feldman, D. and Langberg, M. A unified framework for approximating and clustering data. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 569–578, 2011.
- Feldman, D., Faulkner, M., and Krause, A. Scalable training of mixture models via coresets. In *NIPS*, pp. 2142–2150. Citeseer, 2011.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pp. 1126–1135. PMLR, 2017.
- Franceschi, L., Donini, M., Frasconi, P., and Pontil, M. Forward and reverse gradient-based hyperparameter optimization. In *International Conference on Machine Learning*, pp. 1165–1173. PMLR, 2017.
- Gao, J., Xu, H., Ren, X., Yu, P. L., Liang, X., Jiang, X., Li, Z., et al. Autobert-zero: Evolving bert backbone from scratch. *arXiv preprint arXiv:2107.07445*, 2021.
- Gao, J., Zhou, Y., Yu, P. L., Joty, S., and Gu, J. Unison: Unpaired cross-lingual image captioning. 2022.
- Ghadimi, S., Lan, G., and Zhang, H. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155 (1-2):267–305, 2016.
- Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., and Bengio, Y. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- Gou, J., Yu, B., Maybank, S. J., and Tao, D. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
- Grazzi, R., Franceschi, L., Pontil, M., and Salzo, S. On the iteration complexity of hypergradient computation. In *International Conference on Machine Learning*, pp. 3748–3758. PMLR, 2020.
- Grazzi, R., Pontil, M., and Salzo, S. Convergence properties of stochastic hypergradients. In *International Conference on Artificial Intelligence and Statistics*, pp. 3826–3834. PMLR, 2021.
- Güneş, S., Polat, K., and Yosunkaya, Ş. Multi-class f-score feature selection approach to classification of obstructive sleep apnea syndrome. *Expert systems with applications*, 37(2):998–1004, 2010.
- Har-Peled, S. and Kushal, A. Smaller coresets for k-median and k-means clustering. *Discrete & Computational Geometry*, 37(1):3–19, 2007.

- Hayes, T. L., Cahill, N. D., and Kanan, C. Memory efficient experience replay for streaming learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 9769–9776. IEEE, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hinton, G., Vinyals, O., Dean, J., et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.
- Huggins, J., Campbell, T., and Broderick, T. Coresets for scalable bayesian logistic regression. *Advances in Neural Information Processing Systems*, 29:4080–4088, 2016.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Li, J., Cheng, K., Wang, S., Morstatter, F., Trevino, R. P., Tang, J., and Liu, H. Feature selection: A data perspective. *ACM Computing Surveys (CSUR)*, 50(6):1–45, 2017.
- Liu, H., Simonyan, K., and Yang, Y. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018.
- Lopez-Paz, D. and Ranzato, M. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30:6467–6476, 2017.
- Lorraine, J., Vicol, P., and Duvenaud, D. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*, pp. 1540–1552. PMLR, 2020.
- Lucic, M., Faulkner, M., Krause, A., and Feldman, D. Training gaussian mixture models at scale via coresets. *The Journal of Machine Learning Research*, 18(1):5885–5909, 2017.
- MacKay, M., Vicol, P., Lorraine, J., Duvenaud, D., and Grosse, R. Self-tuning networks: Bilevel optimization of hyperparameters using structured best-response functions. *arXiv preprint arXiv:1903.03088*, 2019.
- Maclaurin, D., Duvenaud, D., and Adams, R. Gradient-based hyperparameter optimization through reversible learning. In *International conference on machine learning*, pp. 2113–2122. PMLR, 2015.
- Miao, J. and Niu, L. A survey on feature selection. *Procedia Computer Science*, 91:919–926, 2016.
- Nguyen, T., Chen, Z., and Lee, J. Dataset meta-learning from kernel ridge-regression. *arXiv preprint arXiv:2011.00050*, 2020.
- Nichol, A. and Schulman, J. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2(3):4, 2018.
- Pedregosa, F. Hyperparameter optimization with approximate gradient. In *International conference on machine learning*, pp. 737–746. PMLR, 2016.
- Pham, H., Guan, M., Zoph, B., Le, Q., and Dean, J. Efficient neural architecture search via parameters sharing. In *International conference on machine learning*, pp. 4095–4104. PMLR, 2018.
- Radovic, M., Ghalwash, M., Filipovic, N., and Obradovic, Z. Minimum redundancy maximum relevance feature selection approach for temporal gene expression data. *BMC bioinformatics*, 18(1):1–14, 2017.
- Rebuffi, S.-A., Kolesnikov, A., Sperl, G., and Lampert, C. H. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pp. 4334–4343. PMLR, 2018.
- Sener, O. and Savarese, S. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*, 2017.
- Shao, W., Meng, T., Li, J., Zhang, R., Li, Y., Wang, X., and Luo, P. Ssn: Learning sparse switchable normalization via sparsestmax. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Shi, H., Pi, R., Xu, H., Li, Z., Kwok, J., and Zhang, T. Bridging the gap between sample-based and one-shot neural architecture search with bonas. *Advances in Neural Information Processing Systems*, 33:1808–1819, 2020.
- Shi, H., Gao, J., Ren, X., Xu, H., Liang, X., Li, Z., and Kwok, J. T.-Y. Sparsebert: Rethinking the importance analysis in self-attention. In *International Conference on Machine Learning*, pp. 9547–9557. PMLR, 2021.
- Shu, J., Xie, Q., Yi, L., Zhao, Q., Zhou, S., Xu, Z., and Meng, D. Meta-weight-net: Learning an explicit mapping for sample weighting. *arXiv preprint arXiv:1902.07379*, 2019.

- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2015.
- Sinha, A., Malo, P., and Deb, K. A review on bilevel optimization: from classical to evolutionary approaches and applications. *IEEE Transactions on Evolutionary Computation*, 22(2):276–295, 2017.
- Sulaiman, M. A. and Labadin, J. Feature selection based on mutual information. In *2015 9th International Conference on IT in Asia (CITA)*, pp. 1–6. IEEE, 2015.
- Tsang, I. W., Kwok, J. T., Cheung, P.-M., and Cristianini, N. Core vector machines: Fast svm training on very large data sets. *Journal of Machine Learning Research*, 6(4), 2005.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- Vicol, P., Metz, L., and Sohl-Dickstein, J. Unbiased gradient estimation in unrolled computation graphs with persistent evolution strategies. In *International Conference on Machine Learning*, pp. 10553–10563. PMLR, 2021.
- Vitter, J. S. Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, 11(1): 37–57, 1985.
- Wang, H., Wu, Z., and He, J. Training fair deep neural networks by balancing influence. *arXiv preprint arXiv:2201.05759*, 2022a.
- Wang, L., Zhang, X., Yang, K., Yu, L., Li, C., HONG, L., Zhang, S., Li, Z., Zhong, Y., and Zhu, J. Memory replay with data compression for continual learning. In *International Conference on Learning Representations*, 2022b. URL <https://openreview.net/forum?id=a7H7OucbWaU>.
- Wang, T., Zhu, J.-Y., Torralba, A., and Efros, A. A. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018.
- Wang, Z., Shen, L., Duan, T., Zhan, D., Fang, L., and Gao, M. Learning to learn and remember super long multi-domain task sequence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7982–7992, 2022c.
- Yao, L., Pi, R., Xu, H., Zhang, W., Li, Z., and Zhang, T. G-detkd: Towards general distillation framework for object detectors via contrastive and semantic-guided feature imitation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3591–3600, 2021a.
- Yao, L., Pi, R., Xu, H., Zhang, W., Li, Z., and Zhang, T. Joint-detnas: Upgrade your detector with nas, pruning and dynamic distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10175–10184, 2021b.
- Zenke, F., Poole, B., and Ganguli, S. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pp. 3987–3995. PMLR, 2017.
- Zhou, X., Zhang, W., Chen, Z., Diao, S., and Zhang, T. Efficient neural network training via forward and backward propagation sparsification. *Advances in Neural Information Processing Systems*, 2021a.
- Zhou, X., Zhang, W., Xu, H., and Zhang, T. Effective sparsification of neural networks with global sparsity constraint. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3599–3608, 2021b.
- Zhou, X., Lin, Y., Pi, R., Zhang, W., Xu, R., Peng, C., and Zhang, T. Model agnostic sample reweighting for out-of-distribution learning. In *International Conference on Machine Learning*. PMLR, 2022.
- Zou, F., Shen, L., Jie, Z., Zhang, W., and Liu, W. A sufficient condition for convergences of adam and rmsprop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11127–11135, 2019.

