

Deep Active Learning with Noise Stability

Xingjian Li^{1*}, Pengkun Yang^{2*}, Yangcheng Gu³, Xueying Zhan¹, Tianyang Wang⁴,
Min Xu^{1†}, Chengzhong Xu⁵

¹Computational Biology Department, Carnegie Mellon University

²Center for Statistical Science, Tsinghua University

³School of Software, Tsinghua University

⁴Department of Computer Science, University of Alabama at Birmingham

⁵State Key Lab of IOTSC, University of Macau

lixj04@gmail.com, yangpengkun@tsinghua.edu.cn, guyc19@mails.tsinghua.edu.cn

sinezhan17@gmail.com, tw2@uab.edu, mxu1@cs.cmu.edu, czxu@um.edu.mo

Abstract

Uncertainty estimation for unlabeled data is crucial to active learning. With a deep neural network employed as the backbone model, the data selection process is highly challenging due to the potential over-confidence of the model inference. Existing methods resort to special learning fashions (e.g. adversarial) or auxiliary models to address this challenge. This tends to result in complex and inefficient pipelines, which would render the methods impractical. In this work, we propose a novel algorithm that leverages noise stability to estimate data uncertainty. The key idea is to measure the output derivation from the original observation when the model parameters are randomly perturbed by noise. We provide theoretical analyses by leveraging the small Gaussian noise theory and demonstrate that our method favors a subset with large and diverse gradients. Our method is generally applicable in various tasks, including computer vision, natural language processing, and structural data analysis. It achieves competitive performance compared against state-of-the-art active learning baselines.

Introduction

The success of training a deep neural network highly depends on a huge amount of labeled data. However, much data is unlabeled in the era of big data, and data annotation cost could be extremely high, such as in medical area (Gorritz et al. 2017; Konyushkova, Sznitman, and Fua 2017). Due to limited annotation budget, it might be more feasible to annotate a portion of the given data. Then a question can be naturally raised: which data is deserved for being annotated?

Active learning was proposed to solve this question. One of the main ideas is to select the most uncertain/informative data from an unlabeled pool for annotation. The newly annotated data is then used to train a task model in supervised fashion. Nevertheless, uncertainty estimation with deep neural networks (DNNs) remains challenging, due to the potential over-confidence of deep models (Ren et al. 2021). That is, considering classification problems, the primitive softmax

output (as a score) does not necessarily reflect a reliable uncertainty or confidence of the prediction. Several methods have been proposed to address this challenge. For example, deep Bayesian methods (Gal and Ghahramani 2016) provide a more accurate estimate of uncertainty (e.g. in terms of entropy or mutual information) with the posterior predictive probability over the model parameters. Many other methods propose to employ auxiliary models for a better estimate, such as query-by-committee (QBC) (Freund et al. 1997a; Gorritz et al. 2017), Variational Auto-Encoder (Sinha, Ebrahimi, and Darrell 2019), adversarial learning (Ducoffe and Precioso 2018; Mayer and Timofte 2020) and Graph Convolutional Network (Caramalau, Bhattarai, and Kim 2021).

Despite the advances in uncertainty estimation, they rely on additional models or ad-hoc learning fashions, highly limiting their applications. Furthermore, many of them fail in considering the diversity of the selected data, throwing a risk of selecting redundant or similar examples in practical scenarios of batch acquisition. Motivated to address these issues, our work aims to design a more general and effective active learning algorithm.

In this work, we study the problem of deep active learning from a different perspective named *noise stability*. Given a well-trained model and a data example, we are interested in how the model behavior changes when a slight perturbation is imposed on the model parameters. Specifically, we observe the deviation on the deep features of DNNs. A data example which is sensitive to such perturbations, is regarded as uncertain under the current model. The sensitivity can be measured by the magnitude of the feature deviation. Furthermore, the deviation direction implies common and different patterns among examples, which provides meaningful information to encourage diversity. Imagine the ideal case that, when the network perturbation causes the ablation of recognizing a specific pattern, examples that highly depend on this pattern will have more changes than others on corresponding elements of the feature vector¹.

*These authors contributed equally.

†Corresponding author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹In real cases, a random perturbation leads to a complex change in model parameters, e.g. some extracted patterns are affected more than others, which consequently results in much denser deviation directions among examples.

In light of the aforementioned intuitions, we design a novel deep active learning algorithm by measuring the feature deviation from the original observation, when imposing a small noise on the model parameters. The magnitude of the deviation acts as a metric of uncertainty, and the direction is utilized to constitute a diverse subset of selected unlabeled examples. We provide a theoretical analysis under the condition that the noise magnitude is small. By imposing standard multivariate Gaussian noise to the parameters, we prove that the induced deviation yields a sort of *projected gradient*. The expectation of the magnitude of this gradient is equivalent to the gradient norm w.r.t the model parameters, whereas the direction of this gradient serves as a proxy of the direction of the original gradient.

Our method is easy to implement and free of data-dependent designs or auxiliary models. Therefore, it can be exploited in classification, regression and more complex tasks such as semantic segmentation. It is also applicable to various data domains such as image, natural language and structural data. We evaluate our method on multiple tasks, including (1) image classification on MNIST (LeCun et al. 1998), Cifar10 (Krizhevsky, Hinton et al. 2009), SVHN (Netzer et al. 2011), Cifar100 (Krizhevsky, Hinton et al. 2009), and Caltech101 (Fei-Fei, Fergus, and Perona 2006); (2) regression on the Ames Housing dataset (De Cock 2011); (3) semantic segmentation on Cityscapes (Cordts et al. 2016); and (4) natural language processing on MRPC (Dolan and Brockett 2005). Our method is shown to exceed or be comparable to the state-of-the-art baselines on all these tasks.

The contributions of this work are summarized as follows.

