## A More on the background

## A. 1 SVRG and SCSG

Here we provide the pseudocode for SVRG (Algorithm 2) and SCSG (Algorithm 3) seen in Lei et al. [35]. The idea of SVRG (Algorithm[2] is to reuses past full gradient computations (line 3) to reduce the variance of the current stochastic gradient estimate (line 7) before the parameter update (line 8). Note that $N=1$ corresponds to a GD step (i.e., $v_{k-1}^{(j)} \leftarrow g_{j}$ in line 7). For $N>1, v_{k-1}^{(j)}$ is the corrected gradient in SVRG and is an unbiased estimate of the true gradient $\nabla J(\boldsymbol{\theta})$. SVRG achieves linear convergence $O(1 / T)$ using the semi-stochastic gradient.

```
Algorithm 2 SVRG
    Input: Number of stages \(T\), initial iteratre \(\tilde{\boldsymbol{\theta}}_{0}\), number of gradient steps \(N\), step size \(\eta\)
    for \(t=1\) to \(T\) do
        \(g_{t} \leftarrow \nabla J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)=\frac{1}{n} \sum_{i=1}^{n} \nabla J_{i}\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\)
        \(\boldsymbol{\theta}_{0}^{(t)} \leftarrow \tilde{\boldsymbol{\theta}}_{t-1}\)
        for \(k=1\) to \(N\) do
            Randomly pick \(i_{k} \in[n]\)
            \(v_{k-1}^{(t)} \leftarrow \nabla J_{i_{k}}\left(\boldsymbol{\theta}_{k-1}^{(t)}\right)-\nabla J_{i_{k}}\left(\boldsymbol{\theta}_{0}^{(t)}\right)+g_{t}\)
            \(\boldsymbol{\theta}_{k}^{(t)} \leftarrow \boldsymbol{\theta}_{k-1}^{(t)}-\eta_{t} v_{k-1}^{(t)}\)
        \(\tilde{\boldsymbol{\theta}}_{t} \leftarrow \boldsymbol{\theta}_{\tilde{N}_{t}}^{(t)}\)
    Output: \(\tilde{\boldsymbol{\theta}}_{T}\) (Convex case) or \(\tilde{\boldsymbol{\theta}}_{t}\) uniformly picked from \(\left\{\tilde{\boldsymbol{\theta}}_{t}\right\}_{t=1}^{T}\) (Non-Convex case)
```

More recently, Stochastically Controlled Stochastic Gradient (SCSG) has been proposed [34], to further reduce the computational cost of SVRG. The key difference is that SCSG (Algorithm 3) considers a sequence of time-varying batch sizes ( $B_{t}$ and $b_{t}$ ) and employs geometric sampling to generate the number of parameter update steps $N_{t}$ in each iteration (line 6), instead of fixing the batch sizes and the number of updates as done in SVRG. Particularly when finding an $\epsilon$-approximate solution (Definition [1] for optimizing smooth non-convex objectives, Lei et al. [35] proves that SCSG is never worse than SVRG in convergence rate and significantly outperforms SVRG when the required $\epsilon$ is small.

```
Algorithm 3 SCSG for smooth non-convex objectives
    Input: Number of stages \(T\), initial iteratre \(\tilde{\boldsymbol{\theta}}_{0}\), batch size \(B_{t}\), mini-batch size \(b_{t}\), step size \(\eta_{t}\)
    for \(t=1\) to \(T\) do
        Uniformly sample a batch \(\mathcal{I}_{t} \subset\{1, \cdots, n\}\) with \(\left|\mathcal{I}_{t}\right|=B_{t}\)
        \(g_{t} \leftarrow \nabla J_{\mathcal{L}_{t}}\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\)
        \(\boldsymbol{\theta}_{0}^{(t)} \leftarrow \tilde{\boldsymbol{\theta}}_{t-1}\)
        Generate \(N_{t} \sim \operatorname{Geom}\left(B_{t} /\left(B_{t}+b_{t}\right)\right)\)
        for \(k=1\) to \(N_{t}\) do
            Randomly pick \(\tilde{\mathcal{I}}_{k-1} \subset[n]\) with \(\left|\tilde{\mathcal{I}}_{k-1}\right|=b_{t}\)
            \(v_{k-1}^{(t)} \leftarrow \nabla J_{\tilde{\mathcal{I}}_{k-1}}\left(\boldsymbol{\theta}_{k-1}^{(t)}\right)-\nabla J_{\tilde{\mathcal{I}}_{k-1}}\left(\boldsymbol{\theta}_{0}^{(t)}\right)+g_{t}\)
            \(\boldsymbol{\theta}_{k}^{(t)} \leftarrow \boldsymbol{\theta}_{k-1}^{(t)}+\eta_{t} v_{k-1}^{(t)}\)
        \(\tilde{\boldsymbol{\theta}}_{t} \leftarrow \boldsymbol{\theta}_{N_{t}}^{(t)}\)
    Output: \(\tilde{\boldsymbol{\theta}}_{T}\) (P-L case) or sample \(\tilde{\boldsymbol{\theta}}_{T}^{*}\) from \(\left\{\tilde{\boldsymbol{\theta}}_{t}\right\}_{t=1}^{T}\) with \(P\left(\tilde{\boldsymbol{\theta}}_{T}^{*}=\tilde{\boldsymbol{\theta}}_{t}\right) \propto \eta_{t} B_{t} / b_{t}\) (Smooth case)
```

As a member of the SVRG-like algorithms, SCSG enjoys the same convergence rate of SVRG while being computationally cheaper than SVRG for tasks with small $\epsilon$ requirements [34], which is highly desired in RL, hence the motivation of FedPG-BR to adapt SCSG.

## A. 2 Gradient estimator

Use $g(\tau \mid \boldsymbol{\theta})$ to denote the unbiased estimator to the true gradient $J(\boldsymbol{\theta})$. The common gradient estimators are the REINFORCE and the GPOMDP estimators, which are considered as baseline estimators in [18] and [19]. The REINFORCE [39]:

$$
g(\tau \mid \boldsymbol{\theta})=\left[\sum_{h=0}^{H-1} \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}\left(a_{h} \mid s_{h}\right)\right]\left[\sum_{h=0}^{H-1} \gamma^{h} \mathcal{R}\left(s_{h}, a_{h}\right)-C_{b}\right]
$$

And the GPOMDP [40]

$$
g(\tau \mid \boldsymbol{\theta})=\sum_{h=0}^{H-1}\left[\sum_{t=0}^{h} \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}\left(a_{t} \mid s_{t}\right)\right]\left(\gamma^{h} r\left(s_{h}, a_{h}\right)-C_{b_{h}}\right)
$$

where $C_{b}$ and $C_{b_{h}}$ are the corresponding baselines. Under Assumption 3, whether we use the REINFORCE or the GPOMDP estimator, Proposition 4 holds [18, 19].

```
Algorithm 4 GPOMDP (for federation of K agents)
    Input: number of iterations \(T\), batch size \(B\), step size \(\eta\), initial parameter \(\tilde{\boldsymbol{\theta}}_{0} \in \mathbb{R}^{d}\)
    for \(t=1\) to \(T\) do
        \(\boldsymbol{\theta}^{t} \leftarrow \tilde{\boldsymbol{\theta}}_{t-1} \quad\); broadcast to agents
        for \(k=1\) to \(K\) do
            Sample \(B\) trajectories \(\left\{\tau_{t, i}^{(k)}\right\}\) from \(p\left(\cdot \mid \boldsymbol{\theta}^{t}\right)\)
            \(\mu_{t}^{(k)}=\frac{1}{B} \sum_{i=1}^{B} g\left(\tau_{t, i}^{(k)} \mid \boldsymbol{\theta}^{t}\right) \quad ;\) push \(\mu_{t}^{(k)}\) to server
        \(\mu_{\tilde{\theta}_{t}}=\frac{1}{K} \sum_{k=1}^{K} \mu_{t}^{(k)}\)
        \(\tilde{\boldsymbol{\theta}}_{t} \leftarrow \boldsymbol{\theta}^{t}+\eta \mu_{t}\)
    Output \(\boldsymbol{\theta}_{\text {out }}\) : uniformly randomly picked from \(\left\{\tilde{\boldsymbol{\theta}}_{t}\right\}_{t=1}^{T}\)
```


