Research note on the martingale concentration inequality

Wei Xiong *

Abstract

We are interested in the martingale concentration inequality used in sequential estimation problem.

Contents

1 Introduction

This is a research note when reading chapter 13 of Zhang (2023).

In sequential estimation problem, we will observe a sequence of random variables $Z_t \in \mathcal{Z}_t$, where Z_t may depend on the history $\mathcal{S}_{t-1} = [Z_1, \cdots, Z_{t-1}] \in \mathcal{Z}^{t-1}$. We denote the sigma algebra generated by $\mathcal{S}_t$ as the filtration $\mathcal{F}_t$. We say a sequence $\{\xi_t\}$ is adapted to the filtration $\{\mathcal{F}_t\}$, if each ξ_t is a function of $\mathcal{S}_t$. That is, each ξ_t does not depend on the future $(Z_s, s > t)$. This is also referred to as that ξ_t is measurable in $\mathcal{F}_t$. The sequence

$$\xi'_t(\mathcal{S}_t) := \xi_t(\mathcal{S}_t) - \mathbb{E}[\xi_t(\mathcal{S}_t)|\mathcal{F}_{t-1}] = \xi_t(\mathcal{S}_t) - \mathbb{E}_{Z_t|\mathcal{S}_{t-1}}\xi_t(\mathcal{S}_t),$$

is referred to as a martingale difference sequence, where we have

$$\mathbb{E}[\xi'_t|\mathcal{F}_{t-1}] = 0.$$

The sum of such a martingale difference sequence

$$\sum_{s=1}^{t} \xi'_s = \sum_{s=1}^{t} \xi'_s(\mathcal{S}_s)$$

is referred to as a martingale. In what follows, we further assume that $\mathcal{Z} = \mathcal{Z}^{(x)} \times \mathcal{Z}^{(y)}$ and $Z_t = (Z_t^{(x)}, Z_t^{(y)})$. For instance, the $Z_t^{(x)}$ may be regarded as the context of the contextual bandit in iteration t, while $Z_t^{(y)}$ is the random reward.

*Email: wx13@illinois.edu

2 Martingale Exponential Inequalities

For notation simplicity, we use

$$\mathbb{E}_{Z_t^{(y)}}[\cdot] := \mathbb{E}_{Z_t^{(y)}|Z_t^{(x)},\mathcal{S}_{t-1}}[\cdot].$$

Lemma 1 (Martingale Exponential Inequalities)**.** *Consider a sequence of real-valued random functions $\xi_1(\mathcal{S}_1), \cdots, \xi_T(\mathcal{S}_T)$. Let $\tau \le T$ be a stopping time so that $\mathbb{I}(t \le \tau)$ is measurable in $\mathcal{S}_t$. We have*

$$\mathbb{E}_{\mathcal{S}_T} \exp\big(\sum_{s=1}^{\tau} \xi_i - \sum_{s=1}^{\tau} \log \mathbb{E}_{Z_i^{(y)}} \exp(\xi_i)\big) = 1.$$

where $Z_t = (Z_t^{(x)}, Z_t^{(y)})$ and $\mathcal{Z}_t = (Z_1, \cdots, Z_t)$.

Proof. We prove by induction on T'. When $T' = 0$, the inequality is trivial. Now suppose that it holds for $T' - 1$ for some $T' \ge 1$. We let $\widetilde{\xi}_i = \xi_i \mathbb{I}(i \le \tau)$ which is measurable in $\mathcal{S}_i$. We have

$$\sum_{s=1}^{\tau} \xi_i - \sum_{s=1}^{\tau} \log \mathbb{E}_{Z_i^{(y)}} \exp(\xi_i) = \sum_{s=1}^{\tau} \widetilde{\xi}_i - \sum_{s=1}^{\tau} \log \mathbb{E}_{Z_i^{(y)}} \exp(\widetilde{\xi}_i).$$