1. We propose a novel deep active learning method, namely *NoiseStability*, to select unlabeled data for annotation. It is free of any auxiliary models or customized training fashions. We conduct extensive experiments on diverse tasks and data types to demonstrate the general applicability and superior performance of our method.
2. We prove that selecting unlabeled data w.r.t. noise stability is theoretically equivalent to selecting data according to randomly projected gradients. Conditioned on standard Gaussian noise, the induced feature deviation exhibits decent properties in both estimating uncertainty and ensuring diversity.
3. Connections of our method with predictive variance reduction and gradient-based active learning methods are discussed in detail.

Related Work

Active Learning

Active learning, aiming at determining which data to be annotated, has been a long-term open research topic in the community. Among existing literature, two most important criterion for data selection are uncertainty and diversity.

Uncertainty Estimation. Traditional active learning methods take the data, which leads to a *low predictive confidence*, as most informative for further learning. Typical strategies select data of which predictions are close to the decision boundary (Tong and Koller 2001; Dasgupta and Hsu 2008;

Žliobaitė et al. 2013) or inconsistent among different hypotheses (known as version-space based approach) (Cohn, Atlas, and Ladner 1994; Freund et al. 1997b). On over-confident DNNs, to tackle the challenge of incredible probabilistic predictions, these ideas are implemented through deep Bayesian approximation with MC-Dropout (Gal and Ghahramani 2016), query-by-committee (QBC) with multiple-model training (Gorritz et al. 2017) or adversarial training (Ducoffe and Precioso 2018; Mayer and Timofte 2020).

Some other approaches aim to choose spaces which have a *low data density*. For example, (Sinha, Ebrahimi, and Darrell 2019; Zhang et al. 2020) adopt Variational Autoencoder (VAE) to model data distributions, and then select the data that less likely lies in the labeled set in an adversarial fashion. (Caramalau, Bhattarai, and Kim 2021) uses Graph Convolutional Networks (GCN) to characterize the relationship among examples and distinguish those unlabeled examples which are sufficiently different from labelled ones.

Another typical idea focuses on possible model changes caused by adding candidate examples, and favors those *highly influential data*. For example, (Cohn, Ghahramani, and Jordan 1996) introduces the goal of selecting a new example that is expected to minimize the predictive variance on existing labeled examples. (Yoo and Kweon 2019) incorporates a light-weight module to learn the predictive error of unlabeled examples, namely Learning Loss for Active Learning (LL4AL). (Liu et al. 2021; Wang et al. 2022) instead suggest leveraging the influence function (Koh and Liang 2017) to estimate the potential model change.

For differentiable functions such as DNNs, the Expectation Gradient Length (EGL) has been shown an effective criteria to measure the influence (to the model) of a given example (Settles and Craven 2008), which motivates active learning applications in specific domains, such as speech recognition (Huang et al. 2016) and text classification (Zhang, Lease, and Wallace 2017). However, since labels are not available in active learning, these methods have to resort to pseudo labels or a marginalization over the label space.

Diversification. In a more realistic setting where a batch of examples are selected at one time, active learning algorithms should avoid redundancy among selected examples. To tackle this problem, distribution-based methods (Nguyen and Smeulders 2004; Yang et al. 2015; Elhamifar et al. 2013; Hasan and Roy-Chowdhury 2015) have been designed. Inspired by the theory of set cover, a recent state-of-the-art method Coreset (Sener and Savarese 2018) proposes to select examples with diverse representations. (Loquercio, Segu, and Scaramuzza 2020) further investigates the relevance between data diversity and uncertainty.

Moreover, several specific uncertainty-based approaches successfully take diversity into consideration. VAAL (Sinha, Ebrahimi, and Darrell 2019) selects diverse examples by choosing data from different regions of VAE’s latent space. CoreGCN (Caramalau, Bhattarai, and Kim 2021) proposes to select diverse data through graph embeddings. Cluster-Margin (Citovsky et al. 2021) leverages Hierarchical Agglomerative Clustering (HAC) to diversify a batch of most uncertain examples.

For gradient-based methods, enhancing diversity is a non-

trivial objective. This is because when comparing gradient directions, one has to deal with the extremely huge dimension of the gradients (w.r.t. model parameters). A recent state-of-the-art method BADGE (Ash et al. 2020) proposes a practical solution by using the parameters of the last neural layer.

Posterior Approximation for Bayesian Neural Network

The Bayesian Neural Network (BNN) aims to quantify the uncertainty by considering the probabilistic distribution of parameters, rather than deterministic ones. In terms of predictive variance reduction, our work is relevant to the linearised Laplace approximation of the BNN posterior (Immer et al. 2021; Daxberger et al. 2021) and BAIT (Ash et al. 2021) which involves the negative Hessian matrix.

Despite their theoretical relevance, the resulting algorithms are different: (Immer et al. 2021) is based on the determinant of the Hessian matrix; the method in (Daxberger et al. 2021) involves the diagonal of inverse Hessian. Our method is more efficient and scalable, where choosing examples with larger and diverse projected gradients serves as the cornerstone of our approach.

Noise Stability for Supervised Learning

Noise stability, particularly in terms of the output stability w.r.t. the input noise, has been studied in standard supervised learning. For example, (Bishop 1995) proposes to train a model with perturbed input examples as a regularizer to improve the model performance. (Arora et al. 2018) shows that, the stability of each layer’s computation towards input noise, injected at lower layers, acts as a good indicator of the generalization bounds for DNNs. An intuition behind these achievements is that, a model (more exactly, its parameters) robust to perturbed input examples should perform better in recognizing unseen examples (e.g. test data).

Inspired by the aforementioned intuition, in this work we explore how the robustness of model parameters impact data selection in active learning. In fact, this problem can be regarded as an inverse problem of the classical noise stability. In terms of the objective, our method pursues informative data which eventually leads to an enhanced model generalization.