## A. 3 Federated GPOMDP and SVRPG

Closely following the problem setting of FedPG-BR, we adapt both GPOMDP and SVRPG to the FRL setting. The pseudocode is shown in Algorithm 4 and Algorithm5.

```
Algorithm 5 SVRPG (for federation of K agents)
    Input: number of epochs \(T\), epoch size \(N\), batch size \(B\), mini-batch size \(b\), step size \(\eta\), initial
    parameter \(\tilde{\boldsymbol{\theta}}_{0} \in \mathbb{R}^{d}\)
    for \(t=1\) to \(T\) do
        \(\boldsymbol{\theta}_{0}^{t} \leftarrow \tilde{\boldsymbol{\theta}}_{t-1} \quad\); broadcast to agents
        for \(k=1\) to \(K\) do
            Sample \(B\) trajectories \(\left\{\tau_{t, i}^{(k)}\right\}\) from \(p\left(\cdot \mid \boldsymbol{\theta}_{0}^{t}\right)\)
            \(\mu_{t}^{(k)}=\frac{1}{B} \sum_{i=1}^{B} g\left(\tau_{t, i}^{(k)} \mid \boldsymbol{\theta}_{0}^{t}\right) \quad ;\) push \(\mu_{t}^{(k)}\) to server
        \(\mu_{t}=\frac{1}{K} \sum_{k=1}^{K} \mu_{t}^{(k)}\)
        for \(n=0\) to \(N-1\) do
            Sample \(b\) trajectories \(\left\{\tau_{n, j}^{t}\right\}\) from \(p\left(\cdot \mid \boldsymbol{\theta}_{n}^{t}\right)\)
            \(v_{n}^{t}=\frac{1}{b} \sum_{j=1}^{b}\left[g\left(\tau_{n, j}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-\omega\left(\tau_{n, j}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n, j}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right]+\mu_{t}\)
            \(\boldsymbol{\theta}_{n+1}^{t}=\boldsymbol{\theta}_{n}^{t}+\eta v_{n}^{t}\)
        \(\tilde{\boldsymbol{\theta}}_{t} \leftarrow \boldsymbol{\theta}_{N}^{t}\)
    Output \(\boldsymbol{\theta}_{\text {out }}\) : uniformly randomly picked from \(\left\{\tilde{\boldsymbol{\theta}}_{t}\right\}_{t=1}^{T}\)
```


## B Proof of Theorem 6

In our proof, we follow the suggestion from Lei et al. [35] to set $b_{t}=1$ to derive better theoretical results. Refer to Section $F$ in this appendix for the value of $b_{t}$ used in our experiments.

Proof. From the L-smoothness of the objective function $J(\boldsymbol{\theta})$, we have

$$
\begin{align*}
\mathbb{E}_{\tau_{n}^{t}}\left[J\left(\boldsymbol{\theta}_{n+1}^{t}\right)\right] \geq & \mathbb{E}_{\tau_{n}^{t}}\left[J\left(\boldsymbol{\theta}_{n}^{t}\right)+\left\langle\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right), \boldsymbol{\theta}_{n+1}^{t}-\boldsymbol{\theta}_{n}^{t}\right\rangle-\frac{L}{2}\left\|\boldsymbol{\theta}_{n+1}^{t}-\boldsymbol{\theta}_{n}^{t}\right\|^{2}\right] \\
= & J\left(\boldsymbol{\theta}_{n}^{t}\right)+\eta_{t}\left\langle\mathbb{E}_{\tau_{n}^{t}}\left[v_{n}^{t}\right], \nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\rangle-\frac{L \eta_{t}^{2}}{2} \mathbb{E}_{\tau_{n}^{t}}\left[\left\|v_{n}^{t}\right\|^{2}\right] \\
\geq & J\left(\boldsymbol{\theta}_{n}^{t}\right)+\eta_{t}\left\langle\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)+e_{t}, \nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\rangle \\
& \quad-\frac{L \eta_{t}^{2}}{2}\left[\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2}\right]  \tag{4}\\
= & J\left(\boldsymbol{\theta}_{n}^{t}\right)+\eta_{t}\left(1-L \eta_{t}\right)\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+\eta_{t}\left\langle e_{t}, \nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\rangle \\
& \quad-L \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right)\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}-L \eta_{t}^{2}\left\|e_{t}\right\|^{2}
\end{align*}
$$

where (4) follows from Lemma 11 . Use $\mathbb{E}_{t}$ to denote the expectation with respect to all trajectories $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$, given $N_{t}$. Since $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$ are independent of $N_{t}, \mathbb{E}_{t}$ is equivalently the expectation with respect to $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$. The above inequality gives

$$
\begin{aligned}
\mathbb{E}_{t}\left[J\left(\boldsymbol{\theta}_{n+1}^{t}\right)\right] \geq \mathbb{E}_{t}\left[J\left(\boldsymbol{\theta}_{n}^{t}\right)\right] & +\eta_{t}\left(1-L \eta_{t}\right) \mathbb{E}_{t}\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+\eta_{t} \mathbb{E}_{t}\left\langle e_{t}, \nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\rangle \\
& -L \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}-L \eta_{t}^{2}\left\|e_{t}\right\|^{2}
\end{aligned}
$$

Taking $n=N_{t}$ and using $\mathbb{E}_{N_{t}}$ to denote the expectation w.r.t. $N_{t}$, we have from the above

$$
\begin{aligned}
\mathbb{E}_{N_{t}} \mathbb{E}_{t}\left[J\left(\boldsymbol{\theta}_{N_{t}+1}^{t}\right)\right] \geq \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left[J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right] & +\eta_{t}\left(1-L \eta_{t}\right) \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\|^{2}+\eta_{t} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\langle e_{t}, \nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\rangle \\
& -L \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}-L \eta_{t}^{2}\left\|e_{t}\right\|^{2}
\end{aligned}
$$

Rearrange,

$$
\begin{align*}
\eta_{t}\left(1-L \eta_{t}\right) \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\|^{2} \leq & \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left[J\left(\boldsymbol{\theta}_{N_{t}+1}^{t}\right)\right]+L \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2} \\
& -\eta_{t} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\langle e_{t}, \nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\rangle+L \eta_{t}^{2}\left\|e_{t}\right\|^{2}-\mathbb{E}_{N_{t}} \mathbb{E}_{t}\left[J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right] \\
= & \frac{1}{B_{t}}\left(\mathbb{E}_{t} \mathbb{E}_{N_{t}}\left[J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right]-J\left(\boldsymbol{\theta}_{0}^{t}\right)\right)-\eta_{t} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\langle e_{t}, \nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\rangle \\
& +L \eta_{t}^{2}\left(L g+C_{g}^{2} C_{w}\right) \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+L \eta_{t}^{2}\left\|e_{t}\right\|^{2} \tag{5}
\end{align*}
$$

where (5) follows from Lemma 16 with Fubini's theorem. Note that $\tilde{\boldsymbol{\theta}}_{t}=\boldsymbol{\theta}_{N_{t}}^{t}$ and $\tilde{\boldsymbol{\theta}}_{t-1}=\boldsymbol{\theta}_{0}^{t}$. If we take expectation over all the randomness and denote it by $\mathbb{E}$, we get

