

Tricks and Plugins to GBM on Images and Sequences

Biyi Fang

Department of Engineering Science and Applied Mathematics

Northwestern University

`biyifang2021@u.northwestern.edu`

Jean Utke

Data Discovery & Decision Science

Allstate

`jutke@allstate.com`

Diego Klabjan

Department of Industrial Engineering and Management Sciences

Northwestern University

`d-klabjan@northwestern.edu`

February 16, 2022

Abstract

Convolutional neural networks (CNNs) and transformers, which are composed of multiple processing layers and blocks to learn the representations of data with multiple abstract levels, are the most successful machine learning models in recent years. However, millions of parameters and many blocks make them difficult to be trained, and sometimes several days or weeks are required to find an ideal architecture or tune the parameters. Within this paper, we propose a new algorithm for boosting Deep Convolutional Neural Networks (BoostCNN) to combine the merits of dynamic feature selection and BoostCNN, and another new family of algorithms combining boosting and transformers. To learn these new models, we introduce subgrid selection and importance sampling strategies and propose a set of algorithms to incorporate boosting weights into a deep learning architecture based on a least squares objective function. These algorithms not only reduce the required manual effort for finding an appropriate network architecture but also result in superior performance and lower running time. Experiments show that the proposed methods outperform benchmarks on several fine-grained classification tasks.

1 Introduction

Deep convolutional neural networks (CNNs) and transformers such as BERT have had great recent success in learning image representations for vision tasks and NLP, respectively. Given the outstanding results produced by these networks, they have been widely applied in image classification ([19], [40], [27]), object detection ([15], [21], [39]), speech recognition, ([33], [44], [14]), and language translation ([25], [13], [46]). However, an optimal image or text representation for each task is unique and finding an optimal deep neural network structure is a challenging problem. There are some approaches (neural architecture search) for designing these deep networks such as AutoML for Model Compression (AMC) in [20] and LEAF in [26], however, these methods require weeks of training on thousands of GPUs. Meanwhile, ensemble methods for classification and regression have gained a lot of attention in recent years, which perform, both theoretically and empirically, substantially better than single models in a wide range of tasks, i.e. boosting decision trees [38]. In order to tackle the design challenge specifically for CNNs, an idea of combining boosting and shallow CNNs is proposed in [5]. Their idea is to simplify the complicated design process of deep neural networks by employing the boosting strategy which combines the strengths of multiple CNNs. However, the memory requirement and running time become challenging when the weak learner is not extremely simple. Moreover, very limited contribution has been made to the case when the weak learner is a transformer. Furthermore, no work has been conducted around the idea of only using partial data with weak learners.

In this paper, we propose a family of boosting algorithms for images, namely subgrid BoostCNN, and another family of boosting algorithms for sequences, namely BoostTransformer, which are both based on boosting, deep CNNs and transformers. We select a subset of features for each weak learner, where the concepts are borrowed from random forests. This strategy requires new ideas in order to accommodate unstructured data. Moreover, we apply the concept of importance sampling, which assigns a probability to each sample, to the combination of boosting and a transformer.

Subgrid BoostCNN aims to solve the same problem as deep CNNs but provides higher accuracy with lower running time and memory requirements. Subgrid BoostCNN builds on the previous boosting Deep Convolutional Neural Networks (BoostCNN) [5]. One important new aspect in subgrid BoostCNN is that it does not require a full image for training a weak learner; instead, it selects only important pixels based on the gradient from each image combined with the corresponding residual to train the current weak learner. This implies breaking the original relationship between a pixel and its neighborhood, thus possibly leading to noisier training, but the subgrid BoostCNN does reduce computation to important pixels. Another option to reduce the running time is to skip the optimization process for the original full CNN, which finds the important pixels for the weak learner. Instead, we borrow the CNN portion from the last weak learner concatenated with the fully connected layer used in the first iterate to compute the importance value of each pixel, and train the CNN concatenated with an appropriate fully connected layer. Consequently, subgrid BoostCNN does the optimization process once in each iteration, which is the same as BoostCNN, while subgrid BoostCNN has fewer parameters when compared with BoostCNN. This subgrid trick is essential especially when the training process for the weak learner is computationally demanding. Furthermore, we demonstrate subgrid BoostCNN on three different image datasets

and argue that subgrid BoostCNN outperforms both BoostCNN and deep CNNs. More precisely, subgrid BoostCNN improves the accuracies by 1.16%, 0.82%, 12.10% on CIFAR-10, SVHN and ImageNetSub datasets, respectively, when compared to standard CNN models. In addition, subgrid BoostCNN obtains accuracies 0.34%, 0.50%, 4.19% higher than those generated by BoostCNN on the aforementioned datasets, respectively.

BoostTransformer is an algorithm which combines the merits of boosting and transformers. BoostTransformer incorporates boosting weights with transformers based on least squares objective functions. Motivated by the successful combination of BoostCNN and the subgrid trick, we propose subsequence BoostTransformer, which does not require the full data for training weak learners. In subsequence BoostTransformer, important tokens, which are from the input, are selected for each weak learner based on the attention distribution [43]. Similarly, we might lose the direct connections between consecutive words, while informative words are retained for learning. Consequently, subsequence BoostTransformer takes less time to achieve a better accuracy when compared to vanilla BoostTransformer. Moreover, motivated by the phenomenon that overfitting in BoostTransformer appears early, we propose a new algorithm, namely importance-sampling-based BoostTransformer, which combines the merits of BoostTransformer and importance sampling. Importance-sampling-based BoostTransformer first computes a probability distribution for all the samples in the dataset; then in each iteration, it randomly chooses a subset of samples based on the pre-computed probability distribution; lastly, similar to BoostTransformer, it trains the weak learner on the selected samples. This algorithm not only delays overfitting, but also improves the accuracy and significantly reduces the running time. We present a complete technical proof for importance-sampling-based BoostTransformer showing that the optimal probability distribution is proportional to the norm of the residuals. Lastly, we conduct computational experiments demonstrating a superior performance of the proposed algorithms. More precisely, BoostTransformer provides higher accuracy and more stable solutions when compared to transformers. Moreover, subsequence BoostTransformer and importance-sampling-based BoostTransformer not only provide better and more robust solutions but also dramatically reduce the running time when compared to transformers. Compared to standard transformers, BoostTransformer, subsequence BoostTransformer and importance-sampling-based BoostTransformer provide an average of 0.87%, 0.55%, 0.79% accuracy improvements, respectively, on IMDB, Yelp and Amazon datasets. Furthermore, subsequence BoostTransformer and importance-sampling-based BoostTransformer take only two thirds and one half of time transformers need to learn the datasets, respectively.

In summary, we make the following contributions.

- We provide a better boosting method for deep CNNs, i.e. subgrid BoostCNN, which requires only important pixels from the image dataset where such pixels are selected dynamically for each weak learner.
- We provide a boosting method for sequences, i.e. BoostTransformer, which combines the merits of boosting and transformers.
- We provide a better boosting method for transformers, i.e. subsequence BoostTransformer, which does not require the full sequences but only important tokens.

- We provide another enhancement for BoostTransformer, i.e. importance-sampling-based BoostTransformer, which combines importance sampling and BoostTransformer. Moreover, we provide a proof showing that the optimal probability distribution for the samples is proportional to the norm of the residuals.
- We present numerical results showing that subsequence BoostTransformer and importance-sampling-based BoostTransformer outperform vanilla transformers on select tasks and datasets.

The rest of the paper is organized as follows. In the next section, we review several related works in gradient boosting machine, CNN and transformers. In Section 3, we state the formal optimization problem and provide the exposition of the subgrid BoostCNN. In the subsequent section, we propose BoostTransformer, subsequence BoostTransformer and importance-sampling-based BoostTransformer, followed by the analysis of the optimal probability distribution for importance-sampling-based BoostTransformer. In Section 5, we present experimental results comparing the different algorithms.

2 Related Work

There are many extensions of Gradient Boosting Machine (GBM) [34], however, a full retrospective of this immense literature exceeds the scope of this work. In this section, we mainly state several kinds of variations of GBM which are most related to our new algorithms, together with the two add-ons to our optimization algorithms, i.e. subgrid and importance sampling.

Boosting for CNNs: Deep CNNs, which have recently produced outstanding performance in learning image representations, are capable of learning complex features that are highly invariant and discriminant [16]. The success of deep CNNs in recognizing objects has encouraged recent works to combine boosting together with deep CNNs. Brahim & Aoun [6] propose a new Boosted Convolutional Neural Network architecture, which uses a very deep convolutional neural network reinforced by adding Boosted Blocks which consist of a succession of convolutional layers boosted by using a Multi-Bias Nonlinear Activation function. Nevertheless, the architecture of the proposed Boosted convolutional neural network is fixed; it can not dynamically change the number of Boosted Blocks for a given dataset. Another attempt at combining deep CNNs and boosting is boosted sampling [4], which uses posterior error maps, generated throughout training, to focus sampling on different regions, resulting in a more informative loss. However, boosted sampling applies boosting on selected samples and treats deep CNN as a black box to make a prediction. To enrich the usage of the information generated by deep CNNs, Lee & Chen [24] propose a new BoostCNN structure which employs a trained deep convolutional neural network model to extract the features of the images and then use the AdaBoost algorithm to assemble the Softmax classifiers. However, it remains unclear how to combine different sets of the features extracted and the computational cost is high when training several deep CNNs at the same time. To tackle this problem, Han & Meng [17] propose Incremental Boosting CNN (IB-CNN) to integrate boosting into the CNN via an incremental boosting layer that selects discriminative neurons from a lower layer and is incrementally updated on successive mini-batches. Different from

IB-CNN which only involves one deep CNN, BoostCNN [5] incorporates boosting weights into the neural network architecture based on least squares objective functions, which leads to the aggregation of several CNNs. However, the computational and memory demand of BoostCNN is high when the weak learner is not simple. All of the above train the weak learners on all features.