Deep Active Learning with Noise Stability

Problem Definition

Here we formulate the deep active learning problem. Given a pool of unlabeled data \mathcal{U} , and a labeled pool \mathcal{L} which is initially empty. Active learning aims to select a portion of data from \mathcal{U} depending on an annotation budget (e.g. select 2500 samples out of 50000 at a time). The selected data $\{X_N\}$ is then annotated by human oracles (or equivalent), and added to the labeled pool. That is, $\mathcal{L} \leftarrow \mathcal{L} + \{X_N, Y_N\}$, where $\{Y_N\}$ is the new annotation for $\{X_N\}$. The unlabeled pool is then updated by removing the selected data: $\mathcal{U} \leftarrow \mathcal{U} - \{X_N\}$. The updated \mathcal{L} is then used to train a task model composed of a feature extractor f followed by a task head g . f and g are parameterized by θ and ϕ , respectively. We use $f(\cdot; \theta)$ and $g(\cdot; \phi)$ to denote the corresponding feed-forward operations

Algorithm 1: Active learning with Noise Stability.

Input : \mathcal{T} : random initialized neural network, \mathcal{U} : unlabeled pool of training data, \mathcal{L} : labeled pool of training data;
Output : \mathcal{L} : updated labeled pool;

```

1 begin
2   train  $\mathcal{T} = g \circ f$  with  $\mathcal{L}$ , obtaining the current
   parameters  $\theta, \phi$ ;
3   for  $k \leftarrow 1$  to  $K$  do
4     sample a random direction  $\mathbf{u}^{(k)}$ ;
5     get  $\Delta\theta^{(k)} = \zeta \|\theta\|_2 \mathbf{u}^{(k)}$ ;
6   for every  $\mathbf{x}$  in  $\mathcal{U}$  do
7     for  $k \leftarrow 1$  to  $K$  do
8        $\Delta z(\mathbf{x})^{(k)} = f(\mathbf{x}; \theta + \Delta\theta^{(k)}) - f(\mathbf{x}; \theta)$ ;
9     get  $\widehat{\Delta z}(\mathbf{x})$  by concatenating all  $\Delta z(\mathbf{x})^{(k)}$  for
        $k = 1, 2, \dots, K$ ;
10  compute  $N$  centers  $\{c_j\}_{j=1}^N$  w.r.t.
     $\{\widehat{\Delta z}(\mathbf{x}) | \mathbf{x} \in \mathcal{U}\}$  by greedy k-center algorithm;
11  select samples whose indexes are  $\{c_j\}_{j=1}^N$  in  $\mathcal{U}$  as
     $\{X_N\}$ , and obtain their labels  $\{Y_N\}$ ;
12  update  $\mathcal{L}$  with  $\mathcal{L} = \mathcal{L} \cup \{X_N, Y_N\}$ ;
13  return  $\mathcal{L}$ ;

```

of f and g . Note that the feature extractor f is normally a deep neural network of which parameters $\theta \in \mathcal{R}^n$ are in high dimension. In contrast, the architecture of g is arguably shallow, e.g. a single fully connected layer for classification or regression problems.

Uncertainty Estimation with Noise Stability

We aim at investigating how intermediate representations change when a small perturbation is imposed on model parameters. Specifically, given \mathbf{x} as an input, we consider the output of the feature extractor, i.e. $\mathbf{z} = f(\mathbf{x}; \theta) \in \mathcal{R}^d$. Let $\Delta\theta \in \mathcal{R}^n$ denote the added noise (perturbation) and $\Delta\theta = \zeta \|\theta\|_2 \mathbf{u}$, where ζ controls the relative magnitude of the noise and \mathbf{u} is a random direction with the unit length.² ζ should be relatively small (e.g. 10^{-3}) in order to avoid a catastrophic perturbation to the clean model. Then, for an unlabeled sample $\mathbf{x} \in X_U$, its feature deviation is

$$\Delta z(\mathbf{x}) = \frac{1}{\zeta \|\theta\|_2} (f(\mathbf{x}; \theta + \Delta\theta) - f(\mathbf{x}; \theta)). \quad (1)$$

For brevity, we assume \mathbf{z} and $\Delta\mathbf{z}$ are both row vectors, and use $\Delta\theta^{(k)}$ to denote the k -th sampled parameter noise. We perform K samplings of $\Delta\theta$ and concatenate all the resulting feature deviations as

$$\widehat{\Delta z}(\mathbf{x}) = \sqrt{\frac{n}{K}} (\Delta z(\mathbf{x})^{(1)}, \Delta z(\mathbf{x})^{(2)}, \dots, \Delta z(\mathbf{x})^{(K)}), \quad (2)$$

where $\widehat{\Delta z}(\mathbf{x}) \in \mathcal{R}^{Kd}$ is the empirical result and each $\Delta z(\mathbf{x})^{(k)}$ corresponds to the feature deviation produced by

²In practice, \mathbf{u} can be sampled from a standard multivariate Gaussian distribution and then normalized by its own L^2 -norm.

the parameter noise $\Delta\theta^{(k)}$.

Here we adopt the empirical feature deviation $\widehat{\Delta z}(\mathbf{x})$ as the representation of each unlabeled example \mathbf{x} . The rationale lies in two aspects. Firstly, in active learning, we favor data with higher uncertainty w.r.t. the current learned model. Such data are usually hard examples that tend to be sensitive to small parameter perturbations. Therefore, the length of $\widehat{\Delta z}(\mathbf{x})$ can be used as a criterion for uncertainty estimation. Secondly, the direction of $\widehat{\Delta z}(\mathbf{x})$ implies potential characteristics of an example \mathbf{x} , that is, which learned feature components are sensitive (or not robust) to a parameter perturbation. Diverse currently non-robust features should be focused on, in order to learn more comprehensive knowledge. Due to its relevance to data diversity, the direction of $\widehat{\Delta z}(\mathbf{x})$ can be utilized to select a batch of samples, where the diversity of selected examples is also essential to a good performance (Sener and Savarese 2018; Kirsch, Van Amersfoort, and Gal 2019).

Considering both magnitude and direction, we use Gonzalez’s greedy implementation (Gonzalez 1985) of the k-center algorithm for diverse subset selection. As pointed out in BADGE (Ash et al. 2020), a diverse subset will naturally prefer examples with larger magnitudes as they tend to be far away from each other if measured in a Euclidean space. The pipeline of our method is summarized in Algorithm 1.

Theoretical Understandings

In this section, we present theoretical understandings of our method for deep active learning. Our main conclusions are summarized as follows.