$$
\begin{align*}
\eta_{t}\left(1-L \eta_{t}\right) \mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2}= & \frac{1}{B_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]-\eta_{t} \mathbb{E}\left\langle e_{t}, \nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\rangle \\
& +L \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) \mathbb{E}\left\|\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\|^{2}+L \eta_{t}^{2} \mathbb{E}\left\|e_{t}\right\|^{2} \\
= & \frac{1}{B_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]-\frac{1}{B_{t}} \mathbb{E}\left\langle e_{t}, \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle \\
& +L \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) \mathbb{E}\left\|\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\|^{2}+\eta_{t}\left(1+L \eta_{t}\right) \mathbb{E}\left\|e_{t}\right\|^{2}  \tag{6}\\
\leq & \frac{1}{B_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right] \\
& +\frac{1}{2 \eta_{t} B_{t}}\left[-\frac{1}{B_{t}}+\eta_{t}^{2}\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\right] \mathbb{E}\left\|\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\|^{2} \\
& +\frac{1}{B_{t}} \mathbb{E}\left\langle\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right), \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle+\frac{\eta_{t}}{B_{t}} \mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2}+\frac{\eta_{t}}{B_{t}} \mathbb{E}\left\|e_{t}\right\|^{2} \\
& +L \eta_{t}^{2}\left(L L_{g}+C_{g}^{2} C_{w}\right) \mathbb{E}\left\|\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\|^{2}+\eta_{t}\left(1+L \eta_{t}\right) \mathbb{E}\left\|e_{t}\right\|^{2} \tag{7}
\end{align*}
$$

where (6) follows from Lemma 12 and (7) follows from Lemma 13 . Rearrange,

$$
\begin{align*}
\eta_{t}(1 & \left.-L \eta_{t}-\frac{1}{B_{t}}\right) \mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2}+\frac{1-2 \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}-2 L \eta_{t}^{3}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}^{2}}{2 \eta_{t} B_{t}^{2}} \mathbb{E}\left\|\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\|^{2} \\
& \leq \frac{1}{B_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]+\frac{1}{B_{t}} \mathbb{E}\left\langle\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right), \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle+\eta_{t}\left(1+L \eta_{t}+\frac{1}{B_{t}}\right) \mathbb{E}\left\|e_{t}\right\|^{2} \tag{8}
\end{align*}
$$

Now we can apply Lemma 17 on $\mathbb{E}\left\langle\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right), \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle$ using $a=\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}, b=\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)$, and $\beta=\frac{1-2 \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}-2 L \eta_{t}^{3}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}^{2}}{\eta_{t} B_{t}}$ to get

$$
\begin{align*}
\frac{1}{B_{t}} \mathbb{E}\left\langle\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right), \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle & \leq \frac{1-2 \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}-2 L \eta_{t}^{3}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}^{2}}{2 \eta_{t} B_{t}^{2}} \mathbb{E}\left\|\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\|^{2} \\
& +\frac{\eta_{t}}{2\left[1-2 \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}-2 L \eta_{t}^{3}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}^{2}\right]} \mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2} \tag{9}
\end{align*}
$$

Combining (8) and (9) and rearrange, we have

$$
\begin{align*}
\eta_{t}(1- & \left.L \eta_{t}-\frac{1}{B_{t}}-\frac{1}{2\left[1-2 \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}-2 L \eta_{t}^{3}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}^{2}\right]}\right) \mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2} \\
& \leq \frac{1}{B_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]+\eta_{t}\left(1+L \eta_{t}+\frac{1}{B_{t}}\right) \mathbb{E}\left\|e_{t}\right\|^{2} \\
& \leq \frac{1}{B_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]+\eta_{t}\left(1+L \eta_{t}+\frac{1}{B_{t}}\right)\left[\frac{4 \sigma^{2}}{(1-\alpha)^{2} K B_{t}}+\frac{48 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B_{t}}\right] \tag{10}
\end{align*}
$$

where (10) follows from Lemma 15. We want to choose $\eta_{t}$ such that $1-2 \eta_{t}^{2}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}-$ $2 L \eta_{t}^{3}\left(L_{g}+C_{g}^{2} C_{w}\right) B_{t}^{2}>0$. Denoting $\Phi=L_{g}+C_{g}^{2} C_{w}$, we have

$$
\begin{array}{r}
1-2 \eta_{t}^{2} \Phi B_{t}-2 L \eta_{t}^{3} \Phi B_{t}^{2} \geq 0 \\
2 \eta_{t}^{2} \Phi B_{t}+2 L \eta_{t}^{3} \Phi B_{t}^{2} \leq 1
\end{array}
$$

We can then choose $\eta_{t}$ to have $2 \eta_{t}^{2} \Phi B_{t} \leq \frac{1}{2}$ and $2 L \eta_{t}^{3} \Phi B_{t}^{2} \leq \frac{1}{2}$, which implies:
(i) $\eta_{t} \leq \frac{1}{\left(4 \Phi B_{t}\right)^{1 / 2}}$
(ii) $\eta_{t} \leq \frac{1}{\left(4 L \Phi B_{t}^{2}\right)^{1 / 3}}$

Therefore, we can choose $\eta_{t} \leq \frac{1}{2\left(L \Phi B_{t}^{2}\right)^{1 / 3}}<\frac{1}{\left(4 L \Phi B_{t}^{2}\right)^{1 / 3}}$, and set $B_{t} \geq 4 \Phi L^{-2}$ to ensure both condition (i) and (ii) are satisfied, together with $\frac{1}{\left(4 L \Phi B_{t}^{2}\right)^{1 / 3}} \leq \frac{1}{\left(4 \Phi B_{t}\right)^{1 / 2}}$, and $L \eta_{t}+\frac{1}{B_{t}} \leq 1$. Thus, by choosing $\eta_{t} \leq \frac{1}{2 \Psi B_{t}^{2 / 3}}$, where $\Psi=(L \Phi)^{1 / 3}=\left(L\left(L_{g}+C_{g}^{2} C_{w}\right)\right)^{1 / 3}$, we can obtain the following from (10):

$$
\eta_{t} \mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2} \leq \frac{1}{B_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]+2 \eta_{t}\left[\frac{4 \sigma^{2}}{(1-\alpha)^{2} K B_{t}}+\frac{48 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B_{t}}\right]
$$

Replacing $\eta_{t}=\frac{1}{2 \Psi B_{t}^{2 / 3}}$ and rearranging, we have

$$
\begin{aligned}
\mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2} & \leq \frac{1}{B_{t} \eta_{t}} \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]+2\left[\frac{4 \sigma^{2}}{(1-\alpha)^{2} K B_{t}}+\frac{48 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B_{t}}\right] \\
& \leq \frac{2 \Psi \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{t}\right)-J\left(\tilde{\boldsymbol{\theta}}_{t-1}\right)\right]}{B_{t}^{1 / 3}}+\frac{8 \sigma^{2}}{(1-\alpha)^{2} K B_{t}}+\frac{96 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B_{t}}
\end{aligned}
$$

Replacing $B_{t}$ with constant batch size $B$ and telescoping over $t=1,2, \ldots, T$, we have for $\tilde{\boldsymbol{\theta}}_{a}$ from our algorithm:

$$
\begin{aligned}
\mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{a}\right)\right\|^{2} & \leq \frac{2 \Psi \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}_{T}\right)-J\left(\tilde{\boldsymbol{\theta}}_{0}\right)\right]}{T B^{1 / 3}}+\frac{8 \sigma^{2}}{(1-\alpha)^{2} K B}+\frac{96 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B} \\
& \leq \frac{2 \Psi\left[J\left(\tilde{\boldsymbol{\theta}}^{*}\right)-J\left(\tilde{\boldsymbol{\theta}}_{0}\right)\right]}{T B^{1 / 3}}+\frac{8 \sigma^{2}}{(1-\alpha)^{2} K B}+\frac{96 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B}
\end{aligned}
$$

which completes the proof.

