The Knowledge Gradient for Sequential Decision Making with Stochastic Binary Feedbacks

This note contains supplementary materials to *The Knowledge Gradient for Sequential Decision Making with Stochastic Binary Feedbacks.*

A Proofs of theoretical statements

We provide detailed proofs of Lemma 1 and Theorem 1 in this section.

A.1 Proof of Lemma 1

Lemma 1 Let $\pi \in \Pi$ be a policy, and $\boldsymbol{x}^\pi = \arg\max_{\boldsymbol{x}} \Pr[y = +1|\boldsymbol{x}, \mathcal{D}^N]$ be the implementation decision after the budget N is exhausted. Then

$$\mathbb{E}[\Pr(y = +1|\boldsymbol{x}^\pi, \boldsymbol{w})] = \mathbb{E}[\max_{\boldsymbol{x}} \Pr(y = +1|\boldsymbol{x}, \mathcal{D}^N)],$$

where the expectation is taking over the prior distribution of $\boldsymbol{w}$.

Proof. First notice that for any fixed point $\boldsymbol{x}$,

$$\begin{aligned}
\mathbb{E}_N[\Pr(y = +1|\boldsymbol{x}, \boldsymbol{w})] &= \mathbb{E}_N[\sigma(\boldsymbol{w}^T\boldsymbol{x})] \\
&= \int \sigma(\boldsymbol{w}^T\boldsymbol{x})\Pr(\boldsymbol{w}|\mathcal{D}^N)\mathrm{d}\boldsymbol{w} \\
&= \Pr(y = +1|\boldsymbol{x}, \mathcal{D}^N).
\end{aligned}$$

By the tower property of conditional expectations, and since $\boldsymbol{x}^\pi$ is $\mathcal{F}^N$ measurable,

$$\begin{aligned}
\mathbb{E}[\Pr(y = +1|\boldsymbol{x}^\pi, \boldsymbol{w})] &= \mathbb{E}[\sigma(\boldsymbol{w}^T\boldsymbol{x}^\pi)] \\
&= \mathbb{E}\mathbb{E}_N[\sigma(\boldsymbol{w}^T\boldsymbol{x}^\pi)] \\
&= \mathbb{E}[\Pr(y = +1|\boldsymbol{x}^\pi, \mathcal{D}^N)].
\end{aligned}$$

Then by the definition of $\boldsymbol{x}^\pi$, we have $\Pr(y = +1|\boldsymbol{x}^\pi, \mathcal{D}^N) = \max_{\boldsymbol{x}} \Pr[y = +1|\boldsymbol{x}, \mathcal{D}^N]$. □

A.2 Proof of Theorem 1

Theorem 1 Let $\mathcal{D}^n$ be the n measurements produced by the KG policy and $\boldsymbol{w}^n = \arg\max_{\boldsymbol{w}} \Psi(\boldsymbol{w}|\boldsymbol{m}^0, \Sigma^0)$ with the prior distribution $\Pr(\boldsymbol{w}^*) = \mathcal{N}(\boldsymbol{w}^*|\boldsymbol{m}^0, \boldsymbol{\Sigma}^0)$. Then with probability $P_d(M)$, the expected error of $\boldsymbol{w}^n$ is bounded as

$$\mathbb{E}_{\boldsymbol{y}\sim\mathcal{B}(\mathcal{D}^n,\boldsymbol{w}^*)}||\boldsymbol{w}^n - \boldsymbol{w}^*||_2 \leq \frac{C_{min} + \lambda_{min}\big(\boldsymbol{\Sigma}^{-1}\big)}{2},$$

where the distribution $\mathcal{B}(\mathcal{D}^n, \boldsymbol{w}^*)$ is the vector onBernoulli distribution with $Pr(y^i = +1) = \sigma(\boldsymbol{w}^{*T}\boldsymbol{x}^i)$ of each dimension, $P_d(M)$ is the probability of a d-dimensional standard normal random variable appears in the ball with radius $M = \frac{1}{8}\frac{\lambda_{min}^2}{\sqrt{\lambda_{max}}}$ and $C_{min} = \lambda_{min}\Big(\frac{1}{n}\sum_{i=1}\sigma(\boldsymbol{w}^{*T}\boldsymbol{x}^i)\big(1 - \sigma(\boldsymbol{w}^{*T}\boldsymbol{x}^i)\big)\boldsymbol{x}^i(\boldsymbol{x}^i)^T\Big)$.

Proof. In this proof, we use $\boldsymbol{\Sigma}$ and $\boldsymbol{m}$ to denote $\boldsymbol{\Sigma}^0$ and $\boldsymbol{m}^0$, and use $\boldsymbol{x}_i$, y_i to denote $\boldsymbol{x}^i$ and y^i for notation simplicity.