- When the noise magnitude $\zeta\|\theta\|$ is sufficiently small, for a given example \mathbf{x} , the expectation of $\|\widehat{\Delta z}(\mathbf{x})\|_2$ is the Frobenius norm of the Jacobian of f w.r.t. its parameters θ , when proper prerequisites are satisfied for \mathbf{u} .
- Considering the 1-d case $f : X \rightarrow \mathcal{R}$, the suggested $\widehat{\Delta z}(\mathbf{x})$ can be regarded as a projected gradient, which has a much lower dimension compared to the original gradient $\nabla_{\theta} f$. With the projected gradient being employed, the spatial relationships (e.g. Euclidean distance and inner product) among examples remain the same as if the original gradient is being used, indicating $\widehat{\Delta z}(\mathbf{x})$ a reasonable proxy of the original gradient, when considering diversity.

Noise Stability as Jacobian Norm

Let $f(\mathbf{x}; \theta)$ be differentiable (w.r.t. θ) at $\theta \in \mathcal{R}^n$ given \mathbf{x} as an input. When the imposed noise $\Delta\theta = \zeta\|\theta\|_2\mathbf{u}$ has a sufficiently small magnitude, i.e. in the limit as $\zeta \rightarrow 0$, we use the first-order Taylor expansion to represent the perturbed output $f(\mathbf{x}; \theta + \Delta\theta)$ as

$$f(\mathbf{x}; \theta + \Delta\theta) = f(\mathbf{x}; \theta) + J_{\theta}(\mathbf{x}; \theta)\Delta\theta + o(\zeta\|\theta\|_2), \quad (3)$$

where $J_{\theta}(\mathbf{x}; \theta)$ denotes the Jacobian matrix of f w.r.t the parameters θ as $J_{\theta}(\mathbf{x}; \theta)_{(i,j)} = \partial f(\mathbf{x}; \theta)_{(i)} / \partial \theta_{(j)}$. By applying the Taylor approximation, each block $\Delta z(\mathbf{x})^{(k)}$ can be formulated as

$$\Delta z(\mathbf{x})^{(k)} = J_{\theta}(\mathbf{x}; \theta)\mathbf{u}^{(k)} + o(1). \quad (4)$$

By omitting the higher order term $o(1)$, it results in the output deviation magnitude to be

$$\|\widehat{\Delta z}(\mathbf{x})\|_2^2 = \frac{n}{K} \sum_{k=1}^K \|J_{\theta}(\mathbf{x}; \theta)\mathbf{u}^{(k)}\|_2^2. \quad (5)$$

Conditioned on the elements of $\mathbf{u}^{(k)}$ are zero-mean and independent of each other, it can be derived that the expectation of the data selection metric (i.e. $\|\widehat{\Delta z}(\mathbf{x})\|_2$) is equivalent to the Frobenius norm of Jacobian as

$$\mathbb{E}_{\mathbf{u}} \|\widehat{\Delta z}(\mathbf{x})\|_2^2 = \|J_{\theta}(\mathbf{x}; \theta)\|_F^2. \quad (6)$$

We put a detailed proof of Eq (6) in Appendix A.1 using the Standard multivariate Gaussian distribution as a demonstration. In fact, Eq (6) is satisfied for any appropriate noise distributions with the aforementioned prerequisite of independence and zero mean. We also provide a quantified concentration inequality, showing that $\|\widehat{\Delta z}(\mathbf{x})\|_2^2$ can be close to the Jacobian norm with high probability in Appendix A.1.

Diversified Projected Gradients

Since $\widehat{\Delta z}(\mathbf{x})$ indicates which directions (i.e. elements of the final feature) of an example is more sensitive to a specific parameter perturbation, it can be employed to select diverse examples to constitute a subset. By applying the first order Taylor expansion, each $\widehat{\Delta z}(\mathbf{x})$ can be regarded as a *projected gradient* when the output of f is a scalar (here we ignore the Jacobian matrix for simplicity). That is, the original gradient $\nabla_{\theta} f$ is transformed by random directions $\{\mathbf{u}^{(k)}\}_{k=1}^K$ to its projected format as shown in Eq (4).

The benefit of diverse gradients has been discussed in existing literature. For example, a recent work (Zhang, Kjellström, and Mandt 2017) theoretically shows that, performing SGD with diversified gradients within a mini-batch helps lower the variance of stochastic gradients. BADGE (Ash et al. 2020) selects diverse examples according to their last layer’s gradients rather than considering that of the entire model. Here we show that, leveraging the expectation of the projected gradients $\{\widehat{\Delta z}(\mathbf{x}) | \mathbf{x} \in \mathcal{U}\}$ is equivalent to using $\{\nabla_{\theta} f(\mathbf{x}; \theta) | \mathbf{x} \in \mathcal{U}\}$. Specifically, given two examples \mathbf{x}_i and \mathbf{x}_j , the following two properties hold.

Property 1. Equivalence of Euclidean distance.

$$\mathbb{E}_{\mathbf{u}} \|\widehat{\Delta z}(\mathbf{x}_i) - \widehat{\Delta z}(\mathbf{x}_j)\|_2^2 = \|\nabla_{\theta} f(\mathbf{x}_i; \theta) - \nabla_{\theta} f(\mathbf{x}_j; \theta)\|_2^2. \quad (7)$$

Property 2. Equivalence of inner product.

$$\mathbb{E}_{\mathbf{u}} \langle \widehat{\Delta z}(\mathbf{x}_i), \widehat{\Delta z}(\mathbf{x}_j) \rangle = \langle \nabla_{\theta} f(\mathbf{x}_i; \theta), \nabla_{\theta} f(\mathbf{x}_j; \theta) \rangle. \quad (8)$$

We put a detailed proof of these properties in Appendix A.2. Similarly, we will also show that our estimated Euclidean distance and inner product are close to their respective expectations with high probability.

