# Appendix to "Adam with Bandit Sampling for Deep Learning" 

Rui Liu, Tianyi Wu, Barzan Mozafari
Computer Science and Engineering
University of Michigan, Ann Arbor
\{ruixliu, tianyiwu, mozafari\}@umich.edu

## A Proof of Theorem 1

According to Theorem 4.1 in [1], the convergence rate of Adam is

$$
\begin{align*}
& \sum_{t=1}^{T}\left[f_{t}\left(\theta_{t}\right)-f_{t}\left(\theta^{*}\right)\right] \\
\leq & \frac{D^{2}}{2 \alpha\left(1-\beta_{1}\right)} \sum_{i=1}^{d} \sqrt{T \hat{v}_{T, i}}+ \\
& \frac{\alpha\left(1+\beta_{1}\right) G_{\infty}}{\left(1-\beta_{1}\right) \sqrt{1-\beta_{2}}(1-\gamma)^{2}} \sum_{i=1}^{d}\left\|g_{1: T, i}\right\|_{2}+  \tag{1}\\
& \sum_{i=1}^{d} \frac{D_{\infty}^{2} G_{\infty} \sqrt{1-\beta_{2}}}{2 \alpha\left(1-\beta_{1}\right)(1-\lambda)^{2}}
\end{align*}
$$

First, we show $\sum_{i=1}^{d} \sqrt{\hat{v}_{T, i}} \leq d G_{\infty}$
Proof.

$$
\begin{align*}
& \sum_{i=1}^{d} \sqrt{\hat{v}_{T, i}} \\
= & \sum_{i=1}^{d} \sqrt{\sum_{j=1}^{T} \frac{1-\beta_{2}}{1-\beta_{2}^{T}} \beta_{2}^{T-j} g_{j, i}^{2}} \\
\leq & \sum_{i=1}^{d} \sqrt{\sum_{j=1}^{T} \frac{1-\beta_{2}}{1-\beta_{2}^{T}} \beta_{2}^{T-j} G_{\infty}^{2}}  \tag{2}\\
= & G_{\infty} \sqrt{\frac{1-\beta_{2}}{1-\beta_{2}^{T}} \sum_{i=1}^{d} \sqrt{\sum_{j=1}^{T} \beta_{2}^{T-j}}} \\
= & G_{\infty} \sqrt{\frac{1-\beta_{2}}{1-\beta_{2}^{T}} \sum_{i=1}^{d} \sqrt{\frac{1-\beta_{2}^{T}}{1-\beta_{2}}}} \\
= & d G_{\infty}
\end{align*}
$$

Therefore, the above bound can be rewritten as

$$
\begin{align*}
& \sum_{t=1}^{T}\left[f_{t}\left(\theta_{t}\right)-f_{t}\left(\theta^{*}\right)\right] \\
\leq & \rho_{1} d \sqrt{T}+\sqrt{d} \rho_{2} \sqrt{\sum_{t=1}^{T}\left\|\hat{G}_{t}\right\|^{2}}+\rho_{3} \tag{3}
\end{align*}
$$

where

$$
\begin{align*}
\rho_{1} & =\frac{D^{2} G_{\infty}}{2 \alpha\left(1-\beta_{1}\right)} \\
\rho_{2} & =\frac{\alpha\left(1+\beta_{1}\right) G_{\infty}}{\left(1-\beta_{1}\right) \sqrt{1-\beta_{2}}(1-\gamma)^{2}}  \tag{4}\\
\rho_{3} & =\sum_{i=1}^{d} \frac{D_{\infty}^{2} G_{\infty} \sqrt{1-\beta_{2}}}{2 \alpha\left(1-\beta_{1}\right)(1-\gamma)^{2}}
\end{align*}
$$

Proof.

