A Regret Bounds

The following lemma bounds the expected per-round regret of any randomized algorithm that chooses the perturbed solution in round t, $\tilde{\theta}_t$, as a function of the history.

Lemma 2. *Let* $p_2 \geq \mathbb{P}_t\left(\bar{E}_{2,t}\right)$, $p_3 \leq \mathbb{P}_t\left(E_{3,t}\right)$, *and* $p_3 > p_2$. *Then on event* $E_{1,t}$,

$$\mathbb{E}_t\left[\Delta_{I_t}\right] \leq \dot{\mu}_{\max}(c_1 + c_2)\left(1 + \frac{2}{p_3 - p_2}\right) \times \\ \mathbb{E}_t\left[\|x_{I_t}\|_{G_t^{-1}}\right] + \Delta_{\max}\, p_2\,.$$

Proof. Let $\tilde{\Delta}_i = x_1^\top \theta_* - x_i^\top \theta_*$ and $c = c_1 + c_2$. Let

$$\bar{S}_t = \left\{i \in [K] : c\|x_i\|_{G_t^{-1}} \geq \tilde{\Delta}_i\right\}$$

be the set of *undersampled arms* in round t. Note that $1 \in \bar{S}_t$ by definition. We define the set of *sufficiently sampled arms* as $S_t = [K] \setminus \bar{S}_t$. Let $J_t = \arg\min_{i \in \bar{S}_t} \|x_i\|_{G_t^{-1}}$ be the *least uncertain undersampled arm* in round t.

In all steps below, we assume that event $E_{1,t}$ occurs. In round t on event $E_{2,t}$,

$$\begin{aligned}\Delta_{I_t} &\leq \dot{\mu}_{\max}\,\tilde{\Delta}_{I_t} = \dot{\mu}_{\max}\left(\tilde{\Delta}_{J_t} + x_{J_t}^\top \theta_* - x_{I_t}^\top \theta_*\right) \leq \dot{\mu}_{\max}\left(\tilde{\Delta}_{J_t} + x_{J_t}^\top \tilde{\theta}_t - x_{I_t}^\top \tilde{\theta}_t + c\left(\|x_{I_t}\|_{G_t^{-1}} + \|x_{J_t}\|_{G_t^{-1}}\right)\right) \\ &\leq \dot{\mu}_{\max}\, c\left(\|x_{I_t}\|_{G_t^{-1}} + 2\|x_{J_t}\|_{G_t^{-1}}\right)\,,\end{aligned}$$

where the first inequality holds because $\dot{\mu}_{\max}$ is the maximum derivative of μ, the second is by the definitions of events $E_{1,t}$ and $E_{2,t}$, and the last follows from the definitions of I_t and J_t. Now we take the expectation of both sides and get

$$\mathbb{E}_t\left[\Delta_{I_t}\right] = \mathbb{E}_t\left[\Delta_{I_t}\mathbb{1}\{E_{2,t}\}\right] + \mathbb{E}_t\left[\Delta_{I_t}\mathbb{1}\{\bar{E}_{2,t}\}\right] \leq \dot{\mu}_{\max}\, c\,\mathbb{E}_t\left[\|x_{I_t}\|_{G_t^{-1}} + 2\|x_{J_t}\|_{G_t^{-1}}\right] + \Delta_{\max}\,\mathbb{P}_t\left(\bar{E}_{2,t}\right)\,.$$

The last step is to replace $\mathbb{E}_t\left[\|x_{J_t}\|_{G_t^{-1}}\right]$ with $\mathbb{E}_t\left[\|x_{I_t}\|_{G_t^{-1}}\right]$. To do so, observe that

$$\mathbb{E}_t\left[\|x_{I_t}\|_{G_t^{-1}}\right] \geq \mathbb{E}_t\left[\|x_{I_t}\|_{G_t^{-1}} \,\middle|\, I_t \in \bar{S}_t\right]\mathbb{P}_t\left(I_t \in \bar{S}_t\right) \geq \|x_{J_t}\|_{G_t^{-1}}\,\mathbb{P}_t\left(I_t \in \bar{S}_t\right)\,,$$

where the last inequality follows from the definition of J_t and that $\bar{S}_t$ is $\mathcal{F}_{t-1}$-measurable. We rearrange the inequality as $\|x_{J_t}\|_{G_t^{-1}} \leq \mathbb{E}_t\left[\|x_{I_t}\|_{G_t^{-1}}\right] / \,\mathbb{P}_t\left(I_t \in \bar{S}_t\right)$ and bound $\mathbb{P}_t\left(I_t \in \bar{S}_t\right)$ from below next.