The above analyses lead to the following conclusion: subset selection with $\{\widehat{\Delta z}(\mathbf{x}) | \mathbf{x} \in \mathcal{U}\}$ is approximately equivalent to that with $\{\nabla_{\theta} f(\mathbf{x}; \theta) | \mathbf{x} \in \mathcal{U}\}$, by using either distance-based (e.g. greedy-k-center) or similarity-based algorithms (e.g. k-DPP).

Connections to Variance Reduction

A widely accepted criterion in active learning is to select examples which minimize the predictive variance on existing training data (Cohn, Ghahramani, and Jordan 1996). This idea is established on the theory of variance reduction (MacKay 1992; Settles 2009). For convenience, we reuse previous notations but assume that $\hat{y} = f(\cdot; \theta)$ is a real-valued target function parameterized by θ that results in $\hat{y} : \mathbb{R}^d \rightarrow \mathbb{R}$, and this assumption can naturally generalize to vector-valued functions. Given a labeled training set $\mathcal{L} = \{\mathbf{x}_i, y_i\}_{i=1}^m$, the learning objective is to minimize the Mean Squared Error (S) as $S^2 = \frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i; \theta) - y_i)^2$. According to the theory of variance estimation (Thisted 1988), we derive that, by adding a new sample $\tilde{\mathbf{x}}$, the reduction in the predictive variance of a training example \mathbf{x} is:

$$\Delta \sigma_{\hat{y}}^2(\mathbf{x}) = \frac{\|\mathbf{J}_{\theta}(\mathbf{x}; \theta)^T \mathbf{A}_{S^2}^{-1} \mathbf{u}(\tilde{\mathbf{x}})\|^2 \cdot \|\mathbf{J}_{\theta}(\tilde{\mathbf{x}}; \theta)\|^2}{S^2 + \|\mathbf{u}(\tilde{\mathbf{x}})^T \mathbf{A}_{S^2}^{-1} \mathbf{u}(\tilde{\mathbf{x}})\| \cdot \|\mathbf{J}_{\theta}(\tilde{\mathbf{x}}; \theta)\|^2}, \quad (9)$$

where \mathbf{A}_{S^2} refers to the scaled Fisher Information Matrix: $\mathbf{A}_{S^2} = \frac{1}{S^2} \frac{\partial^2 S^2}{\partial \theta^2}$, and $\mathbf{u}(\tilde{\mathbf{x}})$ of the unit length denotes the direction of $\mathbf{J}_{\hat{y}}(\tilde{\mathbf{x}})$ as $\mathbf{u}(\tilde{\mathbf{x}}) = \frac{\mathbf{J}_{\theta}(\tilde{\mathbf{x}}; \theta)}{\|\mathbf{J}_{\theta}(\tilde{\mathbf{x}}; \theta)\|}$.

Eq (9) implies that, for any direction $\mathbf{u}(\tilde{\mathbf{x}})$, the optimal choice for maximizing $\Delta \sigma_{\hat{y}}^2(\mathbf{x})$ is the ones with the largest $\|\mathbf{J}_{\theta}(\tilde{\mathbf{x}}; \theta)\|^2$. On the other hand, diverse directions $\mathbf{u}(\tilde{\mathbf{x}})$ tend to benefit different training examples more uniformly, which contributes to the overall learning objective. Details for the derivation of Eq (9) is presented in Appendix A.3.

Comparison with to BADGE

Now we discuss more details comparing with BADGE (Ash et al. 2020), which is a state-of-the-art gradient-based active learning algorithm. In a multiclass classification problem with softmax activations, given a loss function L , BADGE yields the i -th block of the gradient of an example \mathbf{x} as

$$\nabla_{\phi_i} L = (p_i - I(\hat{y} = i)) f(\mathbf{x}; \theta), \quad (10)$$

where ϕ_i is the i -th row parameters of ϕ , and p_i is the probability for label i , thus $\hat{y} = \arg \max_i p_i$. Though our method (based on noise stability) is highly relevant to gradients as shown in Eq (4), there are two essential differences with BADGE (as well as other related works).

- Noise stability avoids explicitly approximating gradients, and hence does not require the extensive computation of assuming and marginalizing pseudo labels. This issue can be fatal in complex scenarios. For example, in semantic segmentation (or speech recognition), the number of label assumptions increases exponentially with increase in the number of pixels (or timesteps). Therefore, it is infeasible to gather the gradients w.r.t. all possible labels.
- In terms of final representations, noise stability resorts to *projected gradients* w.r.t. feature, while BADGE actually relies on *feature itself* and probabilistic predictions. In terms of uncertainty estimation, the gradient magnitude $\|\nabla_{\theta} f(\mathbf{x}; \theta)\|_2$ is a more meaningful indicator than the magnitude of $f(\mathbf{x}; \theta)$. Besides, in a regression problem, Eq (10) is meaningless due to the infeasibility of assuming

a ground truth label and has to resort to the original feature only. In contrast, noise stability is still applicable, taking into account both uncertainty and diversity.

Implementation Tips

In practice, for a deep neural network we can compute the perturbation $\hat{\Delta}z(\mathbf{x})$ w.r.t. the final deep feature. However, for a classification network, it is also feasible to employ the final predictive output to measure noise stability, since such an output can be regarded as a high-level semantic feature. Moreover, the predictive output usually compresses the primitive deep feature, thus resulting in a vector with a lower dimension. In applications with a large pool size and a high annotation budget, adopting the predictive output will naturally lead to an efficient implementation. Therefore, we suggest to compute the deviation on the predictive output for classification tasks, and on the final deep feature for other tasks such as regression and semantic segmentation.

Experimental Results

We compare our method with the state-of-the-art active learning baselines. Since Random selection is simple yet effective, we also include it as a baseline. We firstly validate our method on image classification tasks. Then we conduct extensive experiments on multiple tasks in various domains, including regression, semantic segmentation, as well as natural language processing. The model architectures that we use include CNN, MLP and BERT. We perform 7 to 10 active learning cycles, and the annotation budget for each cycle ranges from 20 to more than 3000. All the data selections are carried without replacement, following the common practice in active learning literature. All the reported results are averaged over 3 runs for a reliable evaluation, *unless specified otherwise*. For clear comparison, we plot the average accuracy of each cycle in figures, and additionally calculate AUBC (area under the budget curve) with the mean and variance. We conduct all the experiments using Pytorch (Paszke et al. 2017) and our source code will be publicly available.