$$
\begin{align*}
& \sum_{t=1}^{T}\left[f_{t}\left(\theta_{t}\right)-f_{t}\left(\theta^{*}\right)\right] \\
\leq & \rho_{1} d \sqrt{T}+\rho_{2} \sum_{i=1}^{d}\left\|g_{1: T, i}\right\|_{2}+\rho_{3} \\
= & \rho_{1} d \sqrt{T}+d \rho_{2} \sum_{i=1}^{d} \frac{1}{d} \sqrt{\sum_{t=1}^{T}\left(\hat{G}_{t, i}\right)^{2}}+\rho_{3}  \tag{5}\\
& (\text { due to } \sqrt{\cdot} \text { is concave) } \\
\leq & \rho_{1} d \sqrt{T}+d \rho_{2} \sqrt{\sum_{i=1}^{d} \frac{1}{d} \sum_{t=1}^{T}\left(\hat{G}_{t, i}\right)^{2}}+\rho_{3} \\
= & \rho_{1} d \sqrt{T}+\sqrt{d} \rho_{2} \sqrt{\sum_{t=1}^{T}\left\|\hat{G}_{t}\right\|^{2}}+\rho_{3}
\end{align*}
$$

we could get the following theorem for the convergence of Adam with sampling with replacement with batch size $K$.

$$
\begin{align*}
& \sum_{t=1}^{T}\left[f_{t}\left(\theta_{t}\right)-f_{t}\left(\theta^{*}\right)\right] \\
= & \rho_{1} d \sqrt{T}+\sqrt{d} \rho_{2} \sqrt{\frac{1}{n^{2} K} \sum_{t=1}^{T} \mathbb{E}\left[\sum_{j=1}^{n} \frac{\left\|g_{t}\right\|^{2}}{p_{j}^{t}}\right]}+\rho_{3} \tag{6}
\end{align*}
$$

Proof. Since $\hat{G}_{t}=\frac{1}{K} \sum_{k=1}^{K} \hat{g}_{I_{k}^{t}}=\frac{1}{K} \sum_{k=1}^{K} \frac{g_{I_{k}^{t}}}{n p_{I_{k}^{t}}}$, we have

$$
\begin{align*}
\mathbb{E}\left[\left\|\hat{G}_{t}\right\|^{2}\right] & \leq \frac{1}{n^{2} K^{2}} \sum_{k=1}^{K} \mathbb{E}\left[\frac{\left\|g_{I_{k}^{t}}\right\|^{2}}{p_{I_{k}^{t}}^{2}}\right] \\
& =\frac{1}{n^{2} K^{2}} \sum_{k=1}^{K} \mathbb{E}\left[\sum_{j=1}^{n} \frac{\left\|g_{j}^{t}\right\|^{2}}{\left(p_{j}^{t}\right)^{2}} p_{j}^{t}\right] \\
& =\frac{1}{n^{2} K^{2}} \sum_{k=1}^{K} \mathbb{E}\left[\sum_{j=1}^{n} \frac{\left\|g_{j}^{t}\right\|^{2}}{p_{j}^{t}}\right]  \tag{7}\\
& =\frac{1}{n^{2} K} \mathbb{E}\left[\sum_{j=1}^{n} \frac{\left\|g_{j}^{t}\right\|^{2}}{p_{j}^{t}}\right]
\end{align*}
$$

From previous step, we know

$$
\begin{align*}
& \sum_{t=1}^{T}\left[f_{t}\left(\theta_{t}\right)-f_{t}\left(\theta^{*}\right)\right] \\
\leq & \rho_{1} d \sqrt{T}+\sqrt{d} \rho_{2} \sqrt{\sum_{t=1}^{T}\left\|\hat{G}_{t}\right\|^{2}}+\rho_{3} \\
\leq & \rho_{1} d \sqrt{T}+\sqrt{d} \rho_{2} \sqrt{\sum_{t=1}^{T} \frac{1}{n^{2} K} \mathbb{E}\left[\sum_{j=1}^{n} \frac{\left\|g_{j}^{t}\right\|^{2}}{p_{j}^{t}}\right]}+\rho_{3}  \tag{8}\\
= & \rho_{1} d \sqrt{T}+\sqrt{d} \rho_{2} \sqrt{\frac{1}{n^{2} K} \sum_{t=1}^{T} \mathbb{E}\left[\sum_{j=1}^{n} \frac{\left\|g_{j}^{t}\right\|^{2}}{p_{j}^{t}}\right]}+\rho_{3}
\end{align*}
$$