## C Proof of Corollary 7

Proof. Recall $\Psi=\left(L\left(L_{g}+C_{g}^{2} C_{w}\right)\right)^{1 / 3}$. From Theorem6, we have

$$
\mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{a}\right)\right\|^{2} \leq \underbrace{\frac{2 \Psi \mathbb{E}\left[J\left(\tilde{\boldsymbol{\theta}}^{*}\right)-J\left(\tilde{\boldsymbol{\theta}}_{0}\right)\right]}{T B^{1 / 3}}}_{T=O\left(\frac{1}{\epsilon B^{1 / 3}}\right)}+\underbrace{\frac{8 \sigma^{2}}{(1-\alpha)^{2} K B}}_{B_{K}=O\left(\frac{1}{\epsilon K}\right)}+\underbrace{\frac{96 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B}}_{B_{\alpha}=O\left(\frac{\alpha^{2}}{\epsilon}\right)}
$$

To guarantee that the output of Algorithm 1 is $\epsilon$-approximate, i.e., $\mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{a}\right)\right\|^{2} \leq \epsilon$, we need the number of rounds $T$ and the batch size $B$ to meet the following:

$$
\text { (i) } T=O\left(\frac{1}{\epsilon B^{1 / 3}}\right), \text { (ii) } B_{K}=O\left(\frac{1}{\epsilon K}\right), \text { and (iii) } B_{\alpha}=O\left(\frac{\alpha^{2}}{\epsilon}\right)
$$

By union bound and using $\mathbb{E}[\operatorname{Traj}(\epsilon)]$ to denote the total number of trajectories required by each agent to sample, the above implies that

$$
\begin{aligned}
\mathbb{E}[\operatorname{Traj}(\epsilon)] & \leq T B_{K}+T B_{\alpha} \\
& \leq O\left(\frac{1}{\epsilon^{5 / 3} K^{2 / 3}}+\frac{\alpha^{4 / 3}}{\epsilon^{5 / 3}}\right)
\end{aligned}
$$

in order to obtain an $\epsilon$-approximate policy, which completes the proof for Corollary 7 ( $i$ ). Note that the total number of trajectories generated across the whole FRL system, denoted by $\mathbb{E}\left[\operatorname{Traj}_{\text {total }}(\epsilon)\right]$ is thus bounded by:

$$
\mathbb{E}\left[\operatorname{Traj}_{\text {total }}(\epsilon)\right] \leq O\left(\frac{K^{1 / 3}}{\epsilon^{5 / 3}}+\frac{K \alpha^{4 / 3}}{\epsilon^{5 / 3}}\right)
$$

Now for an ideal system where $\alpha=0$ :

$$
\begin{aligned}
\mathbb{E}[\operatorname{Traj}(\epsilon)] & \leq O\left(\frac{1}{\epsilon^{5 / 3} K^{2 / 3}}\right) \\
\mathbb{E}\left[\operatorname{Traj}_{\text {total }}(\epsilon)\right] & \leq O\left(\frac{K^{1 / 3}}{\epsilon^{5 / 3}}\right)
\end{aligned}
$$

which completes the proof for Corollary 7 (ii). Moreover, when $K=1$, the number of trajectories required by the agent using FedPG-BR is

$$
\mathbb{E}[\operatorname{Traj}(\epsilon)] \leq O\left(\frac{1}{\epsilon^{5 / 3}}\right)
$$

which is Corollary 7 (iii) and is coherent with the recent analysis of SVRPG [19].

## D More on the Byzantine Filtering Step

In this section, we continue our discussion on our Byzantine Filtering Step in Section 3.3 We include the pseudocode for the subroutine FedPG-Aggregate below for ease of reference:
As discussed in Section 3.3, R2 (line 8 in Algorithm 1.1) ensures that $\mathcal{G}_{t}$ always include all good agents and for any Byzantine agents being included, their impact on the convergence of Algorithm 1 is limited since their maximum distance to $\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)$ is bounded by $3 \sigma$. Here we give proofs for the claims.
Claim D.1. Under Assumption 2 and $\forall \alpha<0.5$, the filtering rule $R 2$ in Algorithm 1.1 ensures that, in any round $t$, all gradient estimates sent from non-Byzantine agents are included in $\mathcal{G}_{t}$, i.e., $\left|\mathcal{G}_{t}\right| \geq(1-\alpha) K$.

Proof. First, from Assumption 2

$$
\left\|\mu_{t}^{(k)}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| \leq \sigma, \forall k \in \mathcal{G}
$$

```
Algorithm 1.1 FedPG-Aggregate
    Input: Gradient estimates from \(K\) agents in round \(t\) : \(\left\{\mu_{t}^{(k)}\right\}_{k=1}^{k}\), Variance Bound \(\sigma\), filtering
    threshold \(\mathfrak{T}_{\mu} \triangleq 2 \sigma \sqrt{\frac{V}{B_{t}}}\), where \(V \triangleq 2 \log \left(\frac{2 K}{\delta}\right)\) and \(\delta \in(0,1)\)
    \(S_{1} \triangleq\left\{\mu_{t}^{(k)}\right\}\) where \(k \in[K]\) s.t. \(\left|\left\{k^{\prime} \in[K]:\left\|\mu_{t}^{\left(k^{\prime}\right)}-\mu_{t}^{(k)}\right\| \leq \mathfrak{T}_{\mu}\right\}\right|>\frac{K}{2}\)
    \(\mu_{t}^{\operatorname{mom}} \leftarrow \operatorname{argmin}_{\mu_{t}^{(\tilde{k})}}\left\|\mu_{t}^{(\tilde{k})}-\operatorname{mean}\left(S_{1}\right)\right\|\) where \(\tilde{k} \in S_{1}\)
    \(R 1: \mathcal{G}_{t} \triangleq\left\{k \in[K]:\left\|\mu_{t}^{(k)}-\mu_{t}^{\operatorname{mom}}\right\| \leq \mathfrak{T}_{\mu}\right\}\)
    if \(\left|\mathcal{G}_{t}\right|<(1-\alpha) K \quad\) then
        \(S_{2} \triangleq\left\{\mu_{t}^{(k)}\right\}\) where \(k \in[K]\) s.t. \(\left|\left\{k^{\prime} \in[K]:\left\|\mu_{t}^{\left(k^{\prime}\right)}-\mu_{t}^{(k)}\right\| \leq 2 \sigma\right\}\right|>\frac{K}{2}\)
        \(\mu_{t}^{\text {mom }} \leftarrow \operatorname{argmin}_{\mu_{t}^{(\tilde{k})}}\left\|\mu_{t}^{(\tilde{k})}-\operatorname{mean}\left(S_{2}\right)\right\|\) where \(\tilde{k} \in S_{2}\)
        \(R 2: \mathcal{G}_{t} \triangleq\left\{k \in[K]:\left\|\mu_{t}^{(k)}-\mu_{t}^{\operatorname{mom}}\right\| \leq 2 \sigma\right\}\)
    Return: \(\mu_{t} \triangleq \frac{1}{\left|\mathcal{G}_{t}\right|} \sum_{k \in \mathcal{G}_{t}} \mu_{t}^{(k)}\)
```

it implies that $\left\|\mu_{t}^{\left(k_{1}\right)}-\mu_{t}^{\left(k_{2}\right)}\right\| \leq 2 \sigma, \forall k_{1}, k_{2} \in \mathcal{G}$. So, for any value of the vector median [46] in S2 $=\left\{\mu_{t}^{(k)}\right\}$ (defined in line 6$)$ :

$$
\left\|\mu_{t}^{(k)}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| \leq 3 \sigma, \forall \mu_{t}^{(k)} \in \mathbf{S} 2
$$

An intuitive illustration is provided in Fig. 4. Next, consider the worst case where all values sent by the $K$ agents are included in S 2 : for all $(1-\alpha) K$ good agents, they send the same value $\mu_{t}^{(k)}$, s.t., $\left\|\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)-\mu_{t}^{(k)}\right\|=\sigma, \forall k \in \mathcal{G}$; and for all $\alpha K$ Byzantine agents, they send the same value $\mu_{t}^{\left(k^{\prime}\right)}$ s.t., $\left\|\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)-\mu_{t}^{\left(k^{\prime}\right)}\right\|=3 \sigma, \forall k^{\prime} \in S 2 \backslash \mathcal{G}$. Then the mean of values in S 2 satisfies:

$$
\begin{aligned}
\left\|\mu_{t}^{\text {mean }}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| & =\frac{(1-\alpha) K \cdot \sigma+\alpha K \cdot 3 \sigma}{K} \\
& =(1-\alpha) \sigma+3 \alpha \sigma \\
& =\sigma+2 \alpha \sigma \\
& <2 \sigma
\end{aligned}
$$

where the last inequality holds for $\alpha<0.5$ which is our assumption. Then the value $\mu_{t}^{\text {mom }}$ of Algorithm 1 will be set to any $\mu_{t}^{(k)}$ from S 2 , of which is the closet to $\mu_{t}^{\text {mean }}$.
The selection of $\mu_{t}^{\text {mom }}$ implies $\left\|\mu_{t}^{\text {mom }}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| \leq \sigma$. Therefore, by constructing a region of $\mathcal{G}_{t}$ that is centred at $\mu_{t}^{\mathrm{mom}}$ and $2 \sigma$ in radius (line 8 ), $\mathcal{G}_{t}$ can cover all estimates from non-Byzantine agents and hence ensure $\left|\mathcal{G}_{t}\right| \geq(1-\alpha) K$.