In particular, on event $E_{1,t}$,

$$\begin{aligned}\mathbb{P}_t\left(I_t \in \bar{S}_t\right) &\geq \mathbb{P}_t\left(\exists i \in \bar{S}_t : x_i^\top \tilde{\theta}_t > \max_{j \in S_t} x_j^\top \tilde{\theta}_t\right) \geq \mathbb{P}_t\left(x_1^\top \tilde{\theta}_t > \max_{j \in S_t} x_j^\top \tilde{\theta}_t\right) \\ &\geq \mathbb{P}_t\left(x_1^\top \tilde{\theta}_t > \max_{j \in S_t} x_j^\top \tilde{\theta}_t,\, E_{2,t} \text{ occurs}\right) \geq \mathbb{P}_t\left(x_1^\top \tilde{\theta}_t > x_1^\top \theta_*,\, E_{2,t} \text{ occurs}\right) \\ &\geq \mathbb{P}_t\left(x_1^\top \tilde{\theta}_t > x_1^\top \theta_*\right) - \mathbb{P}_t\left(\bar{E}_{2,t}\right) \geq \mathbb{P}_t\left(x_1^\top \tilde{\theta}_t - x_1^\top \bar{\theta}_t > c_1\|x_1\|_{G_t^{-1}}\right) - \mathbb{P}_t\left(\bar{E}_{2,t}\right)\,.\end{aligned}$$

Note that we require a sharp inequality because $I_t \in \bar{S}_t$ is not guaranteed on event $\left\{\exists i \in \bar{S}_t : x_i^\top \tilde{\theta}_t \geq \max_{j \in S_t} x_j^\top \tilde{\theta}_t\right\}$. The fourth inequality holds because on event $E_{1,t} \cap E_{2,t}$,

$$x_j^\top \tilde{\theta}_t \leq x_j^\top \theta_* + c\|x_j\|_{G_t^{-1}} < x_j^\top \theta_* + \tilde{\Delta}_j = x_1^\top \theta_*$$

holds for any $j \in S_t$. The last inequality holds because $x_1^\top \theta_* \leq x_1^\top \bar{\theta}_t + c_1\|x_1\|_{G_t^{-1}}$ holds on event $E_{1,t}$. Finally, we use the definitions of p_2 and p_3 to complete the proof. □

The regret bound of `GLM-TSL` is proved below.

Theorem 3. *The n-round regret of* GLM-TSL *is bounded as*

$$R(n) \leq \dot{\mu}_{\max}(c_1 + c_2)\left(1 + \frac{2}{0.15 - 1/n}\right) \times \sqrt{2dn\log(2n/d)} + (\tau + 3)\Delta_{\max}\,,$$

where

$$\begin{aligned} a &= c_1\sqrt{\dot{\mu}_{\max}}\,, \\ c_1 &= \sigma\dot{\mu}_{\min}^{-1}\sqrt{d\log(n/d) + 2\log n}\,, \\ c_2 &= c_1\sqrt{2\dot{\mu}_{\min}^{-1}\,\dot{\mu}_{\max}\log(Kn)}\,, \end{aligned}$$

and the number of exploration rounds τ satisfies

$$\lambda_{\min}(G_\tau) \geq \max\left\{\sigma^2\dot{\mu}_{\min}^{-2}(d\log(n/d) + 2\log n),\, 1\right\}\,.$$