Image Classification

Setting. We use three classification datasets for evaluations, including MNIST (LeCun et al. 1998), Cifar10 (Krizhevsky, Hinton et al. 2009) and SVHN (Netzer et al. 2011). MNIST consists of 60000 training samples and we only use a small portion since it is an easy task. The experiments on MNIST are repeated 10 times and we report the averaged results. The Cifar10 dataset includes 50000 training and 10000 testing samples, uniformly distributed across 10 classes. SVHN includes 73257 training and 26032 testing samples, and we do not use the additional training data in SVHN, following the common practice. SVHN is a typical imbalanced dataset.

We use different architectures for these three tasks, including a small CNN for MNIST, VGG19 (Simonyan and Zisserman 2014) for Cifar10 and ResNet-18 (He et al. 2016) for SVHN. The network and training details are described in Appendix A.4.1. Further, in Appendix A.4.2, we present additional results on Cifar10, Cifar100 and Caltech101 with ResNet-18 (He et al. 2016).




Figure 1: Classification performance due to the different AL methods on MNIST (left), Cifar10 (middle) and SVHN (right). AUBC with mean and variance is presented in Appendix A.4.3.

Observations. We compare our method with the state-of-the-art baselines and illustrate the results in Figure 1. The baseline methods include Random selection, Entropy, Least Margin, BALD (Houlsby et al. 2011; Gal, Islam, and Ghahramani 2017), QBC (Kuo et al. 2018), CoreSet (Sener and Savarese 2018), VAAL (Sinha, Ebrahimi, and Darrell 2019), LL4AL (Yoo and Kweon 2019), BADGE (Ash et al. 2020), CoreGCN (Caramalau, Bhattarai, and Kim 2021) and ALFA-Mix (Parvaneh et al. 2022). For our method, we use the following hyper-parameters in all the experiments: $K = 30$ and $\zeta = 10^{-3}$, to ensure the constraint on noise magnitude. The number of Monte-Carlo sampling for BALD is 50. Other hyper-parameters for baseline methods are set as suggested in the original papers.

As shown in Figure 1, our method (*NoiseStability*) outperforms or is comparable to the best baselines at all active learning cycles. In particular, our method usually yields a more rapid and stable increase of accuracy at the first several cycles, implying that noise stability is not sensitive to the predictive accuracy and calibration. In addition, the superior results in the SVHN (and Caltech101 in Appendix) experiments demonstrate that our method is robust to imbalanced datasets. We conjecture this is due to the simplicity of our method. In contrast, training auxiliary models (Kuo et al. 2018; Sinha, Ebrahimi, and Darrell 2019) with imbalanced data could be adverse to the auxiliary models themselves, thus leading to an inferior performance of the task model.

Other Tasks and Data Types

In this section, we validate our method with more diverse tasks, including regression, semantic segmentation and paraphrase recognition. For the regression task, we use data that includes both continuous and discrete variables, and we compare our method with the general algorithms that do not rely on probabilistic predictions. For the semantic segmentation task, we focus on the baselines that have been shown very effective for image data, e.g. LL4AL and VAAL. For the paraphrase recognition task, we use natural language data and make comparisons with the competitive baseline methods designed for classification problems, such as BALD and BADGE. General baselines Random and CoreSet are evalu-

ated on all the three benchmarks. We present a brief result below, and refer readers to Appendix A.5.1 for detailed descriptions of the datasets and experimental configurations.

Regression Problem We use the Ames Housing dataset (De Cock 2011) for this task. An MLP model takes a 215-dimensional feature as input and is trained to predict a real value as the house price. Compared with classification, the predictive output of regression does not explicitly contain information about uncertainty, which makes a part of active learning algorithms unavailable. The reported results are averaged over 50 runs for a reliable evaluation. As illustrated in Figure 2 (left), our method always yields the lowest mean absolute error before the performance saturation. We also observe that, the naive random selection becomes very competitive after the first two cycles, indicating that it is challenging for active learning to outperform passive learning in regression problems.

Semantic Segmentation Semantic segmentation is more complex than image classification due to the requirement of high classification accuracy at pixel level. As a consequence, in several baseline methods, such as VAAL (Sinha, Ebrahimi, and Darrell 2019), the auxiliary models need a careful design in order to achieve an satisfactory performance. This increases the burden of deployment in real scenarios. In contrast, our method is free of extra design and can be easily adopted for this application, demonstrating its task-agnostic advantage. We illustrate the results in Figure 2 (middle), and observe a remarkable improvement of our method. Given the Cityscapes dataset is highly imbalanced, this experiment once again demonstrates the superiority of our method on imbalanced datasets.

Natural Language Processing To further demonstrate the domain-agnostic nature of our method, we conduct an experiment on the Microsoft Research Paraphrase Corpus (MRPC) dataset, a classical NLP task aiming to identify whether two sentences are paraphrases of each other, by fine-tuning a pre-trained BERT model. As shown in Figure 2 (right), our method exhibits significant superiority, demonstrating its wide applicability, whereas other state-of-the-art baselines perform less competitive in this task.




Figure 2: Active learning performance due to the different methods on the regression task: Ames (left), semantic segmentation task: Cityscapes (middle), and natural language processing task: MRPC (right). AUBC results are presented in Appendix A.5.2.

Discussion

In this section, we present further discussions on the efficacy of the proposed method, including the choice of hyperparameters, ablation study and computational efficiency. Here we present brief results and put details in Appendix A.6-A.8.

Sensitivity to Hyper-parameters (Relevant to Noise)

We firstly investigate how our algorithm is sensitive to the hyper-parameters (i.e. noise magnitude and sampling times). We conduct this set of experiments on MNIST.