Supplementary Materials: Probabilistic Bilevel Coreset Selection

This appendix can be divided into the following parts:

- In Section A, we provide the proof of our property to show the convergence of our method.
- In Section B, we give the algorithm for calculating the projection onto our constraint set \mathcal{C} .
- In Section C, we give the detailed configurations of our experiments.
- In Section D, we give more experimental results.
- In Section E, we present discussions on future works.

A. Proof of Property 3.3

At first, we would like to rephrase our theorem into a more formal form below by adding some assumptions following (Pedregosa, 2016).

Theorem A.1. *We assume $\Phi(\mathbf{s})$ is L -smooth, and the policy gradient variance $\mathbb{E}\|\mathcal{L}_{\mathcal{B}}(\boldsymbol{\theta}^*(\mathbf{m}))\nabla_{\mathbf{s}} \ln p(\mathbf{m}|\mathbf{s}) - \nabla_{\mathbf{s}}\Phi(\mathbf{s})\|^2 \leq \sigma^2$. Let $\eta < 1/L$ and we denote the gradient mapping \mathcal{G}^t at t -th iteration as*

$$\mathcal{G}^t = \frac{1}{\eta} (\mathbf{s}^t - \mathcal{P}_{\mathcal{C}}(\mathbf{s}^t - \eta\nabla_{\mathbf{s}}\Phi(\mathbf{s}^t))),$$

then we have

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}\|\mathcal{G}^t\|^2 \leq \frac{8 - 2L\eta}{2 - L\eta} \sigma^2, \quad (7)$$

when $T \rightarrow \infty$.

Remark A.2. We would like to point out the following things:

- We give this theorem just to show that our algorithm works well similarly with the general projected/proximal stochastic gradient descent algorithms for one level optimization problems, e.g., (Ghadimi et al., 2016), instead of to show how fast our algorithm can converge. Therefore, we do not consider the techniques, such as variance reduction, to improve the convergence rate of our algorithm, which are out of the scope of this work.
- σ^2 is actually the variance of our gradient estimator PGE with a single mask. It can become smaller with the techniques, such as sample more masks and use larger batch size in PGE. Therefore, the LHS of Eqn.(7) can converge to a small value.
- Our experimental results show that our algorithm can work well even if we sample only one mask in each iteration. Therefore, in the above theorem, we give the result when only one mask is sampled, to make it consistent with the settings of our experiments.

Before giving the detailed proof, we need the following lemmas about the properties of the projection operator, which can be found in (Ghadimi et al., 2016).

Lemma A.3 (firmly nonexpansive operators). *Given a compact convex set $\mathcal{C} \subset \mathbb{R}^d$ and let $\mathcal{P}_{\mathcal{C}}(\cdot)$ be the projection operator on \mathcal{C} , then for any $\mathbf{u} \in \mathbb{R}^d$ and $\mathbf{v} \in \mathbb{R}^d$, we have*

$$\|\mathcal{P}_{\mathcal{C}}(\mathbf{u}) - \mathcal{P}_{\mathcal{C}}(\mathbf{v})\|^2 \leq (\mathbf{u} - \mathbf{v})^\top (\mathcal{P}_{\mathcal{C}}(\mathbf{u}) - \mathcal{P}_{\mathcal{C}}(\mathbf{v})).$$