## B Proof of Theorem 2

Theorem 2 follows from Theorem 1 and Lemma 1. Therefore, we focus on the proof of Lemma 1 here. We prove Lemma 1 using the framework of online learning with bandit feedback.
Online optimization is interested in choosing $p^{t}$ to solve the following problem

$$
\begin{equation*}
\min _{p^{t} \in \mathcal{P}, 1 \leq t \leq T} \sum_{t=1}^{T} L_{t}\left(p^{t}\right) \tag{9}
\end{equation*}
$$

where $L_{t}\left(p^{t}\right)$ is the loss that incurs at each iteration. Equivalently, the goal is the same as minimizing the pseudo-regret:

$$
\begin{equation*}
\bar{R}_{T}=\mathbb{E} \sum_{t=1}^{T} L_{t}\left(p^{t}\right)-\min _{p \in \mathcal{P}} \mathbb{E} \sum_{t=1}^{T} L_{t}(p) \tag{10}
\end{equation*}
$$

The following Algorithm 1 similar to EXP3 could be used to solve the above problem.
To be clear, the Bregman divergence $B_{\phi}(x, y)=\phi(x)-\phi(y)-\langle\nabla \phi(y), x-y\rangle$. Note that, the updating step of Algorithm 1 is equivalent to

$$
\begin{align*}
w^{t+1} & =\nabla \phi^{*}\left(\nabla \phi\left(p^{t}\right)-\alpha_{p} \hat{h}_{t}\right) \\
p^{t+1} & =\underset{y \in \mathcal{P}}{\operatorname{argmin}} B_{\phi}\left(y, w^{t+1}\right) \tag{11}
\end{align*}
$$

We have the following convergence result for Algorithm 1 .

```
Algorithm 1 Online optimization
    Input: stepsize \(\alpha_{p}\)
    Initialize \(x^{1}\) and \(p^{1}\).
    for \(t=1, \cdots, T\) do
        Play a perturbation \(\tilde{p}^{t}\) of \(p^{t}\) and observe \(\nabla L_{t}\left(\tilde{p}^{t}\right)\)
        Compute an unbiased gradient estimate \(\hat{h}_{t}\) of \(\nabla L_{t}\left(p^{t}\right)\), i.e. as long as \(\mathbb{E}\left[\hat{h}_{t}\right]=\nabla L_{t}\left(p^{t}\right)\)
        Update \(p^{t}: p^{t+1} \leftarrow \operatorname{argmin}_{p \in \mathcal{P}}\left\{\left\langle\hat{h}_{t}, p\right\rangle+\frac{1}{\alpha_{p}} B_{\phi}\left(p, p^{t}\right)\right\}\).
    end for
```

Proposition B.1. The Algorithm 1 has the following convergence result

$$
\begin{align*}
\bar{R}_{T}= & \mathbb{E} \sum_{t=1}^{T} L_{t}\left(\tilde{p}^{t}\right)-\min _{p \in \mathcal{P}} \mathbb{E} \sum_{t=1}^{T} L_{t}(p) \\
\leq & \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{1}{\alpha_{p}} \sum_{t=1}^{T} \mathbb{E} B_{\phi^{*}}\left(\nabla \phi\left(p^{t}\right)-\alpha_{p} \hat{h}_{t}, \nabla \phi\left(p^{t}\right)\right)  \tag{12}\\
& +\sum_{t=1}^{T} \mathbb{E}\left[\left\|p_{t}-\tilde{p}_{t}\right\|\left\|\hat{h}_{t}\right\|_{*}\right]
\end{align*}
$$