Proof. Fix $\tau \in [n]$. Let

$$E_{4,t} = \left\{\|\bar{\theta}_t - \theta_*\|_2 \leq 1\right\}$$

and $p_4 \geq \mathbb{P}\left(\bar{E}_{4,t}\right)$ for $t \geq \tau$. Let $p_1 \geq \mathbb{P}\left(\bar{E}_{1,t}, E_{4,t}\right)$, $p_2 \geq \mathbb{P}_t\left(\bar{E}_{2,t}\right)$ on event $E_{4,t}$, and $p_3 \leq \mathbb{P}_t\left(E_{3,t}\right)$. By elementary algebra, we get

$$\begin{aligned} R(n) &\leq \sum_{t=\tau}^{n}\mathbb{E}\left[\Delta_{I_t}\right] + \tau\Delta_{\max} \\ &\leq \sum_{t=\tau}^{n}\mathbb{E}\left[\Delta_{I_t}\mathbb{1}\{E_{4,t}\}\right] + (\tau + p_4 n)\Delta_{\max} \\ &\leq \sum_{t=\tau}^{n}\mathbb{E}\left[\Delta_{I_t}\mathbb{1}\{E_{1,t}, E_{4,t}\}\right] + (\tau + (p_1 + p_4)n)\Delta_{\max} \\ &= \sum_{t=\tau}^{n}\mathbb{E}\left[\mathbb{E}_t\left[\Delta_{I_t}\right]\mathbb{1}\{E_{1,t}, E_{4,t}\}\right] + (\tau + (p_1 + p_4)n)\Delta_{\max}\,. \end{aligned}$$

To get $p_1 \leq 1/n$, we set c_1 as in Lemma 8. Now we apply Lemma 2 to $\mathbb{E}_t\left[\Delta_{I_t}\right]\mathbb{1}\{E_{1,t}, E_{4,t}\}$ and get

$$R(n) \leq \dot{\mu}_{\max}(c_1 + c_2)\left(1 + \frac{2}{p_3 - p_2}\right)\mathbb{E}\left[\sum_{t=\tau}^{n}\|x_{I_t}\|_{G_t^{-1}}\right] + (\tau + (p_1 + p_2 + p_4)n)\Delta_{\max}\,,$$

where a and c_2 are set as in Lemma 4. For these settings, $p_2 \leq 1/n$ and $p_3 \geq 0.15$. To bound $\sum_{t=\tau}^{n}\|x_{I_t}\|_{G_t^{-1}}$, we use Lemma 2 in Li et al. [2017]. Finally, to get $p_4 \leq 1/n$, we choose τ as in Lemma 9. □

The regret bound of GLM-FPL is proved below.

Theorem 5. *The n-round regret of* GLM-FPL *is bounded as*

$$R(n) \leq \dot{\mu}_{\max}(c_1 + c_2)\left(1 + \frac{2}{0.15 - 2/n}\right) \times \sqrt{2dn\log(2n/d)} + (\tau + 4)\Delta_{\max}\,,$$

where

$$\begin{aligned} a &= c_1\dot{\mu}_{\max}\,, \\ c_1 &= \sigma\dot{\mu}_{\min}^{-1}\sqrt{d\log(n/d) + 2\log n}\,, \\ c_2 &= c_1\dot{\mu}_{\min}^{-1}\,\dot{\mu}_{\max}\sqrt{2\log(Kn)}\,, \end{aligned}$$

and the number of exploration rounds τ satisfies

$$\lambda_{\min}(G_\tau) \geq \max\{4\sigma^2 \dot{\mu}_{\min}^{-2}(d \log(n/d) + 2 \log n), \\ 8a^2 \dot{\mu}_{\min}^{-2} \log n,\ 1\}\,.$$

Proof. The proof is almost identical to that of Theorem 3. There are two main differences. First, a and c_2 are set as in Lemma 6. For these settings, $p_2 \leq 2/n$ and $p_3 \geq 0.15$. Second, τ is set as in Lemma 10. ☐