Lemma A.4. *Given a compact convex set $\mathcal{C} \subset \mathbb{R}^d$ and let $\mathcal{P}_{\mathcal{C}}(\cdot)$ be the projection operator on \mathcal{C} , then for any $\mathbf{c} \in \mathcal{C}$ and $\mathbf{u} \in \mathbb{R}^d, \mathbf{v} \in \mathbb{R}^d$, we have*

$$\|\mathcal{P}_{\mathcal{C}}(\mathbf{c} + \mathbf{u}) - \mathcal{P}_{\mathcal{C}}(\mathbf{c} + \mathbf{v})\| \leq \|\mathbf{u} - \mathbf{v}\|.$$

Proof. of Theorem A.1:

In the following, we denote

$$\mathbf{g}^t = \mathcal{L}_{\mathcal{B}}(\boldsymbol{\theta}^*(\mathbf{m})) \nabla_{\mathbf{s}} \ln p(\mathbf{m} | \mathbf{s}^t).$$

In our algorithm, we update \mathbf{s} as

$$\mathbf{s}^{t+1} = \mathcal{P}_{\mathcal{C}}(\mathbf{s}^t - \eta \mathbf{g}^t).$$

Let the stochastic and deterministic gradient mappings be

$$\begin{aligned} \hat{\mathcal{G}}^t &= \frac{1}{\eta} (\mathbf{s}^t - \mathcal{P}_{\mathcal{C}}(\mathbf{s}^t - \eta \mathbf{g}^t)) = \frac{1}{\eta} (\mathbf{s}^t - \mathbf{s}^{t+1}), \\ \mathcal{G}^t &= \frac{1}{\eta} (\mathbf{s}^t - \mathcal{P}_{\mathcal{C}}(\mathbf{s}^t - \eta \nabla \Phi(\mathbf{s}^t))), \end{aligned}$$

we can have

$$\begin{aligned} \Phi(\mathbf{s}^{t+1}) &\leq \Phi(\mathbf{s}^t) + \langle \nabla \Phi(\mathbf{s}^t), \mathbf{s}^{t+1} - \mathbf{s}^t \rangle + \frac{L}{2} \|\mathbf{s}^{t+1} - \mathbf{s}^t\|^2 \\ &= \Phi(\mathbf{s}^t) - \eta \langle \nabla \Phi(\mathbf{s}^t), \hat{\mathcal{G}}^t \rangle + \frac{L\eta^2}{2} \|\hat{\mathcal{G}}^t\|^2 \\ &= \Phi(\mathbf{s}^t) - \eta \langle \nabla \Phi(\mathbf{s}^t) - \mathbf{g}^t + \mathbf{g}^t, \hat{\mathcal{G}}^t \rangle + \frac{L\eta^2}{2} \|\hat{\mathcal{G}}^t\|^2 \\ &= \Phi(\mathbf{s}^t) - \eta \langle \mathbf{g}^t, \hat{\mathcal{G}}^t \rangle + \frac{L\eta^2}{2} \|\hat{\mathcal{G}}^t\|^2 + \eta \langle \delta^t, \hat{\mathcal{G}}^t \rangle \text{ (here } \delta^t = \mathbf{g}^t - \nabla \Phi(\mathbf{s}^t)) \\ &\leq \Phi(\mathbf{s}^t) - \eta \|\hat{\mathcal{G}}^t\|^2 + \frac{L\eta^2}{2} \|\hat{\mathcal{G}}^t\|^2 + \eta \langle \delta^t, \hat{\mathcal{G}}^t \rangle \text{ (Lemma A.3)} \\ &\leq \Phi(\mathbf{s}^t) - (\eta - \frac{L\eta^2}{2}) \|\hat{\mathcal{G}}^t\|^2 + \eta \langle \delta^t, \hat{\mathcal{G}}^t \rangle \\ &= \Phi(\mathbf{s}^t) - (\eta - \frac{L\eta^2}{2}) \|\hat{\mathcal{G}}^t\|^2 + \eta \langle \delta^t, \mathcal{G}^t \rangle + \eta \langle \delta^t, \hat{\mathcal{G}}^t - \mathcal{G}^t \rangle \\ &\leq \Phi(\mathbf{s}^t) - (\eta - \frac{L\eta^2}{2}) \|\hat{\mathcal{G}}^t\|^2 + \eta \langle \delta^t, \mathcal{G}^t \rangle + \eta \|\delta^t\| \|\hat{\mathcal{G}}^t - \mathcal{G}^t\| \\ &\leq \Phi(\mathbf{s}^t) - (\eta - \frac{L\eta^2}{2}) \|\hat{\mathcal{G}}^t\|^2 + \eta \langle \delta^t, \mathcal{G}^t \rangle + \eta \|\delta^t\|^2. \text{ (Lemma A.4)} \end{aligned}$$