If $L_{t}\left(p^{t}\right)$ is linear (i.e. $L_{t}\left(p^{t}\right)=\left\langle l_{t}, p^{t}\right\rangle$ ), then we have

$$
\begin{align*}
\bar{R}_{T}= & \mathbb{E} \sum_{t=1}^{T} L_{t}\left(\tilde{p}^{t}\right)-\min _{p \in \mathcal{P}} \mathbb{E} \sum_{t=1}^{T} L_{t}(p) \\
\leq & \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{1}{\alpha_{p}} \sum_{t=1}^{T} \mathbb{E} B_{\phi^{*}}\left(\nabla \phi\left(p^{t}\right)-\alpha_{p} \hat{h}_{t}, \nabla \phi\left(p^{t}\right)\right)  \tag{13}\\
& +\sum_{t=1}^{T} \mathbb{E}\left[\left\|p_{t}-\mathbb{E}\left[\tilde{p}_{t} \mid p_{t}\right]\right\|\left\|\hat{h}_{t}\right\|_{*}\right]
\end{align*}
$$

For example, let $\phi(p)=\sum_{j=1}^{n} p_{j} \log p_{j}$. Then its convex conjugate is $\phi^{*}(u)=\sum_{j=1}^{n} \exp \left(u_{j}-1\right)$. In this case, due to Equation 11, the updating step of Algorithm 1 becomes

$$
\begin{equation*}
w_{j}^{t+1}=p_{j}^{t} \exp \left(-\alpha_{p} \hat{h}_{t, j}\right), \forall 1 \leq j \leq n \tag{14}
\end{equation*}
$$

The convergence result can also be simplified because

$$
\begin{align*}
& B_{\phi^{*}}\left(\nabla \phi\left(p^{t}\right)-\alpha_{p} \hat{h}_{t}, \nabla \phi\left(p^{t}\right)\right) \\
= & \sum_{j=1}^{n} p_{j}^{t}\left(\exp \left(-\alpha_{p} \hat{h}_{t, j}\right)+\alpha_{p} \hat{h}_{t, j}-1\right) \tag{15}
\end{align*}
$$

$$
\text { (assume } \hat{h}_{t, j} \geq 0 \text { and due to } e^{z}-z-1 \leq z^{2} / 2 \text { for } z \leq 0 \text { ) }
$$

$$
\leq \frac{\alpha_{p}^{2}}{2} \sum_{j=1}^{n} p_{j}^{t} \hat{h}_{t, j}^{2}
$$

Linear Case: Let's consider a special case where $L_{t}\left(p^{t}\right)$ is linear, i.e. $L_{t}\left(p^{t}\right)=\left\langle l_{t}, p^{t}\right\rangle$. Assume $p^{t}$ is a probability distribution, i.e. $\mathcal{P}=\left\{p: \sum_{i=1}^{n} p_{i}=1,0 \leq p_{i} \leq 1, \forall i\right\}$. In this case, $\nabla L_{t}\left(p^{t}\right)=l_{t}$. At iteration $t$, assume that we can't get the whole vector of $l_{t}$. Instead, we can get only one coordinate $l_{t, J_{t}}$, where $J_{t}$ is sampled according to the distribution $p^{t}$. This is equivalent to $\tilde{p}_{j}^{t}=\mathbb{1}\left(j=J_{t}\right)$. Obviously,

$$
\begin{equation*}
\mathbb{E}\left[\tilde{p}_{j}^{t}\right]=\mathbb{E}\left[\mathbb{1}\left(j=J_{t}\right)\right]=\sum_{i=1}^{n} p_{i}^{t} \mathbb{1}(j=i)=p_{j}^{t} \tag{16}
\end{equation*}
$$