B Technical Lemmas

We need an extension of Theorem 1 in Abbasi-Yadkori et al. [2011], which is concerned with concentration of a certain vector-valued martingale. The setup of the claim is as follows. Let $(\mathcal{F}_t)_{t \geq 0}$ be a filtration, $(\eta_t)_{t \geq 1}$ be a stochastic process such that η_t is real-valued and $\mathcal{F}_t$-measurable, and $(X_t)_{t \geq 1}$ be another stochastic process such that X_t is $\mathbb{R}^d$-valued and $\mathcal{F}_{t-1}$-measurable. We also assume that $(\eta_t)_t$ is conditionally R^2-sub-Gaussian, that is

$$\forall \lambda \in \mathbb{R}: \quad \mathbb{E}\left[\exp[\lambda \eta_t] \,|\, \mathcal{F}_{t-1}\right] \leq \exp\left[\frac{\lambda^2 R^2}{2}\right]. \tag{10}$$

We call the triplet $((X_t)_t, (\eta_t)_t, \mathbb{F})$ "nice" when these conditions hold. The modified claim is stated and proved below.

Lemma 7. *Let $((X_t)_t, (\eta_t)_t, \mathbb{F})$ be a "nice" triplet, $S_t = \sum_{s=1}^t \eta_s X_s$, $V_t = \sum_{s=1}^t X_s X_s^\top$; and for $V \succ 0$, let $\tau_0 = \min\{t \geq 1 : V_t \succeq V\}$. Then for any $\delta \in (0,1)$ and $\mathbb{F}$-stopping time $\tau \geq 1$ such that $\tau \geq \tau_0$ holds almost surely, with probability at least $1 - \delta$,*

$$\|S_\tau\|_{V_\tau^{-1}}^2 \leq 2R^2 \log\left(\frac{\det(V_\tau)^{\frac{1}{2}} \det(V_{\tau_0})^{-\frac{1}{2}}}{\delta}\right).$$

Proof. The proof in Abbasi-Yadkori et al. [2011] can easily adjusted as follows. If $((X_t)_t, (\eta_t)_t, \mathbb{F})$ is a "nice" triplet, then for any $\delta \in (0,1)$, $\mathcal{F}_0$-measurable matrix $V \succ 0$, and stopping time $\tau \geq 1$,

$$\mathbb{P}\left(\|S_\tau\|_{V_\tau^{-1}}^2 \leq 2R^2 \log\left(\frac{\det(V_\tau)^{\frac{1}{2}} \det(V_{\tau_0})^{-\frac{1}{2}}}{\delta}\right) \,\middle|\, \mathcal{F}_0\right) \geq 1 - \delta\,. \tag{11}$$

Now, for $t \geq 0$, let $X'_t = X_{\tau_0 + t}$, $\eta'_t = \eta_{\tau_0 + t}$, and $\mathcal{F}'_t = \mathcal{F}_{\tau_0 + t}$. Then $((X'_t)_{t \geq 1}, (\eta'_t)_{t \geq 1}, (\mathcal{F}'_t)_{t \geq 0})$ is a nice triplet and the result follows from (11). ☐

We use the last lemma to prove the following result.

Lemma 8. *Let $c_1 = \sigma \dot{\mu}_{\min}^{-1} \sqrt{d \log(n/d) + 2 \log n}$ and τ be any round such that $\lambda_{\min}(G_\tau) \geq 1$. Then for any $t \geq \tau$,*

$$\mathbb{P}\left(\bar{E}_{1,t} \text{ occurs},\ \|\bar{\theta}_t - \theta_*\|_2 \leq 1\right) \leq 1/n\,.$$

Proof. Let $S_t = \sum_{\ell=1}^{t-1} (Y_\ell - \mu(X_\ell^\top \theta_*)) X_\ell$. By Lemma 1, where $\mathcal{D}_1 = \left\{(X_\ell, \mu(X_\ell^\top \theta_*))\right\}_{\ell=1}^{t-1}$ and $\mathcal{D}_2 = \{(X_\ell, Y_\ell)\}_{\ell=1}^{t-1}$, we have that

$$S_t = \underbrace{\nabla^2 L(\mathcal{D}_1; \theta')}_{V} (\bar{\theta}_t - \theta_*)\,,$$