Boosting for Recurrent Neural Network (RNN) and Transformer: RNN, long short-term memory (LSTM) and transformers have been firmly established as state of the art approaches in sequence modeling and transduction problems such as language modeling and machine translation ([3], [8], [42], [43]). Some efforts have been made to combine boosting with RNN or LSTM. Chen & Lundberg [7] present feature learning via LSTM networks and prediction via gradient boosting trees (XGB). More precisely, they generate features by performing supervised representation learning with an LSTM network, then augment the original XGB model with these new generated features. However, the selection of the features from LSTM is not determined by XGB, which leads to a disconnect between LSTM and XGB. Another attempt at combining boosting and RNN is the boosting algorithm for regression with RNNs [2]. This algorithm adapts an ensemble method to the problem of predicting future values of time series using RNNs as base learners, and it is based on the boosting algorithm where different points of the time series are emphasized during the learning process by training different base learners on different subsets of time points. However, combining boosting and transformers has not previously been investigated. Although, analyses of attention in Transformer have been explored [9], using the attention distribution in token selection has not been extensively studied.

Given the fact that importance sampling improves the performance by prioritizing training samples, importance sampling has been well studied, both theoretically and empirically, in standard stochastic gradient descent settings [35] [45], in deep learning settings [22], and in minibatches [10]. As stated in these papers, importance sampling theoretically improves the convergence rate and is experimentally effective in reducing the training time and training loss. However, no generalization work has been done in a boosting setting.

3 Algorithms for CNN as Weak Learner

In this section, we provide a summary of BoostCNN and propose a new algorithm, subgrid CNN, which combines BoostCNN and the subgrid trick.

3.1 Background: Standard BoostCNN

We start with a brief overview of multiclass boosting. Given a sample $x_i \in \mathcal{X}$ and its class label $z_i \in \{1, 2, \dots, M\}$, multiclass boosting is a method that combines several multiclass predictors $g_t : \mathcal{X} \rightarrow \mathbb{R}^d$ to form a strong committee $f(x)$ of classifiers, i.e. $f(x) = \sum_{t=1}^N \alpha_t g_t(x)$ where g_t and α_t are the weak learner and coefficient selected at the t^{th} boosting iteration. There are various approaches for multiclass boosting such as [18], [32], [41]; we use the GD-MCBoost method of [41], [5] herein. For simplicity, in the rest of the paper, we assume that $d = M$.

Standard BoostCNN [5] trains a boosted predictor $f(x)$ by minimizing the risk of classification

$$\mathcal{R}[f] = \mathbb{E}_{X,Z} [L(z, f(x))] \approx \frac{1}{|\mathcal{D}|} \sum_{(x_i, z_i) \in \mathcal{D}} L(z_i, f(x_i)), \quad (1)$$

where \mathcal{D} is the set of training samples and

$$L(z, f(x)) = \sum_{j=1, j \neq z}^M e^{\frac{1}{2}[\langle y_z, f(x) \rangle - \langle y_j, f(x) \rangle]},$$

given $y_k = \mathbb{1}_k \in \mathbb{R}^M$, i.e. the k^{th} unit vector. The minimization is via gradient descent in a functional space. Standard BoostCNN starts with $f(x) = \mathbf{0} \in \mathbb{R}^d$ for every x and iteratively computes the directional derivative of risk (1), for updating $f(x)$ along the direction of $g(x)$

$$\begin{aligned} \delta \mathcal{R}[f; g] &= \left. \frac{\partial \mathcal{R}[f + \epsilon g]}{\partial \epsilon} \right|_{\epsilon=0} = -\frac{1}{2|\mathcal{D}|} \sum_{(x_i, z_i) \in \mathcal{D}} \sum_{j=1}^M g_j(x_i) w_j(x_i, z_i) \\ &= -\frac{1}{2|\mathcal{D}|} \sum_{(x_i, z_i) \in \mathcal{D}} g(x_i)^T w(x_i, z_i), \end{aligned} \quad (2)$$

where

$$w_k(x, z) = \begin{cases} -e^{-\frac{1}{2}[f_z(x) - f_k(x)]}, & k \neq z \\ \sum_{j=1, j \neq k}^M e^{-\frac{1}{2}[f_z(x) - f_j(x)]}, & k = z, \end{cases} \quad (3)$$

and $g_j(x_i)$ computes the directional derivative along $\mathbb{1}_j$. Then, standard BoostCNN selects a weak learner g^* that minimizes (2), which essentially measures the similarity between the boosting weights $w(x_i, z_i)$ and the function values $g(x_i)$. Therefore, the optimal network output $g^*(x_i)$ has to be proportional to the boosting weights, i.e.

$$g^*(x_i) = \beta w(x_i, z_i), \quad (4)$$

for some constant $\beta > 0$. Note that the exact value of β is irrelevant since $g^*(x_i)$ is scaled when computing α^* . Consequently, without loss of generality, we assume $\beta = 1$ and convert the problem to finding a network $g(x) \in \mathbb{R}^M$ that minimizes the square error loss

$$\mathcal{L}(w, g) = \sum_{(x_i, z_i) \in \mathcal{D}} \|g(x_i) - w(x_i, z_i)\|^2. \quad (5)$$

After the weak learner is trained, BoostCNN applies a line search to compute the optimal step size along g^* ,

$$\alpha^* = \underset{\alpha \in \mathbb{R}}{\operatorname{argmin}} \mathcal{R}[f + \alpha g^*]. \quad (6)$$

Finally, the boosted predictor $f(x)$ is updated as $f = f + \alpha^* g^*$.

3.2 Subgrid BoostCNN

When considering full-size images, BoostCNN using complex CNNs as weak learners is time-consuming and memory hungry. Consequently, we would like to reduce the size of the images to lower the running time and the memory requirement. A straightforward idea would be downsizing the images directly. A problem of this approach is that the noise would possibly spread out to later learners since a strong signal could be weakened during the downsize process. Another candidate for solving the aforementioned problem is randomly selecting pixels from the original images, however, the fluctuation of the performance of the algorithm would be significant especially when the images are sharp or have a lot of noise. In this paper, we apply the subgrid trick to each weak learner in BoostCNN. The remaining question is how to select a subgrid for each weak learner. Formally, a subgrid is defined by deleting a subset of rows and columns. Moreover, the processed images may not have the same size between iterations, which in turn requires that the new BoostCNN should allow each weak learner to have a at least different dimensions. However, that impedes reusing weak learner model parameters from one weak learner iterate to next.

In order to address these issues, we first separate a standard deep CNN into two parts. We call all layers such as convolutional layers and pooling layers, except the last fully-connected (FC) layers, the *feature extractor*. In contrast, we call the last FC layers the *classifier*. Furthermore, we refer to g_0 as the basic weak learner and all the succeeding g_t as the additive weak learners. Subgrid BoostCNN defines an importance index for each pixel (j, k) in the image as

$$I_{j,k} = \frac{1}{|\mathcal{D}|} \sum_{(x_i, z_i) \in \mathcal{D}} \sum_{c \in C} \left| \frac{\partial \mathcal{L}(w, g)}{\partial x_i^{j,k,c}} \right|, \quad (7)$$

where $x_i^{j,k,c}$ denotes pixel (j, k) in channel c from sample i and C represents the set of all channels. The importance index of a row, column is a summation of the importance indexes in the row, column divided by the number of columns, rows, respectively. This importance index is computed based on the residual of the current predictor. Therefore, a larger importance value means a larger adjustment is needed for this pixel at the current iterate. The algorithm uses the importance index generated based on the feature extractor of the incumbent weak learner and the classifier from g_0 to conduct subgrid selection. The selection strategy we apply in the algorithm is deleting less important columns and rows, which eventually provides the important subgrid. After the subgrid is selected, subgrid BoostCNN creates a new tensor x_i^t at iterate t , and then feeds it into an appropriate feature extractor followed by a proper classifier. The modified minimization problem becomes

$$\mathcal{L}(w, g) = \sum_{(x_i, z_i) \in \mathcal{D}} \|g(x_i^t) - w(x_i, z_i)\|^2, \quad (8)$$

where the modified boosting classifier is

$$f(x) = \sum_{t=1}^N \alpha_t g_t(x^t). \quad (9)$$

In this way, subgrid BoostCNN dynamically selects important subgrids based on the updated residuals. Moreover, subgrid BoostCNN is able to deal with inputs of different sizes by applying different classifiers. Furthermore, we are allowed to pass the feature extractor’s parameters from the previous weak learner since the feature extractor is not restricted to the input size. The proposed algorithm (subgrid BoostCNN) is summarized in Algorithm 1.