Furthermore, to get the unbiased estimate of gradient, we set $\hat{h}_{t, j}=\frac{l_{t, j} \mathbb{1}\left(j=J_{t}\right)}{p_{j}^{t}}$. To verify $\hat{h}_{t, j}$ is indeed unbiased w.r.t. $l_{t, j}$,

$$
\begin{equation*}
\mathbb{E}\left[\hat{h}_{t, j}\right]=\mathbb{E}\left[\frac{l_{t, j} \mathbb{1}\left(j=J_{t}\right)}{p_{j}^{t}}\right]=\frac{l_{t, j} \mathbb{E}\left[\mathbb{1}\left(j=J_{t}\right)\right]}{p_{j}^{t}}=\frac{l_{t, j}}{p_{j}^{t}} p_{j}^{t}=l_{t, j} \tag{17}
\end{equation*}
$$

For sampling multiple actions $I^{t}$, we have the following

$$
\begin{equation*}
\hat{h}_{t, j}=\frac{l_{t, j} \sum_{k=1}^{K} \mathbb{1}\left(j=I_{k}^{t}\right)}{K p_{j}^{t}} \tag{18}
\end{equation*}
$$

And its convergence result is

$$
\begin{align*}
& \bar{R}_{T}= \mathbb{E} \sum_{t=1}^{T} L_{t}\left(\tilde{p}^{t}\right)-\min _{p \in \mathcal{P}} \mathbb{E} \sum_{t=1}^{T} L_{t}(p) \\
& \leq \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{1}{\alpha_{p}} \sum_{t=1}^{T} \mathbb{E} B_{\phi^{*}}\left(\nabla \phi\left(p^{t}\right)-\alpha_{p} \hat{h}_{t}, \nabla \phi\left(p^{t}\right)\right) \\
&+\sum_{t=1}^{T} \mathbb{E}\left[\left\|p_{t}-\mathbb{E}\left[\tilde{p}_{t} \mid p_{t}\right]\right\|\left\|\hat{h}_{t}\right\|_{*}\right] \\
&= \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{1}{\alpha_{p}} \sum_{t=1}^{T} \mathbb{E} B_{\phi^{*}}\left(\nabla \phi\left(p^{t}\right)-\alpha_{p} \hat{h}_{t}, \nabla \phi\left(p^{t}\right)\right) \\
& \leq \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \sum_{t=1}^{T} \mathbb{E} \sum_{j=1}^{n} p_{j}^{t} \hat{h}_{t, j}^{2}  \tag{19}\\
&= \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \sum_{t=1}^{T} \mathbb{E} \sum_{j=1}^{n} p_{j}^{t} l_{t, j}^{2}\left(\sum_{k=1}^{K} \mathbb{1}\left(j=I_{k}^{t}\right)\right)^{2} \\
&\left(K p_{j}^{t}\right)^{2} \\
& \leq \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \sum_{t=1}^{T} \mathbb{E} \sum_{j=1}^{n} p_{j}^{t} \frac{l_{t, j}^{2} \sum_{k=1}^{K}\left(\mathbb{1}\left(j=I_{k}^{t}\right)\right)^{2}}{K\left(p_{j}^{t}\right)^{2}} \\
&= \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \sum_{t=1}^{T} \mathbb{E} \frac{l_{t, J_{t}}^{2}}{p_{J_{t}}^{t}} \\
&= \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \sum_{t=1}^{T} \sum_{j=1}^{n} l_{t, j}^{2}
\end{align*}
$$