where $\theta' = \alpha \theta_* + (1 - \alpha)\bar{\theta}_t$ for $\alpha \in [0,1]$. We rearrange the equality as $V^{-1} S_t = \bar{\theta}_t - \theta_*$ and note that $\dot{\mu}_{\min} G_t \preceq V$ on $\|\bar{\theta}_t - \theta_*\|_2 \leq 1$. Now fix arm i. By the Cauchy-Schwarz inequality and from the above discussion,

$$\begin{aligned} \left|x_i^\top \bar{\theta}_t - x_i^\top \theta_*\right| &\leq \|\bar{\theta}_t - \theta_*\|_{G_t} \|x_i\|_{G_t^{-1}} = (\bar{\theta}_t - \theta_*)^\top G_t (\bar{\theta}_t - \theta_*) \|x_i\|_{G_t^{-1}} \\ &= S_t^\top V^{-1} G_t V^{-1} S_t \|x_i\|_{G_t^{-1}} \leq \dot{\mu}_{\min}^{-2} \|S_t\|_{G_t^{-1}} \|x_i\|_{G_t^{-1}}\,. \end{aligned}$$

By (13) in Lemma 9, which is derived using Lemma 7, $\|S_t\|_{G_t^{-1}} \leq \sigma \sqrt{d \log(n/d) + 2 \log n}$ holds with probability at least $1 - 1/n$ in any round $t \geq \tau$. In this case, event $E_{1,t}$ is guaranteed to occur when c_1 is set as in the claim. It follows that $\bar{E}_{1,t}$ occurs on $\|\bar{\theta}_t - \theta_*\|_2 \leq 1$ with probability of at most $1/n$. ☐

The number of initial exploration rounds in GLM-TSL is set below.

Lemma 9. *Let τ be any round such that*

$$\lambda_{\min}(G_\tau) \geq \max\left\{\sigma^2 \dot{\mu}_{\min}^{-2}(d \log(n/d) + 2 \log n),\, 1\right\}\,.$$

Then for any $t \geq \tau$, $\mathbb{P}\left(\|\bar{\theta}_t - \theta_*\|_2 > 1\right) \leq 1/n$.

Proof. Fix round t and let $S_t = \sum_{\ell=1}^{t-1}(Y_\ell - \mu(X_\ell^\top \theta_*))X_\ell$. By the same argument as in the proof of Theorem 1 in Li et al. [2017], who use Lemma A of Chen et al. [1999], we have that

$$\|S_t\|_{G_t^{-1}} \leq \dot{\mu}_{\min}\sqrt{\lambda_{\min}(G_t)} \implies \|\bar{\theta}_t - \theta_*\|_2 \leq 1$$

Now fix τ such that $\lambda_{\min}(G_\tau) \geq 1$. For any $t \geq \tau$, $G_t \succeq G_\tau$ and thus

$$\|S_t\|_{G_t^{-1}} \leq \dot{\mu}_{\min}\sqrt{\lambda_{\min}(G_\tau)} \implies \|\bar{\theta}_t - \theta_*\|_2 \leq 1\,. \tag{12}$$

In the next step, we bound $\|S_t\|_{G_t^{-1}}$ from above. By Lemma 7,

$$\|S_t\|_{G_t^{-1}}^2 \leq 2\sigma^2 \log(\det(G_t)^{\frac{1}{2}} \det(G_\tau)^{-\frac{1}{2}} n)$$

holds jointly in all rounds $t \geq \tau$ with probability at least $1 - 1/n$. By Lemma 11 in Abbasi-Yadkori et al. [2011] and from $\|X_t\|_2 \leq 1$, we get $\log\det(G_t) \leq d\log(n/d)$. By the choice of τ, $\det(G_\tau)^{-1} \leq 1$. It follows that

$$\|S_t\|_{G_t^{-1}}^2 \leq \sigma^2(d\log(n/d) + 2\log n) \tag{13}$$

for any $t \geq \tau$ with probability at least $1 - 1/n$. Now we combine this claim with (12) and have that $\|\bar{\theta}_t - \theta_*\|_2 \leq 1$ holds with probability at least $1 - 1/n$ when

$$\lambda_{\min}(G_\tau) \geq \sigma^2 \dot{\mu}_{\min}^{-2}(d\log(n/d) + 2\log n)\,.$$

This concludes the proof. □

The number of initial exploration rounds in GLM-FPL is set below.