Algorithm 1 subgrid BoostCNN

- 1: **Inputs:**
 number of classes M , number of boosting iterations N_b , shrinkage
 parameter ν , dataset $\mathcal{D} = \{(x_1, z_1), \dots, (x_n, z_n)\}$ where
 $z_i \in \{1, \dots, M\}$ is the label of sample x_i , and $0 < \sigma < 1$
 - 2: **Initialize:**
 set $f(x) = \mathbf{0} \in \mathbb{R}^M$, $P_0 = \{(j, k) | (j, k) \text{ is a pixel in } x_i\}$
 - 3: compute $w(x_i, z_i)$ for all (x_i, z_i) , using (3)
 - 4: train a deep CNN g_0^* to optimize (5)
 - 5: $f(x) = g_0^*$
 - 6: **for** $t = 1, 2, \dots, N_b$ **do**
 - 7: update importance index $I_{j,k}$ for $(j, k) \in P_{t-1}$, using (7)
 - 8: select the subgrid based on σ fraction of rows and columns with highest importance
 index and let P_t be the set of selected pixels; form a new tensor x_i^t for each sample i
 - 9: construct a new proper weak learner architecture
 - 10: compute $w(x_i, z_i)$ for all i , using (3) and (9)
 - 11: train a deep CNN g_t^* to optimize (8)
 - 12: find the optimal coefficient α_t , using (6) and (9)
 - 13: $f(x) = f(x) + \nu\alpha_t g_t^*$
 - 14: **end for**
-

Subgrid BoostCNN starts by initializing $f(x) = \mathbf{0} \in \mathbb{R}^M$. The algorithm first generates a full-size deep CNN as the basic weak learner, which uses the full image in steps 3-4. After the basic weak learner g_0^* is generated, in each iteration, subgrid BoostCNN first updates the importance index $I_{j,k}$ for each pixel (j, k) , which has been used in the preceding iterate at step 7. In order to mimic the loss of the full-size image, although we only update the importance indexes for the pixels which have been used in the last iterate, we feed the full-size tensor to the deep CNN g to compute the importance index. The deep CNN g used in (7) to compute the importance value is constructed by copying the feature extractor from the preceding weak learner followed by the classifier in the basic weak learner g_0^* . Next, by deleting less important rows and columns based on $I_{j,k}$, which contain $1 - \sigma$ fraction of pixels, it finds the most important subgrid having σ fraction of pixels at position P_t based on the importance index $I_{j,k}$, and forms a new tensor x_i^t in step 8. Note that P_t is not necessary to be a subset of P_{t-1} and actually is rarely to be a subset of P_{t-1} . This only happens when the highest importance index at iterate t is also the highest score at iterate $t - 1$. Next, a new additive weak learner is initialized by borrowing the feature extractor from the preceding weak learner g_{t-1}^* followed by a randomly initialized FC layer with the proper size in step 9. Once the additive weak learner is initialized, subgrid BoostCNN computes

the boosting weights, $w(x) \in \mathbb{R}^M$ according to (3) and (9), trains a network g_t^* to minimize the squared error between the network output and boosting weights using (8), and finds the boosting coefficient α_t by minimizing the boosting loss (6) in steps 10-12. Lastly, the algorithm adds the network to the ensemble according to $f(x) = f(x) + \nu\alpha_t g_t^*$ for $\nu \in [0, 1]$ in step 13.

4 Algorithms for Transformer as Weak Learner

In this section, we propose three algorithms combining boosting and transformers from different perspectives. We assume a BERT-like bidirectional transformer classifier [12] [28]. The first token of each sequence is a special classification token, and the corresponding final hidden state output of this token is used as the aggregated representation for the classification.

4.1 Standard BoostTransformer

Inspired by BoostCNN, we propose BoostTransformer which combines boosting and transformers (encoder) together. For a sequence classification problem, we are given a sample $x_i \in \mathcal{X}$, which contains a sequence of tokens, and its class label $z_i \in \{1, 2, \dots, M\}$. The risk function, the functional gradient and the optimal boosting coefficient α^t are exactly the same as those in (1), (2), and (6), respectively. The algorithm follows standard gradient boosting machine.

4.2 Subsequence BoostTransformer

Combining the subgrid trick and BoostTransformer means applying the subgrid trick to each weak learner in BoostTransformer. Different from deep CNNs, transformers are able to deal with sequences of any length, thus, there is no issue when transferring information from the current weak learner to the succeeding weak learner. Similar to subgrid BoostCNN, we denote g_0 as the basic weak learner, which deals with the whole dataset, and all the succeeding g_t 's as the additive weak learners. Moreover, subsequence BoostTransformer defines an importance index for each token ω in the vocabulary based on the attention distribution. More precisely, the importance value of token ω is computed by adding two parts; the first part is the importance of the token ω itself, and the second part is the importance of token ω to the remaining tokens in the same sample. In an L -layer transformer for a sequence x of length s (following [12] we assume that the first token in x is a placeholder, which indicates that the corresponding token in the final layer is used as the embedding for classification), and positions $1 \leq i, j \leq s$, and layer k for $1 \leq k \leq L$, let the attention from position i to position j between layer $k-1$ and k be denoted by $a(i, j; k; x)$. We have $\sum_{j=1}^s a(j, i; k; x) = 1$ for every i, k, x . Then, given a transformer with L layers, the self-importance of token ω in position p in a sample x_i is

$$I^S(\omega, x_i) = \left(\prod_{k=1}^{L-1} a(p, p; k; x_i) \right) \cdot a(p, 1; L; x_i) \approx a(p, 1; L; x_i), \quad (10)$$

The importance of token ω to others is

$$I^R(\omega, x_i) = \left[\prod_{k=1}^{L-1} \max_{j, j \neq p} a(p_{k-1}, j; k; x_i) \right] \cdot a(p_{L-1}, 1; L; x_i), \quad (11)$$

where $p_{k-1} = \operatorname{argmax}_{j, j \neq p} a(p_{k-2}, j; k-1; x_i)$ for $k = 2, 3, \dots, L-1$, and $p_0 = p$. The first term computes the product of the maximum attention values through the path which does not contain p until the second to last layer. For the second term, as it has been shown in [12], the classification layer only takes the 1st position of the last transformer layer which is corresponding to the classification token, therefore, the formula in (11) does not check all possible attention distributions; instead, it counts the attention value from the position p_{L-1} to the 1st position in the last transformer layer directly. After the aforementioned importance values are computed, the importance value of the vocabulary word $\hat{\omega}$ is

$$I(\hat{\omega}) = \sum_{\substack{x_i, \omega \in x_i \\ \omega = \hat{\omega}}} (I^S(\omega, x_i) + I^R(\omega, x_i)). \quad (12)$$

Then, the algorithm uses the importance index to select the most important tokens. After the tokens are selected, subsequence BoostTransformer creates a new sample x_i^t at iterate t , which contains only the important tokens, and is used by the weak learner. The modified minimization problem and the boosting weak learner are explicitly presented in (8) and (9), respectively. The proposed algorithm (subsequence BoostTransformer) is summarized in Algorithm 2.

Different from standard BoostTransformer, subsequence BoostTransformer first reviews the whole dataset in steps 3-4 and generates the basic weak learner g_0^* . Once the basic weak learner is created, in each iteration, subsequence BoostTransformer first updates the attention-based importance vector I_ω for any $\omega \in V_{t-1}$ in step 7, and selects σ fraction of the tokens to form the vocabulary set V_t , and lastly constructs a new sample x_i^t by deleting any tokens not in V_t in step 8. After the new sample x_i^t is constructed, subsequence BoostTransformer initializes the weights of the current transformer by using the weights in g_{t-1}^* and trains the transformer with x_i^t to minimize the squared error in (8) in steps 9-10. Lastly, the algorithm finds the boosting coefficient α_t by minimizing (6) in step 11 and adds the additive weak learner to the ensemble in step 12.

Algorithm 2 subsequence BoostTransformer

- 1: **Inputs:**
number of classes M , number of boosting iterations N_b , shrinkage parameter ν , dataset $\mathcal{D} = \{(x_1, z_1), \dots, (x_n, z_n)\}$ where $z_i \in \{1, \dots, M\}$ is the label of sample x_i , and $0 < \sigma < 1$
 - 2: **Initialize:**
set $f(x) = \mathbf{0} \in \mathbb{R}^M$, $V_0 = \{\omega | \omega \in x_i \text{ for some } x_i\}$
 - 3: compute $w(x_i, z_i)$ for all (x_i, z_i) , using (3)
 - 4: train a transformer g_0^* to optimize (5)
 - 5: $f(x) = g_0^*$
 - 6: **for** $t = 1, 2, \dots, N_b$ **do**
 - 7: update importance values I_ω for $\omega \in V_{t-1}$, using (10), (11) and (12)
 - 8: form $V_t \subset V_{t-1}$ with $\frac{|V_t|}{|V_{t-1}|} \approx \sigma$ and $I(\omega) > I(\omega') \forall \omega \in V_t, \omega' \in V_{t-1} \setminus V_t$ and form a new sample x_i^t for each sample i
 - 9: compute $w(x_i, z_i)$ for all i , using (3) and (9)
 - 10: train a transformer g_t^* to optimize (8)
 - 11: find the optimal coefficient α_t , using (6) and (9)
 - 12: $f(x) = f(x) + \nu \alpha_t g_t^*$
 - 13: **end for**
-

4.3 Importance-sampling-based BoostTransformer

Importance sampling, a strategy for preferential sampling of more important samples capable of accelerating the training process, has been well studied in stochastic gradient descent (SGD) [1]. However, there is virtually no existing work combining the power of importance sampling with the strength of boosting. Motivated by the phenomenon that overfitting appears early in standard BoostTransformer, we propose importance-sampling-based BoostTransformer, which combines importance sampling and BoostTransformer. Importance-sampling-based BoostTransformer mimics importance sampling SGD by introducing a new loss function and computing a probability distribution for drawing samples. Similarly, importance-sampling-based BoostTransformer computes a probability distribution in each iteration, and draws a subset of samples to train the weak learner based on the distribution. The probability distribution is

$$P(I = i) = \frac{\|w(x_i, z_i)\|}{\sum_{(x_j, z_j) \in \mathcal{D}} \|w(x_j, z_j)\|}, \quad (13)$$

which yields the new loss function for a subset of samples \mathcal{G} to be

$$\bar{\mathcal{L}}_{\mathcal{G}}(w, g) = \sum_{(x_i, z_i) \in \mathcal{G}} \frac{1}{|\mathcal{D}| P(I = i)} \|g(x_i) - w(x_i, z_i)\|^2. \quad (14)$$

To any minimization algorithm one would typically use. We then apply any optimization algorithm with respect to (14) (by further using mini-batches or importance sampling).

The entire algorithm is exhibited in Algorithm 3.