Therefore, we can get

$$(\eta - \frac{L\eta^2}{2}) \|\hat{\mathcal{G}}^t\|^2 \leq \Phi(\mathbf{s}^t) - \Phi(\mathbf{s}^{t+1}) + \eta \langle \delta^t, \mathcal{G}^t \rangle + \eta \|\delta^t\|^2.$$

Thus, we can obtain

$$\sum_{t=1}^T (\eta - \frac{L\eta^2}{2}) \|\hat{\mathcal{G}}^t\|^2 \leq \Phi(\mathbf{s}^1) - \Phi(\mathbf{s}^{T+1}) + \eta \sum_{t=1}^T (\langle \delta^t, \mathcal{G}^t \rangle + \|\delta^t\|^2). \quad (8)$$

Now, we turn to analyze the item $\langle \delta^t, \mathcal{G}^t \rangle$ as follows:

$$\mathbb{E} \langle \delta^t, \mathcal{G}^t \rangle = \mathbb{E}_{\mathbf{s}^t} \mathbb{E}_{|\mathbf{s}^t} (\langle \mathbf{g}^t - \nabla \Phi(\mathbf{s}^t), \mathcal{G}^t \rangle | \mathbf{s}^t) = 0, \quad (9)$$

For $\|\delta^t\|^2$, we have

$$\mathbb{E}\|\delta^t\|^2 = \mathbb{E}\|\mathbf{g}^t - \nabla\Phi(\mathbf{s}^t)\|^2 \leq \sigma^2. \quad (10)$$

Combining the inequalities (9), (8) and (10), we can have

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}\|\hat{\mathcal{G}}^t\|^2 \leq \frac{\Phi(\mathbf{s}^1) - \Phi^*}{(1 - L\eta/2)T} + \frac{\sigma^2}{1 - L\eta/2}. \quad (11)$$

Finally, we bound $\mathbb{E}\|\mathcal{G}^t\|^2$ as follows:

$$\mathbb{E}\|\mathcal{G}^t\|^2 \leq 2\mathbb{E}\|\hat{\mathcal{G}}^t\|^2 + 2\mathbb{E}\|\mathbf{g}^t - \nabla\Phi(\mathbf{s}^t)\|^2 \quad (12)$$

$$\leq 2\mathbb{E}\|\hat{\mathcal{G}}^t\|^2 + 2\sigma^2. \quad (13)$$

Combine inequalities (13) and (11), when $T \rightarrow \infty$, we can obtain

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}\|\mathcal{G}^t\|^2 \leq \frac{2}{1 - L\eta/2} \left(\frac{\Phi(\mathbf{s}^1) - \Phi^*}{T} + (2 - L\eta/2)\sigma^2 \right) \rightarrow \frac{8 - 2L\eta}{2 - L\eta} \sigma^2$$

□

B. Project Calculation

The projection from \mathbf{s} to \mathcal{C} can be calculated by:

Algorithm 2 Projection from \mathbf{z} to \mathcal{C}

Require: a vector \mathbf{z} .

1: Solve v_1 from $\mathbf{1}^\top [\min(1, \max(0, \mathbf{z} - v_1^* \mathbf{1}))] - K = 0$.

2: $v_2^* \leftarrow \max(0, v_1^*)$.