Lemma 10. *Let τ be any round such that*

$$\lambda_{\min}(G_\tau) \geq \max\left\{4\sigma^2 \dot{\mu}_{\min}^{-2}(d \log(n/d) + 2 \log n),\, 8a^2 \dot{\mu}_{\min}^{-2} \log n,\, 1\right\}\,.$$

Then for any $t \geq \tau$, $\mathbb{P}\left(\|\bar{\theta}_t - \theta_*\|_2 > 1/2\right) \leq 1/n$ and $\mathbb{P}_t\left(\|\tilde{\theta}_t - \theta_*\|_2 > 1\right) \leq 1/n$ on event $\|\bar{\theta}_t - \theta_*\|_2 \leq 1/2$.

Proof. Fix round t. Let S_t be defined as in Lemma 9 and τ_1 be any round such that

$$\lambda_{\min}(G_{\tau_1}) \geq \min\left\{4\sigma^2 \dot{\mu}_{\min}^{-2}(d \log(n/d) + 2 \log n),\, 1\right\}\,.$$

Then by the same argument as in Lemma 9, $\mathbb{P}\left(\|\bar{\theta}_t - \theta_*\|_2 > 1/2\right) \leq 1/n$ holds for any $t \geq \tau_1$.

Now fix round t, history $\mathcal{F}_{t-1}$, and assume that $\|\bar{\theta}_t - \theta_*\|_2 \leq 1/2$ holds. Let

$$\bar{S}_t = \sum_{\ell=1}^{t-1}(Y_\ell + Z_\ell - \mu(X_\ell^\top \bar{\theta}_t))X_\ell = \sum_{\ell=1}^{t-1} Z_\ell X_\ell\,,$$

where the last equality holds because $\sum_{\ell=1}^{t-1}(Y_\ell - \mu(X_\ell^\top \bar{\theta}_t))X_\ell = \mathbf{0}$. Since $\|\bar{\theta}_t - \theta_*\|_2 \leq 1/2$, the 0.5-ball centered at $\bar{\theta}_t$ is within the unit ball centered at θ_*. So, the minimum derivative of μ in the 0.5-ball is not larger than that in the unit ball, and we have by a similar argument to Lemma 9 that

$$\|\bar{S}_t\|_{G_t^{-1}} \leq \frac{1}{2}\dot{\mu}_{\min}\sqrt{\lambda_{\min}(G_t)} \implies \|\tilde{\theta}_t - \bar{\theta}_t\|_2 \leq \frac{1}{2}\,. \tag{14}$$

By definition, $\|\bar{S}_t\|_{G_t^{-1}} = \|U\|_2$ for $U = G_t^{-\frac{1}{2}} \sum_{\ell=1}^{t-1} Z_\ell X_\ell$. Since Z_ℓ are i.i.d. random variables that are resampled in each round, we have $U \sim \mathcal{N}(\mathbf{0}, a^2 I_d)$ given $\mathcal{F}_{t-1}$, and that $\|U\|_2 \leq a\sqrt{2 \log n}$ holds with probability at least $1 - 1/n$ given $\mathcal{F}_{t-1}$. Now we combine this claim with (14) and have that $\|\tilde{\theta}_t - \bar{\theta}_t\|_2 \leq 1/2$ holds with probability at least $1 - 1/n$ for any round t such that

$$\lambda_{\min}(G_t) \geq 8 a^2 \dot{\mu}_{\min}^{-2} \log n.$$

For any such round, when $\|\bar{\theta}_t - \theta_*\|_2 \leq 1/2$ holds, $\mathbb{P}_t \left(\|\tilde{\theta}_t - \theta_*\|_2 \leq 1 \right) \geq 1 - 1/n$. This concludes our proof. □