

# Appendices

## A Proof of Theorem 1

Consider the similar strongly convex conditions used in Pegasos [27], where  $F(\mathbf{w})$  satisfies the following conditions:

**Property 1.**  $F(\mathbf{w}) := \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{w})$  satisfies:

- $F$  is  $\mu$ -strongly convex,
- $\arg \min_{\mathbf{w}} F(\mathbf{w}) \in B_D = \{\mathbf{w} \mid \|\mathbf{w}\| \leq D\}$ , and
- $\|\nabla f_i(\mathbf{w})\| \leq G, \forall \mathbf{w} \in B_D$

With the step size as  $\eta_k = \frac{1}{\mu k}$ , shrinking threshold  $T_k = \frac{G}{k}$  and initialization  $w^1 \in B_D$ , we can prove Theorem 1 with above assumption.

[Proof of Theorem 1] From strong convexity, we have:

$$\begin{aligned} F(\mathbf{w}^*) - F(\mathbf{w}^k) &\geq \langle \nabla F(\mathbf{w}^k), \mathbf{w}^* - \mathbf{w}^k \rangle + \frac{\mu}{2} \|\mathbf{w}^k - \mathbf{w}^*\|^2 \\ F(\mathbf{w}^k) - F(\mathbf{w}^*) &\geq \langle \nabla F(\mathbf{w}^*), \mathbf{w}^k - \mathbf{w}^* \rangle + \frac{\mu}{2} \|\mathbf{w}^k - \mathbf{w}^*\|^2. \end{aligned}$$

Adding the above inequalities gives:

$$\begin{aligned} \langle \nabla F(\mathbf{w}^k) - \nabla F(\mathbf{w}^*), \mathbf{w}^k - \mathbf{w}^* \rangle &\geq \mu \|\mathbf{w}^k - \mathbf{w}^*\|^2 \\ \Rightarrow \langle \nabla F(\mathbf{w}^k), \mathbf{w}^k - \mathbf{w}^* \rangle &\geq \mu \|\mathbf{w}^k - \mathbf{w}^*\|^2 \quad (\text{since } \nabla F(\mathbf{w}^*) = 0) \\ \Rightarrow \langle \mathbb{E}(\nabla f_i(\mathbf{w}^k)), \mathbf{w}^k - \mathbf{w}^* \rangle &\geq \mu \|\mathbf{w}^k - \mathbf{w}^*\|^2 \end{aligned} \quad (7)$$

Where  $i$  is the data index sampled at step  $k$ , the update step is defined as:

$$\mathbf{w}^{k+1} = \mathbf{w}^k - \eta_k \mathbf{g}_i^k \quad (8)$$

where

$$\mathbf{g}_i^k = \nabla f_i(\mathbf{w}^k) - \epsilon_i^k \quad (9)$$

$$\epsilon_i^k = \nabla f_i(\mathbf{w}^k) \mathbf{I} \left( \|\nabla f_i(\mathbf{w}^k)\| \leq \frac{G}{k} \right) \quad (10)$$

Next, we have:

$$\begin{aligned} \mathbb{E}(\|\mathbf{w}^{k+1} - \mathbf{w}^*\|^2) &= \mathbb{E}(\|\mathbf{w}^k - \eta_k \mathbf{g}_i^k - \mathbf{w}^*\|^2) \\ &= \mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|^2) - 2\eta_k \mathbb{E} \langle \mathbf{g}_i^k, \mathbf{w}^k - \mathbf{w}^* \rangle + \eta_k^2 \mathbb{E}(\|\mathbf{g}_i^k\|^2) \\ &\leq \mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|^2) - 2\eta_k \mathbb{E} \langle \nabla f_i(\mathbf{w}^k), \mathbf{w}^k - \mathbf{w}^* \rangle \\ &\quad + \eta_k^2 G^2 + 2\eta_k \mathbb{E} \langle \epsilon_i^k, \mathbf{w}^k - \mathbf{w}^* \rangle \end{aligned} \quad (11)$$