We want to apply online optimization/bandit to learn $p_{j}^{t}$, in which case the loss at $t$-th iteration is $l_{t}\left(p^{t}\right)=\sum_{j=1}^{n} \frac{\left\|g_{j}\left(x^{t}\right)\right\|_{*}^{2}}{p_{j}^{t}}$. Because the loss $l_{t}\left(p^{t}\right)$ is a nonlinear function, we could use nonlinear bandit as a general approach to solve this. However, another simpler approach is to convert it to linear problem, where linear bandit could be used. Specifically, we have

$$
\begin{align*}
& \sum_{t=1}^{T} \mathbb{E}\left[\sum_{j=1}^{n}\left\|g_{j}\left(x^{t}\right)\right\|_{*}^{2}\left(\frac{1}{p_{j}^{t}}-\frac{1}{p_{j}}\right)\right] \\
\leq & \sum_{t=1}^{T} \mathbb{E}\left[\left\langle-\left\{\left\|g_{j}\left(x^{t}\right)\right\|_{j, *}^{2} /\left(p_{j}^{t}\right)^{2}\right\}_{j=1}^{n}, p^{t}-p\right\rangle\right]  \tag{20}\\
= & \sum_{t=1}^{T} \mathbb{E}\left[\left\langle-\left\{\left\|g_{j}\left(x^{t}\right)\right\|_{j, *}^{2} /\left(p_{j}^{t}\right)^{2}+\frac{L^{2}}{p_{\min }^{2}}\right\}_{j=1}^{n}, p^{t}-p\right\rangle\right]
\end{align*}
$$

This is equivalent to online linear optimization setting where $l_{t, j}(x)=-\frac{\left\|g_{j}(x)\right\|_{j, *}^{2}}{\left(p_{j}^{t}\right)^{2}}+\frac{L^{2}}{p_{\text {min }}^{2}}$. And it's easy to see that $0 \leq l_{t, j} \leq \frac{L^{2}}{p_{\text {min }}^{2}}$. Then the convergence result is

$$
\begin{align*}
& \mathbb{E} \sum_{t=1}^{T} L_{t}\left(\tilde{p}^{t}\right)-\min _{p \in \mathcal{P}} \mathbb{E} \sum_{t=1}^{T} L_{t}(p) \\
& =\frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \sum_{t=1}^{T} \sum_{j=1}^{n} l_{t, j}^{2} \\
& \leq \frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \sum_{t=1}^{T} \sum_{j=1}^{n} \frac{L^{4}}{p_{\min }^{4}} \\
& =\frac{B_{\phi}\left(p, p^{1}\right)}{\alpha_{p}}+\frac{\alpha_{p}}{2} \frac{n T L^{4}}{p_{\min }^{4}}  \tag{21}\\
& \leq \frac{R^{2}}{\alpha_{p}}+\frac{\alpha_{p}}{2} \frac{n T L^{4}}{p_{\min }^{4}} \\
& \left(\operatorname{set} \alpha_{p}=\sqrt{\left.\frac{2 R^{2} p_{\min }^{4}}{n T L^{4}}\right)}\right. \\
& =\frac{R L^{2}}{p_{\min }^{2}} \sqrt{2 n T}
\end{align*}
$$

## C Proofs about Comparison with Uniform Sampling

## C. 1 Proof of Lemma 2

Proof.

$$
\begin{align*}
& \mathbb{E} \sum_{j=1}^{n} \frac{\left\|g_{j}\right\|^{2}}{p_{j}} \\
= & \mathbb{E} n \sum_{j=1}^{n}\left\|g_{j}\right\|^{2} \\
= & n \sum_{j=1}^{n} \mathbb{E}\left\|g_{j}\right\|^{2} \\
\leq & n \sum_{j=1}^{n} \mathbb{E} \sum_{i=1}^{d} z_{j, i}^{2}  \tag{22}\\
= & n \sum_{j=1}^{n} \sum_{i=1}^{d} \beta_{3} i^{-\gamma} j^{-\gamma} \\
= & \beta_{3} n \sum_{j=1}^{n} j^{-\gamma} \sum_{i=1}^{d} i^{-\gamma} \\
\leq & \beta_{3} n \log n \log d
\end{align*}
$$