Claim D.2. Under Assumption 2 and $\alpha<0.5$, the filtering rule $R 2$ in Algorithm 1.1 ensures that, in any round $t,\left\|\mu_{t}^{(k)}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| \leq 3 \sigma, \forall k \in \mathcal{G}_{t}$.

Proof. This lemma is a straightforward result following the proof of Claim D. 1 .

Remark. Claim $\sqrt{D .2}$ implies that, in any round $t$, if an estimate sent from Byzantine agent is included in $\mathcal{G}_{t}$, then its impact on the convergence of Algorithm 1 is limited since its distance to $\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)$ is bounded by $3 \sigma$. Fig. 4 provides an intuitive illustration for this claim.


Figure 4: Graphical illustration of the Byzantine filtering strategy where $\mu_{t}^{\left(k_{1}\right)}, \mu_{t}^{\left(k_{2}\right)}$ are two good gradients while the red cross represents one Byzantine gradient which falls within $S_{2} . \mu_{t}^{\text {mom }}$ will be chosen at the red diamond.

As discussed above, R2 ensures that all good agents are included in $\mathcal{G}_{t}$, i.e., a region in which all good agents are concentrated. R1 (lines 2-4) is designed in a similar way and aims to improve the practical performance of FedPG-BR by exploiting Lemma 14. all good agents are highly likely to be concentrated in a much smaller region.
Claim D.3. Define $V \triangleq 2 \log (2 K / \delta)$ and $\delta \in(0,1)$, the filtering R1 in Algorithm 1 ensure

$$
\left\|\mu_{t}^{(k)}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| \leq \sigma \sqrt{\frac{V}{B_{t}}}, \forall k \in \mathcal{G}
$$

with probability of at least $1-\delta$.
Proof. From Assumption $2,\left\|\mu_{t}^{(k)}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| \leq \sigma, \forall k \in \mathcal{G}$. We have

$$
\begin{align*}
\left\|\mu_{t}^{(k)}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| & =\left\|\frac{1}{B_{t}} \sum_{i=1}^{B_{t}} g\left(\tau_{t, i}^{(k)} \mid \boldsymbol{\theta}_{0}^{t}\right)-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\| \\
& =\frac{1}{B_{t}} \sqrt{\left\|\sum_{i=1}^{B_{t}} g\left(\tau_{t, i}^{(k)} \mid \boldsymbol{\theta}_{0}^{t}\right)-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right\|^{2}} \tag{11}
\end{align*}
$$

Consider $X_{i} \triangleq g\left(\tau_{t, i}^{(k)}\right)-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)$ and apply Lemma 14 on 11, we have

$$
\begin{aligned}
\operatorname{Pr}\left[\left\|\sum_{i=1}^{B_{t}} X_{i}\right\|^{2} \leq 2 \log \left(\frac{2}{\delta}\right) \sigma^{2} B_{t}\right] & \geq 1-\delta \\
\operatorname{Pr}\left[\frac{1}{B_{t}} \sqrt{\left\|\sum_{i=1}^{B_{t}} X_{i}\right\|^{2}} \leq \frac{1}{B_{t}} \sqrt{2 \log \left(\frac{2}{\delta}\right) \sigma^{2} B_{t}}\right] & \geq 1-\delta
\end{aligned}
$$

With $V \triangleq 2 \log (2 K / \delta)$ and $\delta \in(0,1)$, the above inequality yields the Claim.
Therefore, the first filtering R1 (lines 2-4) of FedPG-BR constructs a region of $\mathcal{G}_{t}$ centred at $\mu_{t}^{\text {mom }}$ with radius of $2 \sigma \sqrt{\frac{V}{B_{t}}}$, which ensures in any round $t$ that, with probability $\geq 1-\delta$, (a) all good agents are included in $\mathcal{G}_{t}$, and (b) if gradients from Byzantine agents are included in $\mathcal{G}_{t}$, their impact is limited since their maximum distance to $\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)$ is bounded by $3 \sigma \sqrt{\frac{V}{B_{t}}}$ (The proof is similar to that of ClaimD.2. Compared to R2, R1 can construct a smaller region that the server believes contains all good agents. If any Byzantine agent is included, their impact is also smaller, with probability of at least $1-\delta$. Therefore, R1 is applied first such that if R1 fails (line 5) which happens with probability $<\delta, \mathrm{R} 2$ is then employed as a backup to ensure that $\mathcal{G}_{t}$ always includes all good agents.

## E Useful technical lemmas

Lemma 8 (Unbiaseness of importance sampling).

$$
\begin{aligned}
\mathbb{E}_{\tau \sim p\left(\cdot \mid \boldsymbol{\theta}_{n}\right)}\left[\omega\left(\tau \mid \boldsymbol{\theta}_{n}, \boldsymbol{\theta}_{0}\right) g\left(\tau \mid \boldsymbol{\theta}_{0}\right)\right] & =\mathbb{E}_{\tau \sim p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)}\left[g\left(\tau \mid \boldsymbol{\theta}_{0}\right)\right] \\
& =\nabla J\left(\boldsymbol{\theta}_{0}\right)
\end{aligned}
$$

Proof. Drop $t$ from notation and use $\tau_{n}$ to denote trajectories sampled from $\boldsymbol{\theta}_{n}$ at step $n$. From the definition of gradient estimation, we have

$$
\begin{aligned}
g\left(\tau_{n} \mid \boldsymbol{\theta}_{0}\right) & =\mathbb{E}_{\tau \sim p\left(\cdot \mid \boldsymbol{\theta}_{n}\right)}\left[\nabla_{\boldsymbol{\theta}_{0}} p\left(\boldsymbol{\theta}_{0}\right) r(\tau)\right] \\
& =\int p\left(\cdot \mid \boldsymbol{\theta}_{n}\right) \nabla_{\boldsymbol{\theta}_{0}} p\left(\boldsymbol{\theta}_{0}\right) r(\tau) d \tau \\
& =\int \frac{p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)}{p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)} p\left(\cdot \mid \boldsymbol{\theta}_{n}\right) \nabla_{\boldsymbol{\theta}_{0}} p\left(\boldsymbol{\theta}_{0}\right) r(\tau) d \tau \\
& =\int p\left(\cdot \mid \boldsymbol{\theta}_{0}\right) \frac{p\left(\cdot \mid \boldsymbol{\theta}_{n}\right)}{p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)} \nabla_{\boldsymbol{\theta}_{0}} p\left(\boldsymbol{\theta}_{0}\right) r(\tau) d \tau \\
& =\mathbb{E}_{\tau \sim p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)}\left[\frac{p\left(\cdot \mid \boldsymbol{\theta}_{n}\right)}{p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)} \nabla_{\boldsymbol{\theta}_{0}} p\left(\boldsymbol{\theta}_{0}\right) r(\tau)\right] \\
& =\frac{p\left(\cdot \mid \boldsymbol{\theta}_{n}\right)}{p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)} g\left(\tau_{0} \mid \boldsymbol{\theta}_{0}\right)
\end{aligned}
$$