3: $\mathbf{s} \leftarrow \min(1, \max(0, \mathbf{z} - v_2^* \mathbf{1}))$.

output \mathbf{s}

Proof. The projection from \mathbf{z} to set \mathcal{C} can be formulated in the following optimization problem:

$$\begin{aligned} & \min_{\mathbf{s} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{s} - \mathbf{z}\|^2, \\ & s.t. \mathbf{1}^\top \mathbf{s} \leq K \text{ and } 0 \leq s_i \leq 1. \end{aligned}$$

Then we solve the problem with Lagrangian multiplier method.

$$L(\mathbf{s}, v) = \frac{1}{2} \|\mathbf{s} - \mathbf{z}\|^2 + v(\mathbf{1}^\top \mathbf{s} - K) \quad (14)$$

$$= \frac{1}{2} \|\mathbf{s} - (\mathbf{z} - v\mathbf{1})\|^2 + v(\mathbf{1}^\top \mathbf{z} - K) - \frac{n}{2} v^2. \quad (15)$$

with $v \geq 0$ and $0 \leq s_i \leq 1$. Minimize the problem with respect to \mathbf{s} , we have

$$\tilde{\mathbf{s}} = \mathbf{1}_{\mathbf{z} - v\mathbf{1} \geq \mathbf{1}} + (\mathbf{z} - v\mathbf{1})_{\mathbf{1} > \mathbf{z} - v\mathbf{1} > \mathbf{0}} \quad (16)$$

Then we have

$$\begin{aligned}
 g(v) &= L(\tilde{\mathbf{s}}, v) \\
 &= \frac{1}{2} \|[z - v\mathbf{1}]_- + [z - (v + 1)\mathbf{1}]_+\|^2 \\
 &\quad + v(\mathbf{1}^\top z - s) - \frac{n}{2}v^2 \\
 &= \frac{1}{2} \|[z - v\mathbf{1}]_-\|^2 + \frac{1}{2} \|[z - (v + 1)\mathbf{1}]_+\|^2 \\
 &\quad + v(\mathbf{1}^\top z - s) - \frac{n}{2}v^2, v \geq 0. \\
 g'(v) &= \mathbf{1}^\top [v\mathbf{1} - z]_+ + \mathbf{1}^\top [(v + 1)\mathbf{1} - z]_- \\
 &\quad + (\mathbf{1}^\top z - s) - nv \\
 &= \mathbf{1}^\top \min(1, \max(0, z - v\mathbf{1})) - K, v \geq 0.
 \end{aligned}$$

It is easy to verify that $g'(v)$ is a monotone decreasing function with respect to v and we can use a bisection method solve the equation $g'(v) = 0$ with solution v_1^* . Then we get that $g(v)$ increases in the range of $(-\infty, v_1^*]$ and decreases in the range of $[v_1^*, +\infty)$. The maximum of $g(v)$ is achieved at 0 if $v_1^* \leq 0$ and v_1^* if $v_1^* > 0$. Then we set $v_2^* = \max(0, v_1^*)$. Finally we have

$$\mathbf{s}^* = \mathbf{1}_{z - v_2^* \mathbf{1} \geq 1} + (z - v_2^* \mathbf{1})_{1 > z - v_2^* \mathbf{1} > 0} \quad (17)$$

$$= \min(1, \max(0, z - v_2^* \mathbf{1})). \quad (18)$$

□

C. Experiment Details

We use the following hyper-parameters during optimization for our experiments. For the inner-loop, the model is trained for 100 epochs using SGD with learning rate of 0.1 and momentum of 0.9. For the outer-loop, the probabilities s are optimized by adam with learning rate of 2.5 and cosine scheduler. The outer-loop is updated for 500-2000 times. Note that gumbel softmax can also be used as an alternative for PGE when calculating the hyper-gradient, which sometimes demonstrates better stability. In implementation, we combine gumbel softmax and PGE to achieve a balance between efficiency and accuracy.