The last term from (11) can be upper bounded as follows:

$$\begin{aligned} &2\eta_k \mathbb{E} \langle \epsilon_i^k, \mathbf{w}^k - \mathbf{w}^* \rangle \\ &= 2\eta_k \mathbb{E} \left( \langle \nabla f_i(\mathbf{w}^k), \mathbf{w}^k - \mathbf{w}^* \rangle \mathbf{I} \left( \|\nabla f_i(\mathbf{w}^k)\| \leq \frac{G}{k} \right) \right) \\ &\leq 2\eta_k \mathbb{E} \left( \|\nabla f_i(\mathbf{w}^k)\| \|\mathbf{w}^k - \mathbf{w}^*\| \mathbf{I} \left( \|\nabla f_i(\mathbf{w}^k)\| \leq \frac{G}{k} \right) \right) \\ &\leq 2\eta_k \frac{G}{k} \mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|) \end{aligned} \quad (12)$$

Substituting (7) and (12) into (11), we have:

$$\mathbb{E}(\|\mathbf{w}^{k+1} - \mathbf{w}^*\|^2) \leq \mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|^2) \left( 1 - \frac{2}{k} \right) + \frac{G^2}{\mu^2 k^2} + \frac{2G}{\mu k^2} \mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|) \quad (13)$$

Letting  $\hat{D} = \max\left(D, \frac{2G}{\mu}\right)$  and letting  $L = 4\hat{D}^2$ , we have:

$$L = 4\hat{D}^2 \geq \frac{G^2}{\mu^2} + \frac{4\hat{D}G}{\mu} \quad (14)$$

The convergence can be established by induction. We want to proof the following condition holds for  $k \geq 1$ :

$$\mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|^2) \leq \frac{L}{k} \quad (15)$$

When  $k = 1$ , the inequality hold:

$$\mathbb{E}(\|\mathbf{w}^1 - \mathbf{w}^*\|^2) \leq \frac{L}{1} \quad (16)$$

Suppose the same holds for  $k$ . Then for  $k + 1$ , using (13) we have:

$$\mathbb{E}(\|\mathbf{w}^{k+1} - \mathbf{w}^*\|^2) \leq \left(1 - \frac{2}{k}\right)\mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|^2) + \frac{G^2}{\mu^2 k^2} + \frac{2G}{\mu k^2}\mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|) \quad (17)$$

Since  $\mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|) \leq \sqrt{\mathbb{E}(\|\mathbf{w}^k - \mathbf{w}^*\|^2)} \leq \sqrt{L} \leq 2\hat{D}$ , we have:

$$\mathbb{E}(\|\mathbf{w}^{k+1} - \mathbf{w}^*\|^2) \leq \left(1 - \frac{2}{k}\right)\frac{L}{k} + L\frac{1}{k^2} \quad (18)$$

$$= \frac{k-1}{k^2}L \leq \frac{L}{k+1} \quad (19)$$

Thus according to the principle of mathematical induction, (15) hold for all  $k \geq 1$  ■

In many machine learning models, the calculation of the objective function value takes many fewer operations than the computation of norm of the gradient. Also, for  $\mu$ -strongly convex  $F(\mathbf{w})$  which is  $M$ -Lipschitz smooth we have:

$$\mu(F(\mathbf{w}) - F(\mathbf{w}^*)) \leq \frac{1}{2}\|\nabla F(\mathbf{w})\|^2 \leq \frac{M^2}{\mu}(F(\mathbf{w}) - F(\mathbf{w}^*)),$$

where the left inequality follows from the Polyak-Lojasiewicz inequality, showing that the gradient norm can be bounded by the loss function value. Due to these reasons, in practice we can also use the loss function value as the utility of shrinking.