Algorithm 3 importance-sampling-based BoostTransformer

- 1: **Inputs:**
number of classes M , number of boosting iterations N_b , shrinkage parameter ν , dataset $\mathcal{D} = \{(x_1, z_1), \dots, (x_n, z_n)\}$ where $z_i \in \{1, \dots, M\}$ is the label of sample x_i , and $0 < \sigma < 1$
 - 2: **Initialize:**
set $f(x) = \mathbf{0} \in \mathbb{R}^M$
 - 3: compute $w(x_i, z_i)$ for all x_i , using (3)
 - 4: train a transformer g_0^* to optimize (5)
 - 5: $f(x) = g_0^*$
 - 6: **for** $t = 1, 2, \dots, N_b$ **do**
 - 7: compute probability distribution P_t , using (13)
 - 8: draw independently $|\mathcal{G}^t|$ samples, which is σ fraction of the samples, based on P_t
 - 9: compute $w(x_i, z_i)$ for $(x_i, z_i) \in \mathcal{G}^t$, using (3)
 - 10: train a transformer g_t^* to optimize (14) on \mathcal{G}^t
 - 11: find the optimal coefficient α_t , using (6) on \mathcal{G}^t
 - 12: $f(x) = f(x) + \nu\alpha_t g_t^*$
 - 13: **end for**
-

Importance-sampling-based BoostTransformer starts with learning the full-size dataset and training a basic weak learner in steps 3-4. In each iteration, the algorithm first computes the probability distribution P_t in step 7 and selects a subset \mathcal{G}^t of samples based on the distribution in step 8. Once the dataset is created, it computes the weights and trains a transformer by using the unbiased loss function (14), following by finding an optimal boosting coefficient in steps 9-12.

In the rest of this section, we provide all analysis of the optimal probability distribution in importance-sampling-based BoostTransformer. Given current aggregated classifier $f_{t-1} = \sum_{i=1}^{t-1} g_i^*$, let us define the expected training progress attributable to iteration t as

$$\mathbb{E}_{P_t} [\Delta^{(t)}] = \|f_{t-1} - f^*\|^2 - \mathbb{E}_{P_t} [\|f_t - f^*\|^2 | \mathcal{F}^{t-1}].$$

Here f^* denotes the solution to (1), and the expectation is taken over the probability distribution P_t , and \mathcal{F}^{t-1} contains the whole history of the algorithm up until iterate $t - 1$. We assume that gradient sampling is unbiased. Inspired by the work in [45], we prove that the optimal probability distribution is proportional to the boosting weight at each iteration.

Theorem 1. *In $\max_{P_t} \mathbb{E}_{P_t} [\Delta^{(t)}]$, the optimal distribution for importance-sampling-based BoostTransformer to select each sample i is proportional to its “boosting weight norm:”, i.e. (13).*

Proof. See Appendix A. □

Based on the fact that the new loss function with respect to the probability distribution is unbiased, we discover that maximizing the improvement of the boosting algorithm is equivalent to minimizing the functional gradient variance. By applying Jensen’s inequality, the optimal probability distribution is essentially proportional to the boosting weights, which are easy to obtain, in boosting algorithms.

5 Experimental Study

In this section, we first compare subgrid BoostCNN with standard BoostCNN and deep CNNs, next, we compare the standard transformer, BoostTransformer, subsequence BoostTransformer and importance-sampling-based BoostTransformer in the second half of the section. We conduct experiments on three different datasets for both CNN related and transformer related algorithms. From all of these datasets, we study the performance of the boosting technique, the subgrid trick and the importance sampling strategy. All the algorithms are implemented in Python with PyTorch [37]. Training is conducted on an NVIDIA Titan XP GPU.

5.1 Image

In this subsection, we illustrate properties of the proposed subgrid BoostCNN and compare its performance with other methods on several image classification tasks. In subgrid BoostCNN, the risk function (1) we employ is cross entropy, and the input of each weak learner is an image with 3 channels which can be handled by standard Conv2d functions in PyTorch. Meanwhile, we implement the subgrid strategy based on (7) with respect to each pixel (j, k) . We delete approximately 10% of the rows and columns, which implies $\sigma = 81\%$ on the total number of pixels, and fix the shrinkage parameter ν to be 0.02. In each weak learner, we apply the ADAM algorithm with the learning rate of 0.0001 and weight decay being 0.0001.

We consider CIFAR-10 [23], SVHN [36] and ImageNetSub [11] datasets as shown in Table 1. For the last dataset, since the original ImageNet dataset is large and takes significant amount of time to train, we select a subset of samples from the original ImageNet dataset. More precisely, we randomly pick 100 labels and select the corresponding samples from ImageNet, which consists of 124,000 images for training and 10,000 images for testing. We denote it as ImageNetSub. Data preprocessing consists of three steps: 1. random resizing and cropping with output size 224×224 , scale uniformly sampled from $[0.08, 1.0]$ and make the aspect ratio uniformly sampled from $[0.75, 1.33]$; 2. random horizontal flipping with flipping probability 0.5; 3. normalization for each channel.

	Number of Training/Testing Samples	Number of Classes
CIFAR-10	50k/10k	10
SVHN	73k/26k	10
ImageNetSub	124k/10k	100

Table 1: Image Datasets

For training, we employ three different deep CNNs, which are ResNet-18, ResNet-50 and ResNet-101. For each combination of dataset/CNN, we first train the deep CNN for a certain number of epochs, and then initialize the weights in the basic weak learner for the boosting algorithms as the weights in the deep CNN. In the subgrid BoostCNN experiments, we use 10 CNN weak learners. We train each weak learner for 15 epochs. For comparison, we train BoostCNN, the ensemble method (without boosting weight update and always using all features) denoted by e-CNN and the subgrid ensemble method named as subgrid e-CNN

(without boosting weight update in step 10 in Algorithm 1) for 10 iterates as well. Notice that subgrid e-CNN essentially mimics random forests. We also train the single deep CNN for 150 epochs to represent approximately the same computational effort as training 10 CNN weak learners for 15 epochs.

We start by applying ResNet-18 as our weak learner for all different ensemble methods. Figures 1, 3 and 5 compare the relative performances with respect to single ResNet-18 vs the running time. The solid lines in green and yellow show the relative performances of BoostCNN and subgrid BoostCNN, respectively, while the dotted lines in green and yellow represent the relative performances of e-CNN and subgrid e-CNN, respectively. As shown in these figures, taking the same amount of time, subgrid BoostCNN outperforms all of the remaining algorithms. Furthermore, we observe that subgrid BoostCNN outperforms BoostCNN, and subgrid e-CNN has the same behavior when compared with e-CNN. In conclusion, the subgrid technique improves the performance of the boosting algorithm. Moreover, Figures 2, 4 and 6 depict subgrid BoostCNN and subgrid e-CNN using three different seeds with respect to their averages. The solid and dotted lines in the same color represent the same seed used in corresponding subgrid BoostCNN and subgrid e-CNN. As the figures show, the solid lines are closer to each other than the dotted lines, which indicates that subgrid BoostCNN is more robust with respect to the variation of the seed when compared with subgrid e-CNN. Furthermore, the standard deviations of the accuracy generated by subgrid e-CNN and subgrid BoostCNN are shown in Table 2. The standard deviations of the accuracy generated by subgrid e-CNN are significant compared to those of subgrid BoostCNN, which in turn indicates that subgrid BoostCNN is less sensitive to the choice of the seed. Therefore, subgrid BoostCNN is more robust than subgrid e-CNN.

	subgrid BoostCNN	subgrid e-CNN
CIFAR-10	0.478	2.519
SVHN	0.385	0.891
ImageNetSub	2.489	7.915

Table 2: Standard deviation times 10^3 of the accuracy results by different seeds