## C. 2 Proof of Theorem 3

Proof. By plugging Lemma 2 into Theorem 2, we have

$$
\begin{align*}
& \sum_{t=1}^{T}\left[f_{t}\left(\theta_{t}\right)-f_{t}\left(\theta^{*}\right)\right]  \tag{23}\\
\leq & \rho_{1} d \sqrt{T}+\frac{\rho_{2} \sqrt{\beta_{3}}}{\sqrt{K}} \sqrt{d \log d} \frac{\sqrt{n \log n}}{n} \sqrt{T}+\rho_{3}
\end{align*}
$$

## C. 3 Proof of Lemma 3

Proof.

$$
\begin{align*}
& \min _{p \geq p_{\text {min }}} \sum_{j=1}^{n} \mathbb{E} \frac{\left\|g_{j}\right\|^{2}}{p_{j}} \\
= & \min _{p \geq p_{\text {min }}} \sum_{j=1}^{n} \frac{1}{p_{j}} \mathbb{E} \sum_{i=1}^{d} z_{j, i}^{2} \\
= & \min _{p \geq p_{\text {min }}} \sum_{j=1}^{n} \frac{1}{p_{j}} \sum_{i=1}^{d} \beta_{3} i^{-\gamma} j^{-\gamma}  \tag{24}\\
= & \min _{p \geq p_{\text {min }}} \sum_{j=1}^{n} \beta_{3} \frac{j^{-\gamma}}{p_{j}} \sum_{i=1}^{d} i^{-\gamma} \\
\leq & \min _{p \geq p_{\text {min }}} \sum_{j=1}^{n} \beta_{3} \frac{j^{-\gamma}}{p_{j}} \log d \\
= & \beta_{3} \log d \min _{p \geq p_{\text {min }}}^{\sum_{j=1}^{n} \frac{j^{-\gamma}}{p_{j}}}
\end{align*}
$$

From Proposition 5 in [2], we know

$$
\begin{equation*}
\min _{p \geq p_{\min }} \sum_{j=1}^{n} \frac{j^{-\gamma}}{p_{j}}=O\left(\log ^{2} n\right) \tag{25}
\end{equation*}
$$

Thus, we have

$$
\begin{aligned}
& \min _{p \geq p_{\text {min }}} \sum_{j=1}^{n} \mathbb{E} \frac{\left\|g_{j}\right\|^{2}}{p_{j}} \\
= & \beta_{3} \log d \min _{p \geq p_{\text {min }}} \sum_{j=1}^{n} \frac{j^{-\gamma}}{p_{j}} \\
= & O\left(\log d \log ^{2} n\right)
\end{aligned}
$$

## C. 4 Proof of Theorem 4

Proof. It follows simply by plugging Lemma 3 into Theorem 2.

## D AMSGrad with Bandit Sampling

We can also use bandit sampling to endow AMSGrad [3] with the ability to adapt to different training examples. Analogous to Theorem 1 in the main text, we have the following theorem regarding AMSGrad with Bandit Sampling.

Theorem D.1. Assume the gradient $\hat{G}_{t}$ is bounded, $\left\|\hat{G}_{t}\right\|_{\infty} \leq G_{\infty}$, and $\alpha_{t}=\alpha / \sqrt{t}, \beta_{1}=\beta_{11}$, and $\gamma=\beta_{1} / \sqrt{\beta_{2}}<1, \beta_{1 t}=\beta_{1} \lambda^{t-1}$. Then, AMSGrad with Bandit Sampling achieves the following convergence rate

$$
\begin{equation*}
R(T) \leq \rho_{1}^{\prime} d \sqrt{T}+\rho_{2}^{\prime} \sqrt{d(1+\log T)} \sqrt{\sum_{t=1}^{T}\left\|\hat{G}_{t}\right\|^{2}}+\rho_{3}^{\prime} \tag{27}
\end{equation*}
$$

where

$$
\begin{align*}
\rho_{1}^{\prime} & =\frac{D_{\infty}^{2} G_{\infty}}{\alpha\left(1-\beta_{1}\right)} \\
\rho_{2}^{\prime} & =\frac{\alpha}{\left(1-\beta_{1}\right)^{2}(1-\gamma) \sqrt{1-\beta_{2}}}  \tag{28}\\
\rho_{3}^{\prime} & =\frac{\beta_{1} D_{\infty}^{2} G_{\infty}}{2\left(1-\beta_{1}\right)(1-\lambda)^{2}}
\end{align*}
$$