Let $f(\boldsymbol{w}) = g(\boldsymbol{w}) + h(\boldsymbol{w})$, where

$$g(\boldsymbol{w}) = \frac{1}{2}(\boldsymbol{w} - \boldsymbol{m})^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{w} - \boldsymbol{m}),$$

and

$$h(\boldsymbol{w}) = -\mathbb{E}_{\boldsymbol{y}\sim\mathcal{B}(\mathcal{D}^n, \boldsymbol{w}^*)}\Big[\frac{1}{n}\sum_{i=1}^{n} \log \sigma(y_i \boldsymbol{w}^T \boldsymbol{x}_i)\Big].$$

Then based on mean value theorem we have

$$h(\boldsymbol{w}) - h(\boldsymbol{w}^*) = (\boldsymbol{w} - \boldsymbol{w}^*)^T \nabla h(\boldsymbol{w}^*) + \frac{1}{2}(\boldsymbol{w} - \boldsymbol{w}^*)^T H\big(\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*)\big)(\boldsymbol{w} - \boldsymbol{w}^*),$$

where H is the Hessian of h.

To analyze the first and second order terms, we use a similar technique adopted in [1]. For the first order term, we have

$$\begin{aligned} \nabla h(\boldsymbol{w}^*) &= \mathbb{E}_{\boldsymbol{y}\sim\mathcal{B}(\mathcal{D}^n, \boldsymbol{w}^*)}\Big[\frac{1}{n}\sum_{i=1} \big(1 - \sigma(y_i \boldsymbol{w}^T \boldsymbol{x}_i)\big) y_i \boldsymbol{x}_i\Big] \\ &= \frac{1}{n}\sum_{i=1} \Big(\big(1 - \sigma(\boldsymbol{w}^T \boldsymbol{x}_i)\big)\sigma(\boldsymbol{w}^T \boldsymbol{x}_i) - \big(1 - \sigma(-\boldsymbol{w}^T \boldsymbol{x}_i)\big)\big(1 - \sigma(\boldsymbol{w}^T \boldsymbol{x}_i)\big)\Big)\boldsymbol{x}_i = 0. \end{aligned}$$

For the second order term, we have

$$\begin{aligned} & H\big(\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*)\big) \\ = \;& \mathbb{E}_{\boldsymbol{y}\sim\mathcal{B}(\mathcal{D}^n, \boldsymbol{w}^*)}\Bigg(\frac{1}{n}\sum_{i=1} \sigma(y_i(\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*))^T \boldsymbol{x}_i)\Big(1 - \sigma(y_i(\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*))^T \boldsymbol{x}_i)\Big)\boldsymbol{x}_i \boldsymbol{x}_i^T\Bigg) \\ = \;& \frac{1}{n}\sum_{i=1} \sigma((\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*))^T \boldsymbol{x}_i)\Big(1 - \sigma((\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*))^T \boldsymbol{x}_i)\Big)\boldsymbol{x}_i \boldsymbol{x}_i^T \\ = \;& \frac{1}{n}\sum_{i=1} J_i(\boldsymbol{w}, \eta)\boldsymbol{x}_i \boldsymbol{x}_i^T, \end{aligned}$$

where $J_i(\boldsymbol{w}, \eta) = \sigma((\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*))^T \boldsymbol{x}_i)\Big(1 - \sigma((\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*))^T \boldsymbol{x}_i)\Big)$.

We expand $J_i(\boldsymbol{w}, \eta)$ to its first order and use mean value theorem again,

$$\begin{aligned} |J_i(\boldsymbol{w}, \eta) - J_i(\boldsymbol{w}^*, \eta)| &= |\eta(\boldsymbol{w} - \boldsymbol{w}^*)^T \boldsymbol{x}_i \sigma'| \\ &\le |\sigma(1-\sigma)(1-2\sigma)||(\boldsymbol{w} - \boldsymbol{w}^*)^T \boldsymbol{x}_i| \\ &\le \|\boldsymbol{w} - \boldsymbol{w}^*\|_2, \end{aligned}$$

where we omit the dependence of σ on $(\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*))^T \boldsymbol{x}_i$ for simplicity and use the fact $\sigma \in (0, 1)$.

Combining the first order and second order analysis and denoting $\|\boldsymbol{w} - \boldsymbol{w}^*\|_2$ as R, we have

$$\begin{aligned} h(\boldsymbol{w}) - h(\boldsymbol{w}^*) &\ge \frac{1}{2}\|\boldsymbol{w} - \boldsymbol{w}^*\|_2^2 \lambda_{min}\Big(H\big(\boldsymbol{w}^* + \eta(\boldsymbol{w} - \boldsymbol{w}^*)\big)\Big) \\ &\ge \frac{1}{2}R^2 \lambda_{min}\Big(\frac{1}{n}\sum_{i=1} \sigma(\boldsymbol{x}_i^T \boldsymbol{w}^*)\big(1 - \sigma(\boldsymbol{x}_i^T \boldsymbol{w}^*)\big)\boldsymbol{x}_i \boldsymbol{x}_i^T - R\Big) \\ &\ge \frac{1}{2}C_{min} R^2 - \frac{1}{2}R^3, \end{aligned}$$

where we use the fact that $\|\boldsymbol{x}_i\|_2 \le 1$ and use C_{min} to denote $\lambda_{min}\Big(\frac{1}{n}\sum_{i=1} \sigma(\boldsymbol{x}_i^T \boldsymbol{w}^*)\big(1 - \sigma(\boldsymbol{x}_i^T \boldsymbol{w}^*)\big)\boldsymbol{x}_i \boldsymbol{x}_i^T\Big)$.