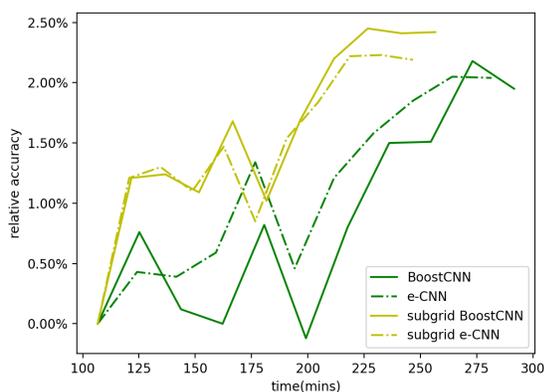


Figure 1: ResNet-18 on CIFAR-10

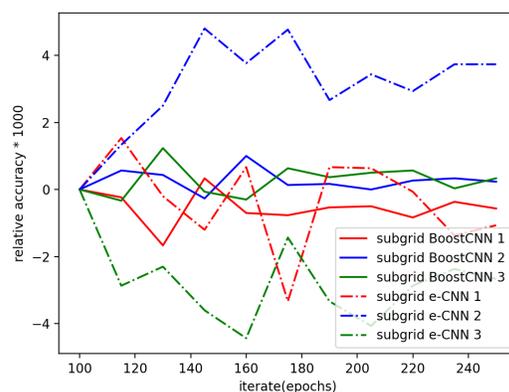


Figure 2: Different Seeds

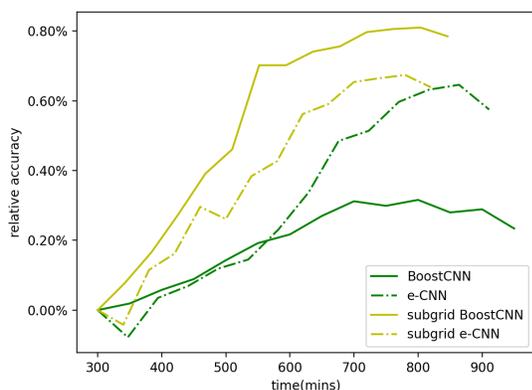


Figure 3: ResNet-18 on SVHN

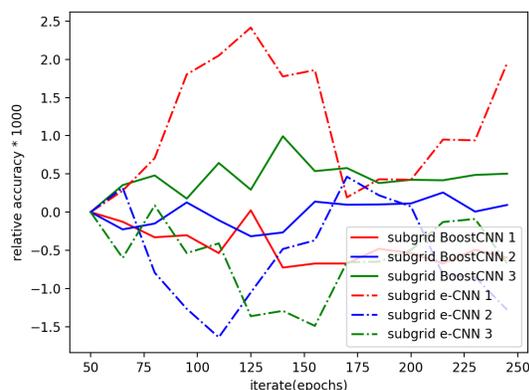


Figure 4: Different Seeds

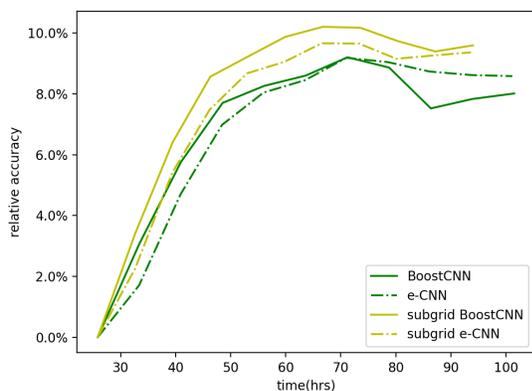


Figure 5: ResNet-18 on ImageNet-Sub

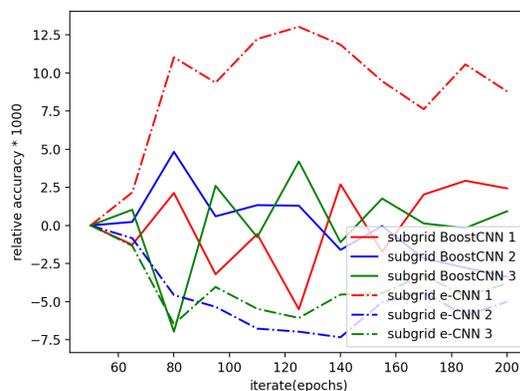


Figure 6: Different Seeds

Next, we evaluate relative performances of subgrid BoostCNN using ResNet-50 as the weak learner on CIFAR-10 and ImageNetSub datasets with respect to the single ResNet-50. We do not evaluate the relative performances on the SVHN dataset since the accuracy

of the single ResNet-50 on the SVHN dataset is over 98%. From Figures 7 and 9, we also observe the benefits of the subgrid technique. Besides, Figures 8 and 10 confirm that subgrid BoostCNN is more stable than subgrid e-CNN since the solid series are closer to each other compared with the dotted series. Furthermore, we establish the relative performances of subgrid BoostCNN using ResNet-50 as the weak learner with respect to the single ResNet-101 in Figure 11. Although single ResNet-101 outperforms single ResNet-50, subgrid BoostCNN using ResNet-50 as the weak learner outperforms single ResNet-101 significantly in Figure 11, which indicates that subgrid BoostCNN with a simpler CNN is able to exhibit a better performance than a single deeper CNN. Lastly, we conduct experiments with ResNet-101 on the ImageNetSub dataset. From Figure 12, we not only discover the superior behaviors of BoostCNN, e-CNN, subgrid BoostCNN and subgrid e-CNN over ResNet-101 as we expect, but also observe the benefit of the subgrid technique.

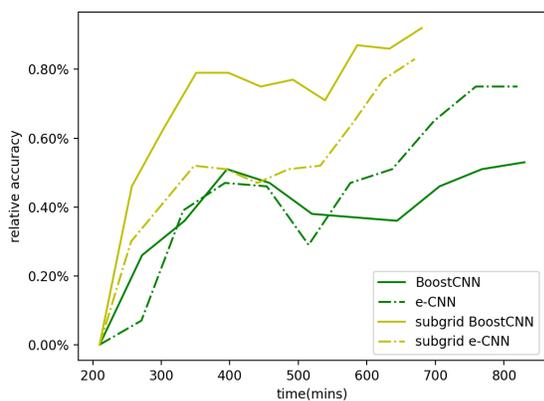


Figure 7: ResNet-50 on CIRFAR-10

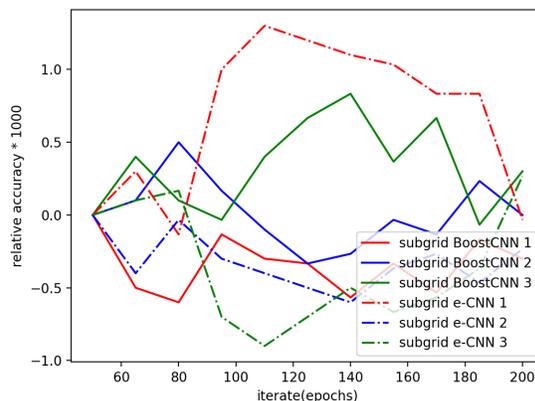


Figure 8: Different Seeds

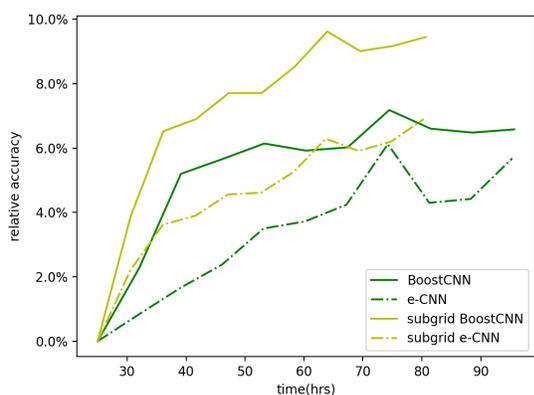


Figure 9: ResNet-50 on ImageNet-Sub

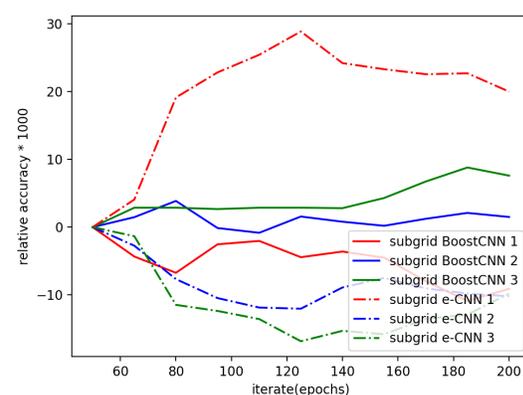


Figure 10: Different Seeds

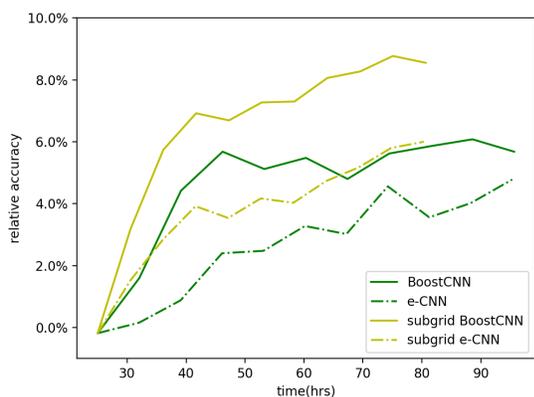


Figure 11: ResNet-50 on ImageNetSub compared to ResNet-101

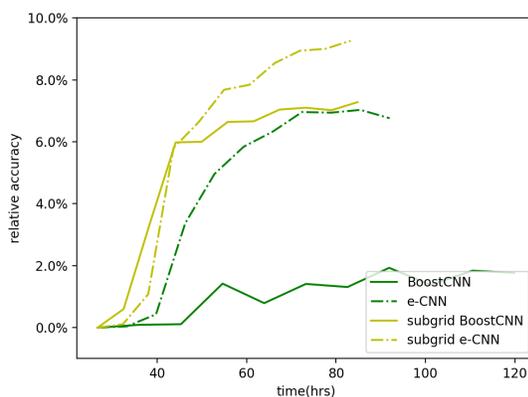


Figure 12: ResNet-101 on ImageNetSub

5.2 Text

In this section, we explore properties of the proposed Boost Transformer, subsequence Boost Transformer and importance-sampling-based Boost Transformer, and compare their performances with other methods on several text classification tasks. In the following experiments, the weak learner used is RoBERTa-based [28] from the HuggingFace library with only word embeddings to be pre-trained weights. Using transformer based boosting algorithms, we train an ensemble of 6 transformers each with 5 epochs (these numbers yield good performance). In subsequence BoostTransformer, we pick the most important 80% of the tokens in the vocabulary and reconstruct the dataset based on this new vocabulary. In importance-sampling-based BoostTransformer, the first flavor, in each iteration, we select 80% of the samples based on the probability distribution in (13) without further subsequence technique. In subsequence importance-sampling-based BoostTransformer, we first select 80% of the samples based on the probability distribution in (13), and then pick the most important 80% of the tokens in the current vocabulary given by the selected 80% samples, after that, we reconstruct the dataset based on this modified vocabulary. For comparison, we train the vanilla transformer and subsequence transformer, which randomly removes 20% of the tokens and trains the network on the dataset for 30 epochs. To train the model, we use AdamW [29] with learning rate 10^{-5} , weight decay 0.01 and batch size 16. We use linear learning rate decay with warmup ratio 0.06.

We start by presenting the three public datasets used: IMDB [30], Yelp polarity reviews and Amazon polarity reviews [31]. The IMDB dataset, which is for binary sentiment classification, contains a set of 25,000 highly polar movie reviews for training, and 25,000 for testing. The Yelp polarity reviews dataset, which is a subset of the dataset obtained from the Yelp Dataset Challenge in 2015, consists of 100,000 training samples and 38,000 testing samples. The classification task for this dataset is predicting a polarity label by considering stars 1 and 2 negative, and 3 and 4 positive for each review text. The last dataset we use is the Amazon polarity reviews dataset, which is a subset of the original Amazon reviews dataset from the Stanford Network Analysis Project (SNAP). Dealing with the same classification task as the Yelp polarity review dataset, the Amazon polarity reviews dataset contains 100,000 training samples and 25,000 testing samples. The subsampled datasets are standard, i.e. we did not create our own subsamples. Empirically we found that a weak learner with 6 heads and 6 layers achieves good robust performance.