Proof. According to Corrollary 1 from [3], we have

$$
\begin{equation*}
R(T) \leq \frac{\rho_{1}^{\prime} \sqrt{T}}{G_{\infty}} \sum_{i=1}^{d} \hat{v}_{T, i}^{1 / 2}+\rho_{2}^{\prime} \sqrt{(1+\log T)} \sum_{i=1}^{d}\left\|g_{1: T, i}\right\|_{2}+\rho_{3}^{\prime} \tag{29}
\end{equation*}
$$

From section "Proof of Theorem 1", we know that

$$
\begin{equation*}
\sum_{i=1}^{d} \sqrt{\hat{v}_{T, i}} \leq d G_{\infty} \tag{30}
\end{equation*}
$$

Therefore,

$$
\begin{align*}
R(T) \leq & \rho_{1}^{\prime} d \sqrt{T}+\rho_{2}^{\prime} \sqrt{1+\log T} \sum_{i=1}^{d}\left\|g_{1: T, i}\right\|_{2}+\rho_{3}^{\prime} \\
= & \rho_{1}^{\prime} d \sqrt{T}+d \rho_{2}^{\prime} \sqrt{1+\log T} \sum_{i=1}^{d} \frac{1}{d} \sqrt{\sum_{t=1}^{T}\left(\hat{G}_{t, i}\right)^{2}}+\rho_{3}^{\prime} \\
& \text { (due to } \sqrt{\cdot} \text { is concave) }  \tag{31}\\
\leq & \rho_{1}^{\prime} d \sqrt{T}+d \rho_{2}^{\prime} \sqrt{1+\log T} \sqrt{\sum_{i=1}^{d} \frac{1}{d} \sum_{t=1}^{T}\left(\hat{G}_{t, i}\right)^{2}}+\rho_{3}^{\prime} \\
= & \rho_{1}^{\prime} d \sqrt{T}+\sqrt{d(1+\log T)} \rho_{2}^{\prime} \sqrt{\sum_{t=1}^{T}\left\|\hat{G}_{t}\right\|^{2}}+\rho_{3}^{\prime}
\end{align*}
$$

Simiarly, we could derive theorems about AMSGrad with Bandit Sampling that are analogous to Theorem 2, 3, 4 for ADAMBS. We omit the details here.

## E Additional Experiments

## E. 1 Comparison with Basic SGD Variants

In the main paper, we compared our method with Adam and Adam with importance sampling. Here, we compare with some basic SGD variants (e.g., SGD, Momentum, Adagrad and RMSprop). As shown in Figure 1 , these basic SGD variants have worse performance than Adam-based methods.


Figure 1: Experiments comparing with SGD variants.

## E. 2 Comparison on Error Rate

In the main paper, we have shown the plots of loss value vs. wall clock time. Some might also be interested in the convergence curves of loss values. Here, we include some plots of error rate vs. wall clock time, as shown in Figure 2 and Figure 3 . They demonstrated the faster convergence of our method in terms of the error rate.


Figure 2: Error Plot of CNN models


Figure 3: Error Plot of RNN and RCNN models

## References

[1] D Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In International Conference on Learning Representations (ICLR), volume 5, 2015.
[2] Hongseok Namkoong, Aman Sinha, Steve Yadlowsky, and John C Duchi. Adaptive sampling probabilities for non-smooth optimization. In International Conference on Machine Learning, pages 2574-2583, 2017.
[3] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In International Conference on Learning Representations (ICLR), 2018.