Then,

$$
\begin{aligned}
\omega\left(\tau \mid \boldsymbol{\theta}_{n}, \boldsymbol{\theta}_{0}\right) g\left(\tau_{n} \mid \boldsymbol{\theta}_{0}\right) & =\frac{p\left(\cdot \mid \boldsymbol{\theta}_{0}\right)}{p\left(\cdot \mid \boldsymbol{\theta}_{n}\right)} g\left(\tau_{n} \mid \boldsymbol{\theta}_{0}\right) \\
& =g\left(\tau_{0} \mid \boldsymbol{\theta}_{0}\right)
\end{aligned}
$$

which gives the lemma.
Lemma 9 (Adapted from [19|). Let $\omega\left(\tau \mid \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)=p\left(\tau \mid \boldsymbol{\theta}_{1}\right) / p\left(\tau \mid \boldsymbol{\theta}_{2}\right)$, under Assumptions 3 and 5$]$ it holds that

$$
\operatorname{Var}\left(\omega\left(\tau \mid \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)\right) \leq C_{w}\left\|\boldsymbol{\theta}_{1}-\boldsymbol{\theta}_{2}\right\|^{2}
$$

where $C_{w}=H\left(2 H G^{2}+M\right)(W+1)$. Furthermore, we have

$$
\begin{aligned}
& \mathbb{E}_{\tau_{n}^{t}}\left\|1-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2} \\
& \quad=\operatorname{Var}_{\boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}}\left(\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right)\right) \\
& \quad \leq C_{w}\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}
\end{aligned}
$$

Proof. The proof can be found in Xu et al. [19].
Lemma 10. For $X_{1}, X_{2} \in \mathbb{R}^{d}$, we have

$$
\left\|X_{1}+X_{2}\right\|^{2} \leq 2\left\|X_{1}\right\|^{2}+2\left\|X_{2}\right\|^{2}
$$

## Lemma 11.

$$
\mathbb{E}_{\tau_{n}^{t}}\left[\left\|v_{n}^{t}\right\|^{2}\right] \leq\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2}
$$

Proof. We follow the suggestion of Lei et al. [35] to set $b_{t}=1$ to deliver better theoretical results. However in our experiments, we do allow $b_{t}$ to be sampled from different values. With $b_{t}=1$ and $\mu_{t}=\frac{1}{\left|\mathcal{G}_{t}\right|} \sum_{k \in \mathcal{G}_{t}} \mu_{t}^{(k)}$, we have the flowing definition according to Algorithm 1 .

$$
\begin{equation*}
v_{n}^{t} \triangleq g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)+u_{t} \tag{11-12}
\end{equation*}
$$

which is the SCSG update step. Define $e_{t} \triangleq u_{t}-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)$, we then have

$$
\begin{align*}
\mathbb{E}_{\tau_{n}^{t}}\left[v_{n}^{t}\right] & =\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)+e_{t}+\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right) \\
& =\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)+e_{t} \tag{11.13}
\end{align*}
$$

Note that $\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)=\mathbb{E}_{\tau_{n}^{t}}\left[g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right]$ as we have showed that the importance weighting term results in unbiased estimation of the true gradient in Lemma 8 . Then from $\mathbb{E}\|X\|^{2}=\mathbb{E}\|X-\mathbb{E} X\|^{2}+\|\mathbb{E} X\|^{2}$,

$$
\begin{align*}
& \mathbb{E}_{\tau_{n}^{t}}\left[\left\|v_{n}^{t}\right\|^{2}\right]=\mathbb{E}_{\tau_{n}^{t}}\left\|v_{n}^{t}-\mathbb{E}_{\tau_{n}^{t}}\left[v_{n}^{t}\right]\right\|^{2}+\left\|\mathbb{E}_{\tau_{n}^{t}}\left[v_{n}^{t}\right]\right\|^{2} \\
& =\mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)+u_{t}-\left(\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)+e_{t}\right)\right\|^{2}+\left\|\mathbb{E}_{\tau_{n}^{t}}\left[v_{n}^{t}\right]\right\|^{2} \\
& =\mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)-\left(\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)-\nabla J\left(\boldsymbol{\theta}_{0}^{t}\right)\right)\right\|^{2}+\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)+e_{t}\right\|^{2} \\
& \leq \mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2} \tag{11.14}
\end{align*}
$$

where (11-14) follows from $\mathbb{E}\|X-\mathbb{E} X\|^{2} \leq \mathbb{E}\|X\|^{2}$ and Lemma 10 . Note that

$$
\begin{align*}
& \mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2} \\
= & \mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)+g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)-g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2} \\
= & \mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)+\left(1-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right)\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2} \\
\leq & 2 \mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2}+2 \mathbb{E}_{\tau_{n}^{t}}\left\|\left(1-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right)\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2} \tag{11-15}
\end{align*}
$$

where 11-15) follows from Lemma 10 Combining 11F14) and 11F15, we have

$$
\begin{align*}
& \mathbb{E}_{\tau_{n}^{t}}\left[\left\|v_{n}^{t}\right\|^{2}\right] \leq 2 \mathbb{E}_{\tau_{n}^{t}}\left\|g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}\right)-g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2} \\
& \quad+2 \mathbb{E}_{\tau_{n}^{t}}\left\|\left(1-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right)\right) g\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{0}^{t}\right)\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2} \\
& \leq 2 L_{g}\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2 C_{g}^{2} \mathbb{E}_{\tau_{n}^{t}}\left\|\left(1-\omega\left(\tau_{n}^{t} \mid \boldsymbol{\theta}_{n}^{t}, \boldsymbol{\theta}_{0}^{t}\right)\right)\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2}  \tag{11.16}\\
& \leq 2 L_{g}\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2 C_{g}^{2} C_{w}\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2} \\
&=\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2} \tag{11.18}
\end{align*}
$$

where (11-16) is from Lemma 4 and (11-17) follows from Lemma 9

## Lemma 12.

$$
\eta_{t} \mathbb{E}\left\langle e_{t}, \mathbb{E} \nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\rangle=\frac{1}{B_{t}} \mathbb{E}\left\langle e_{t}, \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle-\eta_{t} \mathbb{E}\left\|e_{t}\right\|^{2}
$$

Proof. Consider $M_{n}^{t}=\left\langle e_{t}, \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle$. We have

$$
M_{n+1}^{t}-M_{n}^{t}=\left\langle e_{t}, \boldsymbol{\theta}_{n+1}^{t}-\boldsymbol{\theta}_{n}^{t}\right\rangle=\eta_{t}\left\langle e_{t}, v_{n}^{t}\right\rangle
$$

Taking expectation with respect to $\tau_{n}^{t}$, we have

$$
\begin{aligned}
\mathbb{E}_{\tau_{n}^{t}}\left[M_{n+1}^{t}-M_{n}^{t}\right] & =\eta_{t}\left\langle e_{t}, \mathbb{E}_{\tau_{n}^{t}}\left[v_{n}^{t}\right]\right\rangle \\
& =\eta_{t}\left\langle e_{t}, \nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\rangle+\eta_{t}\left\|e_{t}\right\|^{2}
\end{aligned}
$$

following from (11-13). Use $\mathbb{E}_{t}$ to denote the expectation with respect to all trajectories $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$, given $N_{t}$. Since $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$ are independent of $N_{t}, \mathbb{E}_{t}$ is equivalently the expectation with respect to $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$. We have

$$
\mathbb{E}_{t}\left[M_{n+1}^{t}-M_{n}^{t}\right]=\eta_{t}\left\langle e_{t}, \mathbb{E}_{t} \nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\rangle+\eta_{t}\left\|e_{t}\right\|^{2}
$$

Taking $n=N_{t}$ and denoting $\mathbb{E}_{N_{t}}$ the expectation w.r.t. $N_{t}$, we have

$$
\mathbb{E}_{N_{t}} \mathbb{E}_{t}\left(M_{N_{t}+1}^{t}-M_{N_{t}}^{t}\right)=\eta_{t}\left\langle e_{t}, \mathbb{E}_{N_{t}} \mathbb{E}_{t} \nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\rangle+\eta_{t}\left\|e_{t}\right\|^{2}
$$

Using Fubini's theorem, Lemma 16 and using the fact $\boldsymbol{\theta}_{N_{t}}^{t}=\tilde{\boldsymbol{\theta}}_{t}$ and $\boldsymbol{\theta}_{0}^{t}=\tilde{\boldsymbol{\theta}}_{t-1}$,

$$
\begin{aligned}
\mathbb{E}_{N_{t}} \mathbb{E}_{t}\left(M_{N_{t}+1}^{t}-M_{N_{t}}^{t}\right) & =-\mathbb{E}_{t} \mathbb{E}_{N_{t}}\left(M_{N_{t}}^{t}-M_{N_{t}+1}^{t}\right) \\
& =-\left(\frac{1}{B_{t} /\left(B_{t}+1\right)}-1\right)\left(M_{0}^{t}-\mathbb{E}_{N_{t}} \mathbb{E}_{t} M_{N_{t}}^{t}\right) \\
& =\frac{1}{B_{t}} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\langle e_{t}, \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle \\
& =\eta_{t}\left\langle e_{t}, \mathbb{E}_{N_{t}} \mathbb{E}_{t} \nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\rangle+\eta_{t}\left\|e_{t}\right\|^{2}
\end{aligned}
$$

Taking expectation with respect to the whole past yields the lemma.