Given the architecture of the weak learner, we start by discussing experiments on IMDB. In Figure 13, we compare the relative performances of the algorithms with respect to the vanilla transformer. As it shows, all versions of BoostTransformer do not perform as good as the standard transformer and subsequence transformer in the first few epochs. However, they catch up quickly and dominate the performance in the remaining training epochs. Even more, based on Figure 14, which represents each model’s relative improvement with respect to its initial weights, all versions of BoostTransformer maintain their performances as the number of epochs increases, while the performances of the standard transformer and the subsequence transformer start decreasing and fluctuating dramatically after the first few epochs, which implies that all versions of BoostTransformer are more robust than the standard and subsequence transformer.

Next, we evaluate relative performances with respect to the vanilla transformer on the

Yelp and Amazon polarity review datasets. From Figures 15-18, we discover that the superior and more robust behavior of boosting algorithms over transformer is vigorous.

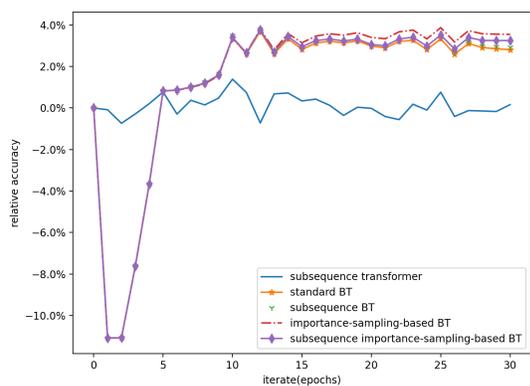


Figure 13: Relative Accuracy on IMDB

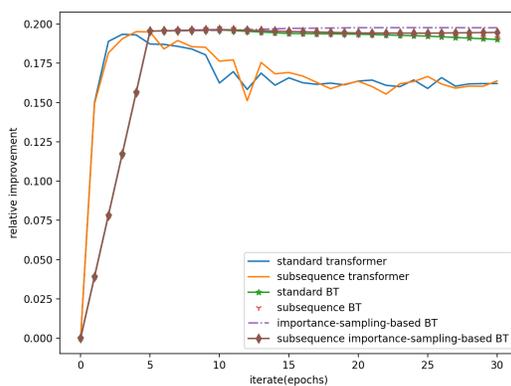


Figure 14: Improvement on IMDB

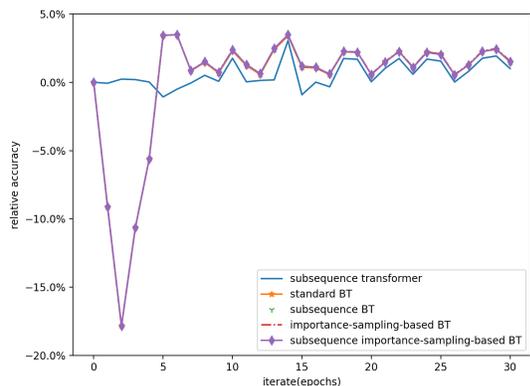


Figure 15: Relative Accuracy on Yelp

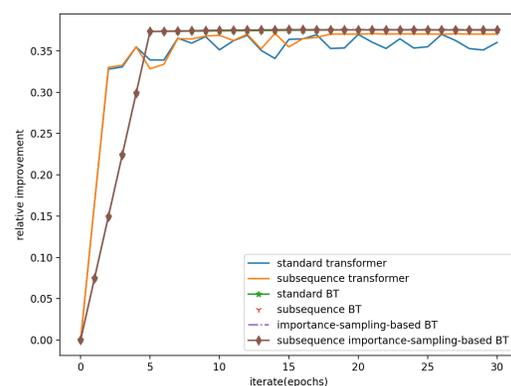


Figure 16: Improvement on Yelp

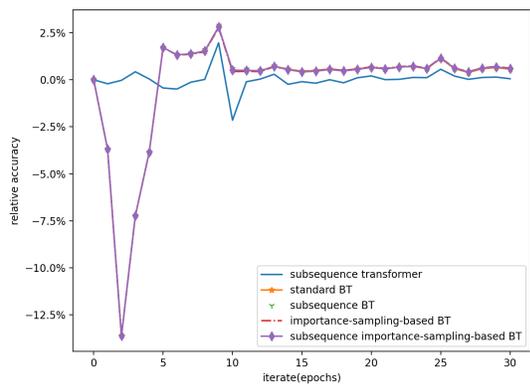


Figure 17: Relative Accuracy on Amazon

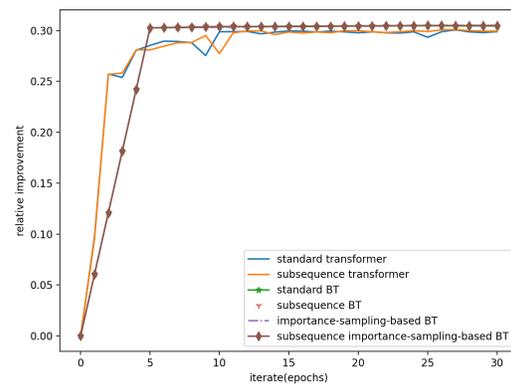


Figure 18: Improvement on Amazon

Furthermore, we zoom in on the performances at iterates larger than 2. In Figures 19-21,

compared with the standard BoostTransformer, we observe that the subsequence BoostTransformer, importance-sampling-based BoostTransformer and subsequence importance-sampling-based BoostTransformer demonstrate a superior performance. Therefore, we conclude that the subsequence and importance sampling techniques are beneficial for the boosting algorithms. Moreover, we observe that the importance-sampling-based BoostTransformer gradually improves its performance and maintains its performance later on, while the subsequence BoostTransformer hits its best accuracy in early epochs and then starts fluctuating and decaying. The gap between the importance-sampling-based BoostTransformer and the subsequence BoostTransformer is more significant on the IMDB dataset, which has a much smaller size than the Yelp and Amazon polarity review datasets. For the subsequence importance-sampling-based BoostTransformer, compared to the subsequence BoostTransformer, although the subsequence importance-sampling-based BoostTransformer does not fluctuate and decrease as much as the subsequence BoostTransformer, which is more obvious in a small dataset (i.e. the IMDB dataset), its best accuracy is lower than that of the subsequence BoostTransformer, which is more obvious in larger datasets (i.e. the Yelp and Amazon datasets). On the other hand, compared to the importance-sampling-based BoostTransformer, although the subsequence importance-sampling-based BoostTransformer obtains its best accuracy earlier than the importance-sampling-based BoostTransformer, its overall performance fluctuates while the importance-sampling-based BoostTransformer keeps increasing and maintains its high-quality performance in all of the datasets, which implies that the subsequence importance-sampling-based BoostTransformer is less stable than the importance-sampling-based BoostTransformer. In conclusion, the subsequence BoostTransformer fits well for datasets with enough samples and the importance-sampling-based BoostTransformer is more suitable for datasets with a limited number of samples.