Label Noise and Class Imbalance For pairwise noise, the label of a particular class has a probability p to be flipped to the adjacent class; for symmetric noise, a class has a probability $\frac{p}{n-1}$ to be changed to any other $n - 1$ classes. For the class imbalance experiment, we adopt similar setting as in (Cui et al., 2019). Specifically, the number of training samples per class is exponentially reduced according to the function $n'_i = n_i \sigma^i$, where i is the class index. We define dataset imbalance factor as $\frac{n_{max}}{n_{min}}$, where n_{max} and n_{min} are the number of samples in the largest and smallest classes, respectively.

Data Summarization We follow (Borsos et al., 2020) and use a convolutional neural network stacked with two blocks of convolution, dropout, max-pooling and ReLU activation for MNIST, and use ResNet18(He et al., 2016) for the CIFAR10 experiment. For baselines K-center clustering using the embedding from last layer (Sener & Savarese, 2017), iCaRL’s selection (Rebuffi et al., 2017) and Hardest sampling (Aljundi et al., 2019a), which depend on the model embedding, we pretrain a feature extractor using 1000 uniformly sampled data.

Continual learning To make a fair comparison, we follow (Borsos et al., 2020) to keep 1000 subsamples for each task other than SplitCIFAR100 (where all data is used). For PermMNIST, we use a fully-connected network with two hidden layers of 100 neurons, followed by ReLU activation and dropout with probability 0.2, the memory size is set to 100. For SplitMNIST, we use the same CNN architecture as in data summarization task, the memory size is 500. For SplitCIFAR10, we adopt ResNet18 similar to Section 4.1, the memory is set to 200. As for SplitCIFAR100, all the training data is used and the memory size is set to 2000. For experiments with label noise and class imbalance, a balanced held-out validation dataset of size 100 is used.

Streaming We follow (Borsos et al., 2020) to modify PermMNIST and SplitMNIST used in continual learning by concatenating all the tasks and stream the data in subsets of 125. The memory size in this setting is set to 100 and separated into 10 slots. Merge reduced introduced in (Borsos et al., 2020) is also adopted for fair comparison.

In feature selection experiments, we uniformly sample 1000 training data, while all the testing data are used to evaluate the trained model. All the experiments are repeated with 5 different random seeds.

D. More Experiments

D.1. Comparison with Training on Entire Dataset under Label Noise and Class Imbalance Settings




Figure 7. We compare the performance of the trained model on the coreset selected by our method with that trained on the entire dataset, where the dataset contains label noise and class imbalance. We can see that the training with the selected coreset surpasses the entire dataset by a large margin, which is because the coreset can effectively remove the label noise and automatically balance the data in different classes.

As mentioned in Section 1 of the main paper, training on the coreset may sometimes even achieve better performance than training on the entire dataset. We verify this via conducting experiment on dataset with label noise and class imbalance. We can see that training the model using the entire dataset in these cases lead to failure due to the poor data quality. On the other hand, our coreset selection method has the effect of removing the label noise and automatically balancing the data in each class.

D.2. Transferability of found coresets with various sizes among different networks

Table 3. Transferability of found coresets with various sizes among different networks.

Dataset / Size	Net1		Net2	
	Random	Coreset	Random	Coreset
MNIST (100)	87.3	95.01	70.90	75.60
CIFAR10 (4000)	63.50	78.35	67.57	79.05

We use the coreset selected by one network to train different networks. In Table 3, we show the results of the following experiment: The coresets is searched by Net1, which is then used to train both Net1 itself and Net2. The found coreset outperforms uniform sampling for both networks, which verifies the transferability of the found coreset. For CIFAR10, Net1 is ResNet18 and Net2 is ResNet32; for MNIST, Net1 is Convnet and Net2 is MLP.

E. Future Directions

Our coreset selection also performs rather faster than Borsos et al. (2020), it stills faces computational difficulties when applied to larger datasets like ImageNet-1K. One possible solution is to employ model sparsity to speed-up the training process or make the trained model smaller (Shao et al., 2019; Zhou et al., 2021a;b).