## Lemma 13.

$$
\begin{array}{r}
-2 \eta_{t} \mathbb{E}\left\langle e_{t}, \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle \leq\left[-\frac{1}{B_{t}}+\eta_{t}^{2}\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\right] \mathbb{E}\left\|\tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\|^{2}+2 \eta_{t}^{2} \mathbb{E}\left\|e_{t}\right\|^{2} \\
+2 \eta_{t} \mathbb{E}\left\langle\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}, \tilde{\boldsymbol{\theta}}_{t}-\tilde{\boldsymbol{\theta}}_{t-1}\right\rangle+2 \eta_{t}^{2} \mathbb{E}\left\|\nabla J\left(\tilde{\boldsymbol{\theta}}_{t}\right)\right\|^{2}\right.
\end{array}
$$

Proof. We have from the update equation $\boldsymbol{\theta}_{n+1}^{t}=\boldsymbol{\theta}_{n}^{t}+\eta_{t} v_{n}^{t}$, then,

$$
\begin{align*}
\mathbb{E}_{\tau_{n}^{t}}\left\|\boldsymbol{\theta}_{n+1}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}= & \mathbb{E}_{\tau_{n}^{t}}\left\|\boldsymbol{\theta}_{n}^{t}+\eta_{t} v_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2} \\
= & \left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+\eta_{t}^{2} \mathbb{E}_{\tau_{n}^{t}}\left\|v_{n}^{t}\right\|^{2}+2 \eta_{t}\left\langle\mathbb{E}_{\tau_{n}^{t}}\left[v_{n}^{t}\right], \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle \\
\leq & \left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+\eta_{t}^{2}\left[\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2\left\|e_{t}\right\|^{2}\right] \\
& \quad+2 \eta_{t}\left\langle e_{t}, \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle+2 \eta_{t}\left\langle\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right), \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle  \tag{13-19}\\
= & {\left[1+\eta_{t}^{2}\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\right]\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2 \eta_{t}\left\langle\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right), \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle } \\
& \quad+2 \eta_{t}\left\langle e_{t}, \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle+2 \eta_{t}^{2}\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2 \eta_{t}^{2}\left\|e_{t}\right\|^{2}
\end{align*}
$$

where $13-19$ follows the result of $11-18$. Use $\mathbb{E}_{t}$ to denote the expectation with respect to all trajectories $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$, given $N_{t}$. Since $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$ are independent of $N_{t}, \mathbb{E}_{t}$ is equivalently the expectation with respect to $\left\{\tau_{1}^{t}, \tau_{2}^{t}, \ldots\right\}$. We have

$$
\begin{aligned}
\mathbb{E}_{t}\left\|\boldsymbol{\theta}_{n+1}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2} \leq\left[1+\eta_{t}^{2}\left(2 L_{g}\right.\right. & \left.\left.+2 C_{g}^{2} C_{w}\right)\right] \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}+2 \eta_{t} \mathbb{E}_{t}\left\langle\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right), \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle \\
& +2 \eta_{t} \mathbb{E}_{t}\left\langle e_{t}, \boldsymbol{\theta}_{n}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle+2 \eta_{t}^{2} \mathbb{E}_{t}\left\|\nabla J\left(\boldsymbol{\theta}_{n}^{t}\right)\right\|^{2}+2 \eta_{t}^{2}\left\|e_{t}\right\|^{2}
\end{aligned}
$$

Now taking $n=N_{t}$ and denoting $E_{N_{t}}$ the expectation w.r.t. $N_{t}$ we have

$$
\begin{align*}
& -2 \eta_{t} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\langle e_{t}, \boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle \\
& \leq\left[1+\eta_{t}^{2}\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\right] \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2}-\mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{N_{t}+1}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2} \\
& \quad+2 \eta_{t} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\langle\nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right), \boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle+2 \eta_{t}^{2} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\|^{2}+2 \eta_{t}^{2}\left\|e_{t}\right\|^{2} \\
& =\left[-\frac{1}{B_{t}}+\eta_{t}^{2}\left(2 L_{g}+2 C_{g}^{2} C_{w}\right)\right] \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\|^{2} \\
& \quad+2 \eta_{t} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\langle\nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right), \boldsymbol{\theta}_{N_{t}}^{t}-\boldsymbol{\theta}_{0}^{t}\right\rangle+2 \eta_{t}^{2} \mathbb{E}_{N_{t}} \mathbb{E}_{t}\left\|\nabla J\left(\boldsymbol{\theta}_{N_{t}}^{t}\right)\right\|^{2}+2 \eta_{t}^{2}\left\|e_{t}\right\|^{2} \tag{13-20}
\end{align*}
$$

where (13-20) follows Lemma 16 using Fubini's theorem. Rearranging, replacing $\boldsymbol{\theta}_{N_{t}}^{t}=\tilde{\boldsymbol{\theta}}_{t}$ and $\boldsymbol{\theta}_{0}^{t}=\tilde{\boldsymbol{\theta}}_{t-1}$ and taking expectation w.r.t the whole past yields the lemma.
Lemma 14 (Pinelis' inequality [76]; Lemma 2.4 [46]). Let the sequence of random variables $X_{1}, X_{2}, \ldots, X_{N} \in \mathbb{R}^{d}$ represent a random process such that we have $\mathbb{E}\left[X_{n} \mid X_{1}, \ldots, X_{n-1}\right]$ and $\left\|X_{n}\right\| \leq M$. Then,

$$
\mathbb{P}\left[\left\|X_{1}+\ldots+X_{N}\right\|^{2} \leq 2 \log (2 / \delta) M^{2} N\right] \geq 1-\delta
$$

Lemma 15 (Adapted from [48]). If we choose $\delta$ and $B_{t}$ in Algorithm 1 such that:
(i) $\mathrm{e}^{\frac{\delta B_{t}}{2(1-2 \delta)}} \leq \frac{2 K}{\delta} \leq \mathrm{e}^{\frac{B_{t}}{2}}$
(ii) $\delta \leq \frac{1}{5 K B_{t}}$
then we have the following bound for $\mathbb{E}\left\|e_{t}\right\|^{2}$ :

$$
\mathbb{E}\left\|e_{t}\right\|^{2} \leq \frac{4 \sigma^{2}}{(1-\alpha)^{2} K B_{t}}+\frac{48 \alpha^{2} \sigma^{2} V}{(1-\alpha)^{2} B_{t}}
$$

Proof. The proof of this lemma is similar to that of Lemma 7 of Khanduri et al. [48]. The key difference lays on the base conditions used to define the probabilistic events.
In FedPG-BR, the following refined conditions (results of Claims D.1 and D.2) are used,

$$
\left\|\mu_{t}^{\mathrm{mom}}-\nabla J(\boldsymbol{\theta})\right\| \leq \sigma,\left\|\mu_{t}^{(k)}-\mu_{t}^{\mathrm{mom}}\right\| \leq 2 \sigma,\left\|\mu_{t}^{(k)}-\nabla J(\boldsymbol{\theta})\right\| \leq 3 \sigma, \forall k \in \mathcal{G}_{t}
$$

whereas Khanduri et al. [48] needs the following:

$$
\left\|\mu_{t}^{\mathrm{med}}-\nabla J(\boldsymbol{\theta})\right\| \leq 3 \sigma,\left\|\mu_{t}^{(k)}-\mu_{t}^{\mathrm{med}}\right\| \leq 4 \sigma,\left\|\mu_{t}^{(k)}-\nabla J(\boldsymbol{\theta})\right\| \leq 7 \sigma, \forall k \in \mathcal{G}_{t}
$$

The detailed proof of Lemma 15 can be obtained following the derivation of Lemma 7 of Khanduri et al. [48] by modifying the base conditions.