## B Algorithms for CPU/GPU parallelization

---

### Algorithm 2 Assistant (CPU)

---

- **Input:** Training dataset  $D = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$ , BossQueue min size  $c$
  - **Initialize:** BossQueue, AssistantQueue, Assistant
  - **While True:**
    - If BossQueue.size() <  $c$ :
      - \*  $\mathbf{B} = \text{Assistant.sample\_batch}()$
      - \* BossQueue.enqueue( $\mathbf{B}$ )
    - Else if not AssistantQueue.empty():
      - \*  $\mathbf{M} = \text{AssistantQueue.pop}()$
      - \* grad = Assistant.gradient( $\mathbf{M}$ )
      - \* Assistant.update(grad)
- 

## C Discussion on Computational Overhead of Importance Sampling

The case of importance sampling in Figure 1 gives an example that too much overhead could ruin the acceleration effect. The importance sampling algorithm uses a precise estimation of instance utility (normalized gradient norm) which contains both local and global information. However, as the model changes after every parameter update, importance weights need to be updated, which

---

**Algorithm 3** Boss (GPU)

---

- **Input:**
  - **Initialize:** Boss
  - **While True:**
    - If not BossQueue.empty():
      - \* B = BossQueue.pop()
      - \* M = Boss.Forward(B)
      - \* AssistantQueue.enqueue(M)
      - \* grad = Boss.Backward(M)
      - \* Boss.update(grad)
- 

Table 1: Image Classification datasets statistics

Dataset	# instances	image size	# classes
MNIST	60,000	$28 \times 28$	10
Rotated MNIST	60000	$28 \times 28$	10
CIFAR10	60,000	$32 \times 32 \times 3$	10
ImageNet	1,281,167	$224 \times 224 \times 3$	1,000

introduces large computational overhead. A pre-sampling technique is sometimes used to tackle this issue. First a subset of data  $C \subset [N]$  (with  $|C| \ll N$ ) is uniformly sampled and importance scores are evaluated only on  $C$ . After training enough number of batches on  $C$ , another chunk is sampled and evaluated. This can reduce the computational overhead but may introduce new issues. Firstly, the importance weights are fixed once evaluated and may be outdated after parameter updates. In real applications with large data, the model can evolve substantially even within one epoch through the data. Secondly, substantial computational overhead is introduced even with the sub-sampling technique. This indicates that any acceleration method will fail on SGD if it requires access to global information.

## D Experimental parameter setting and further results

**Experimental settings for image classification.** For the image classification experiments reported in Figure 3, we set the learning rate for MNIST, rotated MNIST and raw ImageNet to be 0.0005, 0.005 for CIFAR10. The batch size is chosen to be 16 for MNIST/rotated MNIST and CIFAR10 datasets and 64 for raw ImageNet. As raw ImageNet dataset cannot be load into memory at once, to reduce the image preprocessing cost, we preprocessed all the images and pre-stored them into 12 data chunks. During training, the model loads and trains on one chunk at a time. For SPL algorithm, we implemented the same self-paced scheme in [22] with  $q(t) = 4$  and  $\text{maxgen} = 20$ .

**Experimental setting for NMT task.** We train WMT14 dataset on 8 Tesla V100 GPUs with around 40k tokens per training batch (5k tokens per batch per GPU). Each model is trained until 6 billion tokens are seen. The BLEU scores are evaluated with the averaged model of last 10 checkpoints (epochs). For the Transformer big model, to allocate memory for 3000 tokens per GPU, we use half precision (16 bit) floating point type.

Table 2: Statistics of the Transformers on WMT14 dataset. The model used for computing BLEU score is the averaged model of last 10 epochs.

model	# tokens per epoch	time per epoch(s)	final BLEU score
<b>SGD(base)</b>	112.1M	1115	27.25
<b>SPL(base)</b>	112.1M	1142	27.32
<b>AutoAssist (base)</b>	72.6M	973	<b>27.42</b>
<b>SGD(big)</b>	112.1M	1190	27.41
<b>SPL(big)</b>	112.1M	1221	27.40
<b>AutoAssist (big)</b>	70.8M	1016	<b>27.63</b>




Figure 5: Comparison of various SGD acceleration approaches on image classification.  $x$ -axis is the training time in seconds, while  $y$  axis is the test accuracy. ResNets with varied depth ranging from  $\{18, 34, 101\}$  are considered. Each block column of figures show the results of the ResNet with a specific number of layers. Each block row of images show the results on a specific dataset.




Figure 6: The Assistant predictive accuracy and safeguarding rate  $(1 - \gamma)$ . Even with the simplest logistic regression model, the Assistant is able to reach up to 90% accuracy while predicting the shrinking label, resulting in a confident Assistant with low safe-guarding rate  $(1 - \gamma)$ .