Noise Magnitude. An appropriate setting of the noise magnitude will benefit our method. Specifically, we observe that noise with a too large magnitude tends to destroy the learned knowledge in DNNs. We evaluate the choices of $\{10^{-6}, 10^{-4}, 10^{-2}, 1, 10\}$ respectively for the ζ . The results in Appendix A.6 (left) show that a ζ of between 10^{-4} and 10^{-2} is likely to yield good performance. A larger ζ leads to dramatic performance decrease at early cycles. Interestingly, we observe that a too small ζ also tends to be sub-optimal. We conjecture the reason may lie in that, if the noise is too trivial, only those activated neurons have output perturbations, making the final output deviation less informative.

Noise Sampling Times. Accurate approximations such as that for Eq (6) usually require considerable sampling times. However, we surprisingly find that our method is not that sensitive to the sampling times K . We study the choices of $\{1, 3, 10, 30, 50\}$ respectively for K . As shown in Appendix A.6 (middle), using $K = 10$ can deliver a very competitive performance. We conjecture this is because a random noise is sufficient to produce rich pattern perturbations for the parameters in DNNs, making the output deviation (due to only a few noise samplings) a reasonable criterion for unlabeled data selection (considering both magnitude and direction).

Subset Selection. The topic of subset selection has been well studied in the machine learning community. Other classical algorithms can be easily integrated into our framework. We validate this by replacing our greedy k-center algorithm with kmeans++ (Arthur and Vassilvitskii 2006) (suggested by BADGE) and k-DPP (Kulesza and Taskar 2011). Results in Appendix A.6 (right) show that the three algorithms yield quite similar results in our framework.

Ablation Study

Here we demonstrate that both the gradient magnitude and direction are informative in subset selection. Specifically, we compare our method with its two degenerated versions including (1) *NoiseStability-M* which selects samples with maximum magnitude $\|\Delta z(\mathbf{x})\|_2$; and (2) *NoiseStability-D* which normalizes each $\Delta z(\mathbf{x})$ to unit magnitude before performing the greedy k-center selection. We observe in Appendix A.7 (left) that, neither of the two sub-optimal implementations can compete with the full version of our method.

We then investigate the effectiveness of the magnitude $\|\Delta z(\mathbf{x})\|_2$ on uncertainty estimation by evaluating single sample selection. We compare our method with the popular deep Bayesian active learning baselines on the MNIST dataset. For fair comparison, we perform 50 samplings for all the methods. As shown in Appendix A.7 (right), ours is superior to the baselines, demonstrating that noise stability is a reliable indicator of example uncertainty.

Time Efficiency

The time efficiency of our method depends on the cost of noise sampling and diverse subset selection. With the same hyperparameters as used in experiments, our method exhibits not only comparable computational efficiency among state-of-the-art algorithms, but also good scalability in terms of time complexity, as the annotation budget or category size increases (Appendix A.8).

Conclusion

We propose a simple yet robust method that leverages noise stability for uncertainty estimation in deep active learning. For an input sample, noise stability measures the stability of the output when a random perturbation is imposed on the model parameters. The direction of the feature deviation is also utilized to ensure the diversity of selected examples. Our theoretical analyses prove that, the resulting deviation under a small parameter noise acts as a reasonable proxy of the gradient, in terms of both magnitude and direction. The extensive evaluations on multiple tasks demonstrate the superiority of the proposed method.

References

- Arora, S.; Ge, R.; Neyshabur, B.; and Zhang, Y. 2018. Stronger generalization bounds for deep nets via a compression approach. *ArXiv*, abs/1802.05296.
- Arthur, D.; and Vassilvitskii, S. 2006. k-means++: The advantages of careful seeding. Technical report, Stanford.
- Ash, J.; Goel, S.; Krishnamurthy, A.; and Kakade, S. 2021. Gone fishing: Neural active learning with fisher embeddings. *Advances in Neural Information Processing Systems*, 34: 8927–8939.
- Ash, J. T.; Zhang, C.; Krishnamurthy, A.; Langford, J.; and Agarwal, A. 2020. Deep Batch Active Learning by Diverse, Uncertain Gradient Lower Bounds. In *International Conference on Learning Representations*.
- Bishop, C. M. 1995. Training with noise is equivalent to Tikhonov regularization. *Neural computation*, 7(1): 108–116.
- Caramalau, R.; Bhattarai, B.; and Kim, T.-K. 2021. Sequential Graph Convolutional Network for Active Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9583–9592.
- Citovsky, G.; DeSalvo, G.; Gentile, C.; Karydas, L.; Rajagopalan, A.; Rostamizadeh, A.; and Kumar, S. 2021. Batch active learning at scale. *Advances in Neural Information Processing Systems*, 34: 11933–11944.
- Cohn, D.; Atlas, L.; and Ladner, R. 1994. Improving generalization with active learning. *Machine learning*, 15(2): 201–221.
- Cohn, D. A.; Ghahramani, Z.; and Jordan, M. I. 1996. Active Learning with Statistical Models. *The Journal of artificial intelligence research*, 4: 129–145.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Dasgupta, S.; and Hsu, D. 2008. Hierarchical sampling for active learning. In *Proceedings of the 25th international conference on Machine learning*, 208–215.
- Daxberger, E.; Nalisnick, E.; Allingham, J. U.; Antorán, J.; and Hernández-Lobato, J. M. 2021. Bayesian deep learning via subnetwork inference. In *International Conference on Machine Learning*, 2510–2521. PMLR.
- De Cock, D. 2011. Ames, Iowa: Alternative to the Boston housing data as an end of semester regression project. *Journal of Statistics Education*, 19(3).
- Dolan, W. B.; and Brockett, C. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Ducoffe, M.; and Precioso, F. 2018. Adversarial Active Learning for Deep Networks: a Margin Based Approach. *arXiv Preprint arXiv:1802.09841*.
- Elhamifar, E.; Sapiro, G.; Yang, A.; and Shankar Sasrty, S. 2013. A Convex Optimization Framework for Active Learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 209–216.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2006. One-Shot Learning of Object Categories. *IEEE TPAMI*, 28(4): 594–611.
- Freund, Y.; Seung, H. S.; Shamir, E.; and Tishby, N. 1997a. Selective Sampling Using the Query by Committee Algorithm. *Machine Learning*, 28(2-3): 133–168.
- Freund, Y.; Seung, H. S.; Shamir, E.; and Tishby, N. 1997b. Selective sampling using the query by committee algorithm. *Machine learning*, 28(2): 133–168.
- Gal, Y.; and Ghahramani, Z. 2016. Dropout as A Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *International Conference on Machine Learning*.
- Gal, Y.; Islam, R.; and Ghahramani, Z. 2017. Deep Bayesian Active Learning with Image Data. In *International Conference on Machine Learning*.
- Gonzalez, T. F. 1985. Clustering to minimize the maximum intercluster distance. *Theoretical computer science*, 38: 293–306.
- Gorritz, M.; Carlier, A.; Faure, E.; and Giro-i Nieto, X. 2017. Cost-effective active learning for melanoma segmentation. *arXiv Preprint arXiv:1711.09168*.
- Hasan, M.; and Roy-Chowdhury, A. K. 2015. Context aware active learning of activity recognition models. In *Proceedings of the IEEE International Conference on Computer Vision*, 4543–4551.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Houlsby, N.; Huszár, F.; Ghahramani, Z.; and Lengyel, M. 2011. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*.
- Huang, J.; Child, R.; Rao, V.; Liu, H.; Satheesh, S.; and Coates, A. 2016. Active learning for speech recognition: the power of gradients. *arXiv preprint arXiv:1612.03226*.
- Immer, A.; Bauer, M.; Fortuin, V.; Rätsch, G.; and Khan, M. E. 2021. Scalable marginal likelihood estimation for model selection in deep learning. *arXiv preprint arXiv:2104.04975*.
- Kirsch, A.; Van Amersfoort, J.; and Gal, Y. 2019. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *Advances in neural information processing systems*, 32: 7026–7037.
- Koh, P. W.; and Liang, P. 2017. Understanding Black-box Predictions via Influence Functions. In *International Conference on Machine Learning*.
- Konyushkova, K.; Sznitman, R.; and Fua, P. 2017. Learning active learning from data. In *Proceedings of the Advances in Neural Information Processing Systems*, 4225–4235.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning Multiple Layers of Features from Tiny Images.
- Kulesza, A.; and Taskar, B. 2011. k-DPPs: Fixed-size determinantal point processes. In *ICML*.
- Kuo, W.; Häne, C.; Yuh, E.; Mukherjee, P.; and Malik, J. 2018. Cost-Sensitive Active Learning for Intracranial Hemorrhage Detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*.

- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Liu, Z.; Ding, H.; Zhong, H.; Li, W.; Dai, J.; and He, C. 2021. Influence selection for active learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9274–9283.
- Loquercio, A.; Segu, M.; and Scaramuzza, D. 2020. A general framework for uncertainty estimation in deep learning. *IEEE Robotics and Automation Letters*, 5(2): 3153–3160.
- MacKay, D. J. C. 1992. Information-Based Objective Functions for Active Data Selection. *Neural computation*, 4(4): 590–604.
- Mayer, C.; and Timofte, R. 2020. Adversarial sampling for active learning. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 3071–3079.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading Digits in Natural Images with Unsupervised Feature Learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*.
- Nguyen, H. T.; and Smeulders, A. 2004. Active learning using pre-clustering. In *International Conference on Machine Learning*.
- Parvaneh, A.; Abbasnejad, E.; Teney, D.; Haffari, R.; van den Hengel, A.; and Shi, J. Q. 2022. Active Learning by Feature Mixing.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic Differentiation in Pytorch.
- Ren, P.; Xiao, Y.; Chang, X.; Huang, P.-Y.; Li, Z.; Gupta, B. B.; Chen, X.; and Wang, X. 2021. A survey of deep active learning. *ACM Computing Surveys (CSUR)*, 54(9): 1–40.
- Sener, O.; and Savarese, S. 2018. Active Learning for Convolutional Neural Networks: A Core-Set Approach. In *International Conference on Learning Representations*.
- Settles, B. 2009. Active learning literature survey.
- Settles, B.; and Craven, M. 2008. An analysis of active learning strategies for sequence labeling tasks. In *proceedings of the 2008 conference on empirical methods in natural language processing*, 1070–1079.
- Simonyan, K.; and Zisserman, A. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv Preprint arXiv:1409.1556*.
- Sinha, S.; Ebrahimi, S.; and Darrell, T. 2019. Variational Adversarial Active Learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 5972–5981.
- Thisted, R. A. R. A. 1988. *Elements of statistical computing*. New York: Chapman and Hall. ISBN 0412013711.
- Tong, S.; and Koller, D. 2001. Support Vector Machine Active Learning with Applications to Text Classification. *Journal of Machine Learning Research*, 2(Nov): 45–66.
- Wang, T.; Li, X.; Yang, P.; Hu, G.; Zeng, X.; Huang, S.; Xu, C.-Z.; and Xu, M. 2022. Boosting active learning via improving test performance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8566–8574.
- Yang, Y.; Ma, Z.; Nie, F.; Chang, X.; and Hauptmann, A. G. 2015. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113(2): 113–127.
- Yoo, D.; and Kweon, I. S. 2019. Learning Loss for Active Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 93–102.
- Zhang, B.; Li, L.; Yang, S.; Wang, S.; Zha, Z.-J.; and Huang, Q. 2020. State-Relabeling Adversarial Active Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8756–8765.
- Zhang, C.; Kjellström, H.; and Mandt, S. 2017. Determinantal point processes for mini-batch diversification. In *33rd Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, 11 August 2017 through 15 August 2017*. AUAI Press Corvallis.
- Zhang, Y.; Lease, M.; and Wallace, B. 2017. Active discriminative text representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Žliobaitė, I.; Bifet, A.; Pfahringer, B.; and Holmes, G. 2013. Active learning with drifting streaming data. *IEEE transactions on neural networks and learning systems*, 25(1): 27–39.