Lemma 16. If $N \sim \operatorname{Geom}(\Gamma)$ for $\Gamma>0$. Then for any sequence $D_{0}, D_{1}, \ldots$ with $\mathbb{E}\left\|D_{N}\right\| \leq \infty$, we have

$$
\mathbb{E}\left[D_{N}-D_{N+1}\right]=\left(\frac{1}{\Gamma}-1\right)\left(D_{0}-\mathbb{E} D_{N}\right)
$$

Proof. The proof can be found in Lei et al. [35].
Lemma 17 (Young's inequality (Peter-Paul inequality)). For all real numbers $a$ and $b$ and all $\beta>0$, we have

$$
a b \leq \frac{a^{2}}{2 \beta}+\frac{\beta b^{2}}{2}
$$

## F Experimental details

## F. 1 Hyperparameters

We follow the setups of SVRPG [18] to parameterize the policies using neural networks. For all the algorithms under comparison in the experiments (Section 5]), Adam[77] is used as the gradient optimizer. The 10 random seeds are $[0-9]$. All other hyperparameters used in all the experiments are reported in Table 2

Table 2: Hyperparameters used in the experiments.

| Hyperparameters | Algorithms | CartPole-v1 | LunarLander-v2 | HalfCheetach-v2 |
| :---: | :---: | :---: | :---: | :---: |
| NN policy | - | Categorical MLP | Categorical MLP | Gaussian MLP |
| NN hidden weights | - | 16,16 | 64,64 | 64,64 |
| NN activation | - | ReLU | Tanh | Tanh |
| NN output activation | - | Tanh | Tanh | Tanh |
| Step size (Adam) $\eta$ | - | $1 \mathrm{e}-3$ | $1 \mathrm{e}-3$ | $8 \mathrm{e}-5$ |
| Discount factor $\gamma$ | - | 0.999 | 0.990 | 0.995 |
| Maximum trajectories | - | 5000 | 10000 | 10000 |
| Task horizon $H$ (for training) | - | 500 | 1000 | 500 |
| Task horizon $H$ (for test) | - | 500 | 1000 | 1000 |
| $\alpha$ (for practical setup) | - | 0.3 | 0.3 | 0.3 |
| Number of runs | - | 10 | 10 | 10 |
| Batch size $B_{t}$ | GPOMDP | 16 | 32 | 48 |
|  | SVRPG | 16 | 32 | 48 |
|  | FedPG-BR | sampled from [12, 20] | sampled from [26, 38] | sampled from [46, 50] |
| Mini-Batch size $b_{t}$ | GPOMDP | - | - | - |
|  | SVRPG | 4 | 8 | 16 |
|  | FedPG-BR | 4 | 8 | 16 |
| Number of steps $N_{t}$ | GPOMDP | 1 | 1 | 1 |
|  | SVRPG | 3 | 3 | 3 |
|  | FedPG-BR | $N_{t} \sim \operatorname{Geom}\left(\frac{B_{t}}{B_{t}+b_{t}}\right)$ | $N_{t} \sim \operatorname{Geom}\left(\frac{B_{t}}{B_{t}+b_{t}}\right)$ | $N_{t} \sim \operatorname{Geom}\left(\frac{B_{t}}{B_{t}+b_{t}}\right)$ |
| Variance bound $\sigma$ (Estimated by server) | GPOMDP | - | 碞 | 碞 |
|  | SVRPG | - | - | - |
|  | FedPG-BR | 0.06 | 0.07 | 0.9 |
| Confidence parameter $\delta$ | GPOMDP | - | - | - |
|  | SVRPG | - | - | - |
|  | FedPG-BR | 0.6 | 0.6 | 0.6 |



Figure 5: Performance of FedPG-BR in practical systems with $\alpha>0$ for CartPole. Each subplot corresponds to a different type of Byzantine failure exercised by the 3 Byzantine agents.


Figure 6: Performance of FedPG-BR in practical systems with $\alpha>0$ for LunarLander. Each subplot corresponds to a different type of Byzantine failure exercised by the 3 Byzantine agents.

## F. 2 Computing Infrastructure

All experiments are conducted on a computing server without GPUs. The server is equipped with 14 cores (28 threads) Intel(R) Core(TM) i9-10940X CPU @ 3.30 GHz and 64 G memory. The average runtime for each run of FedPG-BR $(\mathrm{K}=10 \mathrm{~B}=3)$ is 2.5 hours for the CartPole task, 4 hours for the HalfCheetah task, and 12 hours for the LunarLander task.

## G Additional experiments

## G. 1 Performance of FedPG-BR in practical systems with $\alpha>0$ for the CartPole and the LunarLander tasks

The results for the CartPole and the LunarLander tasks which yield the same insights as discussed in experiments (Section 5) are plotted in Figure 5 and Figure 6 As discussed earlier, for both GPOMDP and SVRPG, the federation of more agents in practical systems which are subject to the presence of Byzantine agents, i.e., random failures or adversarial attacks, causes the performance of their federation to be worse than that in the single-agent setting. In particular, RA agents (middle figure) and SF agents (right figure) render GPOMDP and SVRPG unlearnable, i.e., unable to converge at all. This is in contrast to the performance of FedPG-BR. That is, FedPG-BR $(K=10 B=3)$ is able to deliver superior performances even in the presence of Byzantine agents for all three tasks: CartPole (Figure 5), LunarLander (Figure 6), and HalfCheetah (Figure 2 in Section 55). This provides an assurance on the reliability of our FedPG-BR algorithm to promote its practical deployment, and significantly improves the practicality of FRL.

## G. 2 Performance of FedPG-BR against the Variance Attack

We have discussed in Section 3.2 where the high variance in PG estimation renders the FRL system vunlnerable to variance-based attacks such as the Variance Attack (VA) proposed by Baruch et al. [47]. The VA attackers collude together to estimate the population mean and the standard-deviation of gradients at each round, and move the mean by the largest value such that their values are still within the population variance. Intuitively, this non-omniscient attack works by exploiting the high variance in gradient estimation of the population and crafting values that contribute most to the population variance, hence gradually shifting the population mean. According to Cao et al. [20],


Figure 7: Performance of FedPG-BR in practical systems with $\alpha>0$ for CartPole. Among the $K=10$ participating agents, 3 Byzantine agents are colluding together to launch the VA attack.
existing defenses will fail to remove those non-omniscient attackers and the convergence will be significantly worsened if the population variance is large enough.

We are thus motivated to look for solutions that theoretically reduce the variance in policy gradient estimation. Inspired by the variance-reduced policy gradient works [e.g., 18, 19], we adapt the SCSG optimization [35] to our federated policy gradient framework for a refined control over the estimation variance. Through our adaptation, we are able to control the variance by the semi-stochastic gradient (line 11 in Algorithm 1), hence resulting in the fault-tolerant FRL system that can defend the VA attackers. Each plot in Figure 7 shows the experiment for each of the three tasks correspondingly, where 3 Byzantine agents are implemented as the VA attackers [20] ( $z^{\max }$ is 0.18 in our setup). We again include the corresponding single-agent performance ( $K=1$ ) and the federation of 10 good agents $(K=10)$ in the plots for reference. The results show that in all three tasks, FedPG-BR ( $K=10 B=3$ ) still manages to significantly outperform FedPG-BR $(K=1)$ in the single-agent setting. Furthermore, the performance of FedPG-BR $(K=10 B=3)$ is barely worsened compared with FedPG-BR $(K=10)$ with 10 good agents. This shows that, with the adaptation of SCSG, our fault-tolerant FRL system can perfectly defend the VA attack from the literature, which further corroborates our analysis on our Byzantine filtering step (Section 3.3) showing that if gradients from Byzantine agents are not filtered out, their impact is limited since their maximum distance to $\nabla J\left(\boldsymbol{\theta}_{0}^{\boldsymbol{t}}\right)$ is bounded by $3 \sigma$ (ClaimD.2).