On the other hand, for $g(\boldsymbol{w})$ we have

$$\begin{aligned} g(\boldsymbol{w}) - g(\boldsymbol{w}^*) &= (\boldsymbol{w} - \boldsymbol{w}^*)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{w}^* - \boldsymbol{m}) + \frac{1}{2} (\boldsymbol{w} - \boldsymbol{w}^*)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{w} - \boldsymbol{w}^*) \\ &\geq -DR\sqrt{\lambda_{max}\big(\boldsymbol{\Sigma}^{-1}\big)} + \frac{1}{2}\lambda_{min}\big(\boldsymbol{\Sigma}^{-1}\big)R^2, \end{aligned}$$

where$D = \|\boldsymbol{\Sigma}^{-\frac{1}{2}}(\boldsymbol{w}^* - \boldsymbol{m})\|_2$.

Now define function $F(\Delta)$ as $F(\Delta) = f(\boldsymbol{w}^* + \Delta) - f(\boldsymbol{w}^*)$, then we have

$$F(\Delta) = g(\boldsymbol{w}) - g(\boldsymbol{w}^*) + h(\boldsymbol{w}^n) - h(\boldsymbol{w}^*) \geq -\sqrt{\lambda_{max}\big(\boldsymbol{\Sigma}^{-1}\big)}DR + \frac{1}{2}\lambda_{min}\big(\boldsymbol{\Sigma}^{-1}\big)R^2 + \frac{1}{2}C_{min}R^2 - \frac{1}{2}R^3.$$

From now on we will use the simplified symbols λ_{max} and λ_{min} instead of $\lambda_{min}\big(\boldsymbol{\Sigma}^{-1}\big)$ and $\lambda_{max}\big(\boldsymbol{\Sigma}^{-1}\big)$. It is easy to check that in the case when

$$D \leq \frac{1}{8}\frac{\lambda_{min}^2}{\sqrt{\lambda_{max}}}, \tag{A.1}$$

$F(\Delta) \geq 0$ at for all Δ with $\|\Delta\|_2 = \frac{1}{2}(C_{min} + \lambda_{min})$.

Notice that $F(0) = 0$ and recall that $\boldsymbol{w}^n$ minimizes $f(\boldsymbol{w})$ so we have

$$F(\boldsymbol{w}^n - \boldsymbol{w}^*) = f(\boldsymbol{w}) - f(\boldsymbol{w}^*) \leq 0.$$

Then based on the convexity of F we know that $\|\boldsymbol{w}^n - \boldsymbol{w}^*\|_2 \leq \frac{1}{2}(C_{min} + \lambda_{min})$, otherwise the values of F at 0, $\boldsymbol{w}^n - \boldsymbol{w}^*$ and the intersection between the all $\|\Delta\|_2 = \frac{1}{2}(C_{min} + \lambda_{min})$ and line from 0 to $\boldsymbol{w}^n - \boldsymbol{w}^*$ form a concave pattern, which is contradictory.

Now we start to calculate the probability for Eq. (A.1) to hold. Recall that $\boldsymbol{w}^*$ has a prior distribution $\boldsymbol{w}^* \sim \mathcal{N}(\boldsymbol{m}, \boldsymbol{\Sigma})$. Then by denoting the right hand side of Eq. (A.1) as M, we have

$$\begin{aligned} &\Pr\Big(D \leq \frac{1}{8}\frac{\lambda_{min}^2}{\sqrt{\lambda_{max}}}\Big) \\ = &\int_{\|\boldsymbol{\Sigma}^{-\frac{1}{2}}(\boldsymbol{a}-\boldsymbol{m})\| \leq M} (2\pi)^{-\frac{d}{2}} |\boldsymbol{\Sigma}|^{-\frac{1}{2}} \exp\big(-\frac{1}{2}(\boldsymbol{a} - \boldsymbol{m})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{a} - \boldsymbol{m})\big) \mathrm{d}\boldsymbol{a} \\ = &\int_{\|\boldsymbol{b}\|_2 \leq M} (2\pi)^{-\frac{d}{2}} \exp\big(-\frac{1}{2}\boldsymbol{b}^T \boldsymbol{b}\big) \mathrm{d}\boldsymbol{b}, \end{aligned}$$

which is the probability of a d-dimension standard normal random variable appears in the ball with radius M, $P_d(M)$. This completes the proof. □

References

[1] Pradeep Ravikumar, Martin J Wainwright, John D Lafferty, et al. High-dimensional ising model selection using 1-regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.