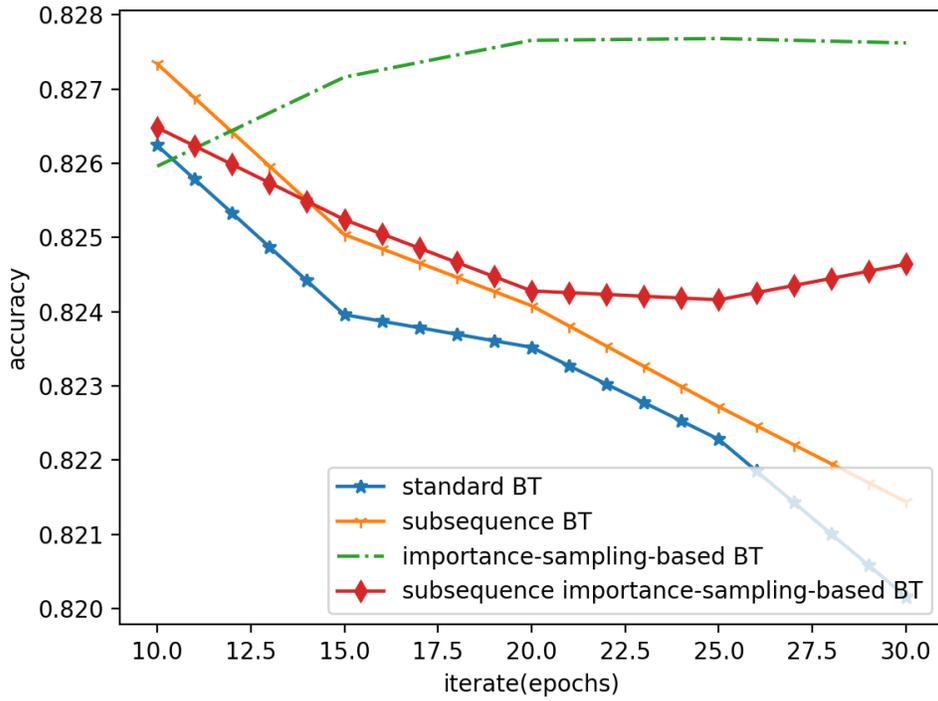


Figure 19: IMDB

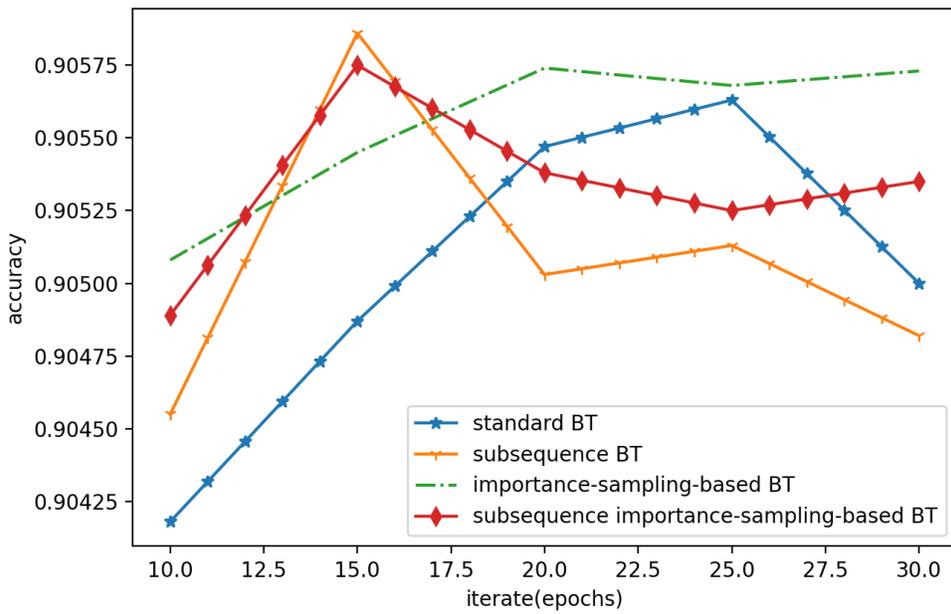


Figure 20: Yelp

Table 3 illustrates the running time of each algorithm on the different datasets in minutes.

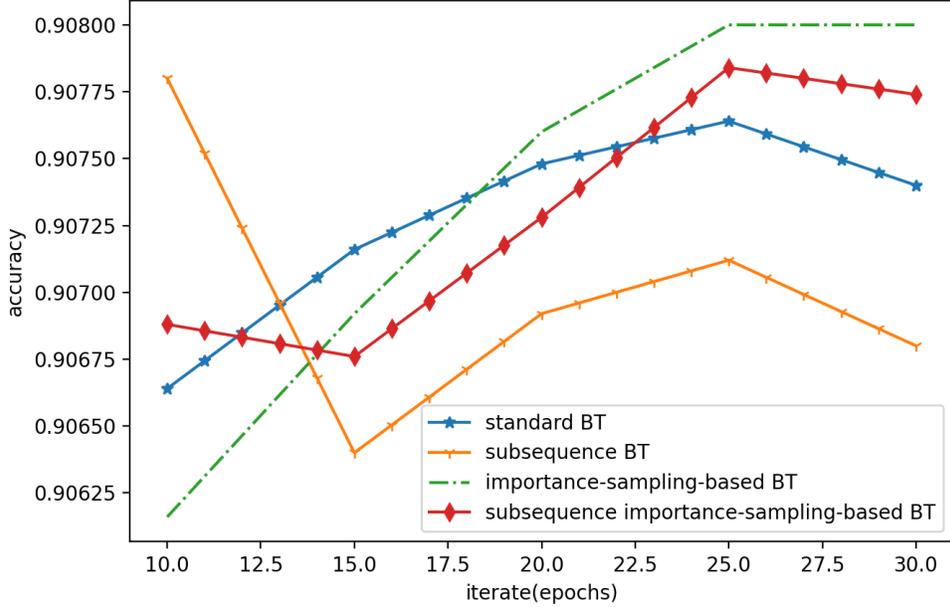


Figure 21: Amazon

As we see in the table, the subsequence technique not only improves the performance of the boosting algorithms but also reduces the running time. Furthermore, the importance sampling technique reduces the running time significantly without hurting the performance.

	Trans.	subsequence Trans.	BT	subsequence BT	importance-sampling-based BT	subsequence importance-sampling-based BT
IMDB	21	14	23	17	13	12
Yelp	76	52	84	66	52	46
Amazon	75	53	79	58	46	41

Table 3: Running time for different algorithms

In conclusion, a subsequence transformer is a good choice if the running time cost is the most important concern, however, if accuracy performance is as crucial as the running time, then the subsequence BoostTransformer is the go-to option since it requires a slight increase in the running time but provides superior and more robust performance when compared to the subsequence transformer. In addition, if a dataset has a limited number of samples, i.e., it is easy to cause overfitting, then the importance-sampling-based BoostTransformer can outperform.

References

- [1] Guillaume Alain, Alex Lamb, Chinnadhurai Sankar, Aaron C. Courville, and Yoshua Bengio. Variance reduction in SGD by distributed importance sampling. *ArXiv*, abs/1511.06481, 2015.
- [2] Mohammad Assaad, Romuald Boné, and Hubert Cardot. A new boosting algorithm for improved time-series forecasting with recurrent neural networks. *Information Fusion*, 9:41–55, 2008.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473, 2015.
- [4] Lorenz Berger, Eoin Hyde, Matt Gibb, Nevil Pavithran, Garin Kelly, Faiz Mumtaz, and Sébastien Ourselin. Boosted training of convolutional neural networks for multi-class segmentation. *ArXiv*, abs/1806.05974, 2018.
- [5] Sourour Brahimi, Najib Ben Aoun, and Chokri Ben Amar. Boosted convolutional neural networks. In *BMVC*, 2016.
- [6] Sourour Brahimi, Najib Ben Aoun, and Chokri Ben Amar. Boosted convolutional neural network for object recognition at large scale. *Neurocomputing*, 330:337–354, 2019.
- [7] Hugh Chen, Scott Lundberg, and Su-In Lee. Hybrid gradient boosting trees and neural networks for forecasting operating room data. *ArXiv*, abs/1801.07384, 2018.
- [8] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN Encoder-Decoder for statistical machine translation. In *EMNLP*, 2014.
- [9] Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. What does BERT look at? An analysis of BERT’s attention. *ArXiv*, abs/1906.04341, 2019.
- [10] Dominik Csiba and Peter Richtárik. Importance sampling for minibatches. *ArXiv*, abs/1602.02283, 2018.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- [13] Mattia A Di Gangi, Matteo Negri, and Marco Turchi. Adapting transformer to end-to-end spoken language translation. In *INTERSPEECH*, 2019.
- [14] Linhao Dong, Shuang Xu, and Bo Xu. Speech-Transformer: A no-recurrence sequence-to-sequence model for speech recognition. In *ICASSP*, 2018.

- [15] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [16] Jiuxiang Gu, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu, Xingxing Wang, Gang Wang, and Jianfei Cai. Recent advances in convolutional neural networks. *Pattern Recognition*, 77:354–377, 2018.
- [17] Shizhong Han, Zibo Meng, Ahmed-Shehab Khan, and Yan Tong. Incremental boosting convolutional neural network for facial action unit recognition. In *NIPS*, 2016.
- [18] Trevor Hastie, Saharon Rosset, Ji Zhu, and Hui Zou. Multi-class AdaBoost. *Statistics and Its Interface*, 2:349–360, 2009.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [20] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. AMC: AutoML for model compression and acceleration on mobile devices. In *ECCV*, 2018.
- [21] Forrest Iandola, Matt Moskewicz, Sergey Karayev, Ross Girshick, Trevor Darrell, and Kurt Keutzer. DenseNet: Implementing efficient ConvNet descriptor pyramids. *ArXiv*, abs/1404.1869, 2014.
- [22] Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. *ArXiv*, abs/1803.00942, 2018.
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Citeseer*, 2009.
- [24] Shin-Jye Lee, Tonglin Chen, Lun Yu, and Chin-Hui Lai. Image classification based on the boost convolutional neural network. *IEEE Access*, 6:12755–12768, 2018.
- [25] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu. Neural speech synthesis with transformer network. In *AAAI*, 2019.
- [26] Jason Liang, Elliot Meyerson, Babak Hodjat, Dan Fink, Karl Mutch, and Risto Miikkilainen. Evolutionary neural AutoML for deep learning. In *GECCO*, 2019.
- [27] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear CNN models for fine-grained visual recognition. In *ICCV*, 2015.
- [28] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- [30] Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *ACL*, 2011.

- [31] Julian McAuley and Jure Leskovec. Hidden factors and hidden topics: understanding rating dimensions with review text. In *RecSys*, 2013.
- [32] Indraneel Mukherjee and Robert E Schapire. A theory of multiclass boosting. *Journal of Machine Learning Research*, 14:437–497, 2013.
- [33] Tomohiro Nakatani. Improving transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration. In *INTER-SPEECH*, 2019.
- [34] Alexey Natekin and Alois Knoll. Gradient boosting machines, a tutorial. *Frontiers in Neurorobotics*, 7, 2013.
- [35] Deanna Needell, Rachel Ward, and Nati Srebro. Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm. *Mathematical Programming*, 155:549–573, 2014.
- [36] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [37] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [38] Ross J. Quinlan. Induction of decision trees. *Machine Learning*, 1:81–106, 2004.
- [39] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:1137–1149, 2015.
- [40] Alex rizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *CACM*, 2017.
- [41] Mohammad J Saberian and Nuno Vasconcelos. Multiclass Boosting: Theory and algorithms. In *NIPS*, 2011.
- [42] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. *ArXiv*, abs/1409.3215, 2014.
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, A. Gomez, L. Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.
- [44] Ching-Feng Yeh, Jay Mahadeokar, Kaustubh Kalgaonkar, Yongqiang Wang, Duc Le, Mahaveer Jain, Kjell Schubert, Christian Fuegen, and Michael L Seltzer. Transformer-Transducer: End-to-end speech recognition with self-attention. *ArXiv*, abs/1910.12977, 2019.
- [45] Peilin Zhao and Tong Zhang. Stochastic optimization with importance sampling for regularized loss minimization. In *ICML*, 2015.

- [46] Shiyu Zhou, Linhao Dong, Shuang Xu, and Bo Xu. Syllable-based sequence-to-sequence speech recognition with the transformer in mandarin chinese. *ArXiv*, abs/1804.10752, 2018.

6 Appendix

A Proof of Theorem 1

Proof. Given a probability distribution P for dataset (x_i, z_i) , by assumption, the stochastic gradient of the loss function is unbiased, i.e.

$$\mathbb{E}_P \left[\frac{\partial \bar{L}_i(z_i, f(x_i))}{\partial f} \right] = \frac{\partial \mathcal{R}}{\partial f}, \quad (15)$$

with

$$\mathcal{R}[f] = \frac{1}{|\mathcal{D}|} \sum_{(x_i, z_i) \in \mathcal{D}} L(z_i, f(x_i)) = \mathbb{E}_P [\bar{L}_i(z_i, f(x_i))]$$

and

$$\bar{L}_i(z_i, f(x_i)) = \frac{1}{|\mathcal{D}| P(I=i)} L(z_i, f(x_i)).$$

At iterate t , in importance-sampling-based Boosting algorithms, given probability distribution P_t and $P_t^i = P_t(I=i)$, the current gradient given a subset \mathcal{G}^t of samples is

$$\begin{aligned} \bar{g}_t^{g^t} &= \frac{1}{|\mathcal{G}^t|} \sum_{(x_i, z_i) \in \mathcal{G}^t} \left. \frac{\partial \bar{L}_i(z_i, f_{t-1}(x_i) + \epsilon g(x_i))}{\partial g} \right|_{\epsilon=0} \\ &= \frac{1}{|\mathcal{G}^t|} \sum_{(x_i, z_i) \in \mathcal{G}^t} \left. \frac{1}{|\mathcal{D}| P_t^i} \frac{\partial L_i(z_i, f_{t-1}(x_i) + \epsilon g(x_i))}{\partial g} \right|_{\epsilon=0} \\ &= \frac{1}{|\mathcal{G}^t|} \sum_{(x_i, z_i) \in \mathcal{G}^t} \frac{1}{|\mathcal{D}| P_t^i} g_t^i = \frac{1}{|\mathcal{G}^t|} \sum_{k=1}^{|\mathcal{G}^t|} G_k, \end{aligned} \quad (16)$$

where $g_t^i = \left. \frac{\partial L_i(z_i, f_{t-1}(x_i) + \epsilon g(x_i))}{\partial g} \right|_{\epsilon=0}$ and G_k is the random variable corresponding to sample k . Note that

$$g_t = \left. \frac{\partial \mathcal{R}[f_{t-1}; g]}{\partial g} \right|_{\epsilon=0} = \left. \frac{\partial \mathcal{R}[f_{t-1} + \epsilon g]}{\partial g} \right|_{\epsilon=0} \quad (17)$$

and

$$\mathbb{E}_{P_t}(G_t) = \mathbb{E}_{P_t} \left[\bar{g}_t^{g^t} \right] = g_t, \quad (18)$$

due to the unbiased gradient in (15). Given $\bar{g}_t^{g^t}$ computed on a subset \mathcal{G}^t with probability distribution P_t , we consider

$$\begin{aligned} \mathbb{E}_{P_t} [\Delta^{(t)}] &= \|f_{t-1} - f^*\|^2 - \mathbb{E}_{P_t} [\|f_t - f^*\|^2 | \mathcal{F}^{t-1}] \\ &= \|f_{t-1} - f^*\|^2 - \mathbb{E}_{P_t} \left[\left\| f_{t-1} + \alpha_t \bar{g}_t^{g^t} - f^* \right\|^2 \middle| \mathcal{F}^{t-1} \right] \\ &= -2\alpha_t \left\langle f_{t-1} - f^*, \mathbb{E}_{P_t} \left[\bar{g}_t^{g^t} \middle| \mathcal{F}^{t-1} \right] \right\rangle - \alpha_t^2 \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g^t} \right\|^2 \middle| \mathcal{F}^{t-1} \right]. \end{aligned} \quad (19)$$

By inserting (18) into (19), we have

$$\mathbb{E}_{P_t} [\Delta^{(t)}] = -2\alpha_t \langle f_{t-1} - f^*, g_t \rangle - \alpha_t^2 \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} \right\|^2 \middle| \mathcal{F}^{t-1} \right]. \quad (20)$$

Thus, maximizing $\mathbb{E}_{P_t} [\Delta^{(t)}]$ is equivalent to minimizing the variance of the gradient, i.e. $\mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} \right\|^2 \middle| \mathcal{F}^{t-1} \right]$. Consequently, consider

$$\begin{aligned} \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} \right\|^2 \middle| \mathcal{F}^{t-1} \right] &= \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} - g_t + g_t \right\|^2 \middle| \mathcal{F}^{t-1} \right] \\ &= \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} - g_t \right\|^2 \middle| \mathcal{F}^{t-1} \right] + \mathbb{E}_{P_t} \left[2 \langle \bar{g}_t^{g_t} - g_t, g_t \rangle \middle| \mathcal{F}^{t-1} \right] + \|g_t\|^2 \\ &= \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} - g_t \right\|^2 \middle| \mathcal{F}^{t-1} \right] + 2 \langle \mathbb{E}_{P_t} [\bar{g}_t^{g_t} \middle| \mathcal{F}^{t-1}] - g_t, g_t \rangle + \|g_t\|^2 \\ &= \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} - g_t \right\|^2 \middle| \mathcal{F}^{t-1} \right] + \|g_t\|^2, \end{aligned} \quad (21)$$

where the last equality holds due to (18). Continuing, we have

$$\begin{aligned} \mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} - g_t \right\|^2 \middle| \mathcal{F}^{t-1} \right] &= \mathbb{E}_{P_t} \left[\left\| \frac{1}{|g_t|} \sum_{i \in g_t} \left(\frac{1}{|\mathcal{D}| P_t^i} g_t^i - g_t \right) \right\|^2 \middle| \mathcal{F}^{t-1} \right] \\ &= \frac{1}{|g_t|^2} \mathbb{E}_{P_t} \left[\left\| \sum_{i \in g_t} \left(\frac{1}{|\mathcal{D}| P_t^i} g_t^i - g_t \right) \right\|^2 \middle| \mathcal{F}^{t-1} \right] \\ &= \frac{1}{|g_t|^2} \mathbb{E}_{P_t} \left[\sum_{i \in g_t} \left\| \frac{1}{|\mathcal{D}| P_t^i} g_t^i - g_t \right\|^2 + \sum_{(i,j) \in g_t, i \neq j} \left\langle \frac{1}{|\mathcal{D}| P_t^i} g_t^i - g_t, \frac{1}{|\mathcal{D}| P_t^j} g_t^j - g_t \right\rangle \middle| \mathcal{F}^{t-1} \right] \\ &= \frac{1}{|g_t|^2} \left(\mathbb{E}_{P_t} \left[\sum_{i \in g_t} \left\| \frac{1}{|\mathcal{D}| P_t^i} g_t^i - g_t \right\|^2 \middle| \mathcal{F}^{t-1} \right] + \mathbb{E}_{P_t} \left[\sum_{(i,j) \in g_t, i \neq j} \left\langle \frac{1}{|\mathcal{D}| P_t^i} g_t^i - g_t, \frac{1}{|\mathcal{D}| P_t^j} g_t^j - g_t \right\rangle \middle| \mathcal{F}^{t-1} \right] \right) \\ &= \frac{|g_t|}{|g_t|^2} \mathbb{E}_{P_t} [\|G_1 - g_t\|^2 \middle| \mathcal{F}^{t-1}] + \frac{2}{|g_t|^2} \binom{|g_t|}{2} \langle \mathbb{E}_{P_t} [G_1 - g_t \middle| \mathcal{F}^{t-1}], \mathbb{E}_{P_t} [G_2 - g_t \middle| \mathcal{F}^{t-1}] \rangle \\ &= \frac{1}{|g_t|} \mathbb{E}_{P_t} [\|G_1 - g_t\|^2 \middle| \mathcal{F}^{t-1}] \\ &= \frac{1}{|g_t|} (\mathbb{E}_{P_t} [\|G_1\|^2 \middle| \mathcal{F}^{t-1}] - \|g_t\|^2). \end{aligned} \quad (22)$$

The fifth equality holds since G_i and G_j are independent, moreover, the seventh equality is valid due to (18). Inserting (22) into (21) yields

$$\mathbb{E}_{P_t} \left[\left\| \bar{g}_t^{g_t} \right\|^2 \middle| \mathcal{F}^{t-1} \right] = \frac{1}{|g_t| |\mathcal{D}|^2} \sum_{(x_i, z_i) \in \mathcal{D}} \frac{1}{P_t^i} \|g_t^i\|^2 - \frac{1}{|g_t|} \|g_t\|^2 + \|g_t\|^2. \quad (23)$$

As (23) shows, maximizing $\mathbb{E}_{P_t} [\Delta^{(t)}]$ is equivalent to minimizing $\frac{1}{P_t^i} \|g_t^i\|^2$. By using the Jensen's inequality, it follows that

$$\sum_{(x_i, z_i) \in \mathcal{D}} \frac{1}{P_t^i} \|g_t^i\|^2 = \sum_{(x_i, z_i) \in \mathcal{D}} P_t^i \left(\frac{\|g_t^i\|}{P_t^i} \right)^2 \geq \left(\sum_{(x_i, z_i) \in \mathcal{D}} \|g_t^i\| \right)^2, \quad (24)$$

and the equality holds when $P_t^i = \|g_t^i\| / \sum_{(x_j, z_j) \in \mathcal{D}} \|g_t^j\|$. Note that $g_t^i = \left. \frac{\partial L(z_i, f_{t-1}(x_i) + \epsilon g(x_i))}{\partial g} \right|_{\epsilon=0}$ is proportional to the boosting weights $w_t(x_i, z_i)$ of sample (x_i, z_i) as stated in (4), therefore, the claim in (13) follows. \square