It follows that

$$\begin{aligned}
&\mathbb{E}_{Z_1,\cdots,Z_{T'}} \exp\big(\sum_{s=1}^{\tau} \xi_i - \sum_{s=1}^{\tau} \log \mathbb{E}_{Z_i^{(y)}} \exp(\xi_i)\big) \\
&= \mathbb{E}_{Z_1,\cdots,Z_{T'}} \exp\big(\sum_{s=1}^{T'} \widetilde{\xi}_i - \sum_{s=1}^{T'} \log \mathbb{E}_{Z_i^{(y)}} \exp(\widetilde{\xi}_i)\big) \\
&= \mathbb{E}_{Z_1,\cdots,Z_{T'}^{(x)}} \Big[\exp\big(\sum_{s=1}^{T'-1} \widetilde{\xi}_i - \sum_{s=1}^{T'-1} \log \mathbb{E}_{Z_i^{(y)}} \exp(\widetilde{\xi}_i)\big) \mathbb{E}_{Z_{T'}^{(y)}} \exp\big(\widetilde{\xi}_{T'} - \log \mathbb{E}_{Z_{T'}^{(y)}} \exp(\widetilde{\xi}_{T'})\big)\Big] \\
&= \mathbb{E}_{Z_1,\cdots,Z_{T'-1}} \Big[\exp\big(\sum_{s=1}^{T'-1} \widetilde{\xi}_i - \sum_{s=1}^{T'-1} \log \mathbb{E}_{Z_i^{(y)}} \exp(\widetilde{\xi}_i)\big)\Big] \\
&= \mathbb{E}_{Z_1,\cdots,Z_{\min(\tau,T'-1)}} \Big[\exp\big(\sum_{s=1}^{\min(\tau,T'-1)} \xi_i - \sum_{s=1}^{\min(\tau,T'-1)} \log \mathbb{E}_{Z_i^{(y)}} \exp(\xi_i)\big)\Big] = 1.
\end{aligned}$$

Here, the last inequality follows from $\min(\tau, T'-1)$ is a stopping time $\le T' - 1$ so we can use the induction hypothesis. □

As a corollary, we have the following counterpart of Chernoff inequality.

Lemma 2 (Martingale Concentration inequality)**.** *Consider a sequence of real-valued random functions $\xi_1(\mathcal{S}_1), \cdots, \xi_T(\mathcal{S}_T)$ adapted to the filtration $\mathcal{F}_t$. We have for any $\delta \in (0,1)$ and $\lambda > 0$:*

$$\mathbb{P}\Big[\exists n > 0 : -\sum_{i=1}^{n} \xi_i \ge \frac{\log(1/\delta)}{\lambda} + \frac{1}{\lambda} \sum_{i=1}^{n} \log \mathbb{E}_{Z_i^{(y)}} \exp(-\lambda \xi_i)\Big] \le \delta.$$

Proof. The proof can be found in Chapter 13 of Zhang (2023). □

Remark 1. *One interesting observation is that Lemma 2 already ensures that the inequality holds for arbitrary $n > 0$ with high probability, while in the i.i.d. setting, this may come from an additional uniform convergence argument. However, we note that the hyper-parameter $\lambda > 0$ is fixed across $n > 0$, which leads to subtle difference in application of the above lemma. We present several examples in the next section.*

3 Examples with application

3.1 Azuma-Hoeffding's inequality

Lemma 3 (Martingale Sub-Gaussian inequality)**.** *Consider a sequence of random functions* $\xi_1(\mathcal{S}_1), \cdots, \xi_t(\mathcal{S}_t), \cdots$. *Assume each* ξ_i *is sub-Gaussian with respect to* $Z_i^{(y)}$*:*

$$\log \mathbb{E}_{Z_i^{(y)}} \leq \lambda \mathbb{E}_{Z_i^{(y)}} \xi_i + \frac{\lambda^2 \sigma_i^2}{2}$$

for some σ_i *that may depend on* $\mathcal{S}_{i-1}$ *and* $Z_i^{(x)}$*. Then for all* $\sigma > 0$*, with probability at least* $1-\delta$*,*

$$\forall n > 0 : \sum_{i=1}^{n} \mathbb{E}_{Z_i^{(y)}} \xi_i < \sum_{i=1}^{n} \xi_i + \big(\sigma + \frac{\sum_{i=1}^{n} \sigma_i^2}{\sigma}\big) \sqrt{\frac{\log(1/\delta)}{2}}.$$

Proof. We set $\lambda = \sqrt{2\log(1/\delta)}/\sigma$ and apply Lemma 2. □

The main problem is that the σ (essentially, the λ in Lemma 2) has to be fixed for all n. Therefore, one cannot set $\sigma = \sqrt{\sum_{i=1}^{n} \sigma_i^2}$ to achieve the best bound for all n. Tuning the parameter requires us to pay for another $\log T$.

Lemma 4 (Azuma-Hoeffding's inequality)**.** *Consider a sequence of random functions* $\xi_1(\mathcal{S}_1), \cdots, \xi_n(\mathcal{S}_n)$ *with a fixed* $n > 0$*. If for each* $i : \sup \xi_i - \inf \xi_i \leq M_i$ *for some constant* M_i*, then with probability at least* $1-\delta$*,*

$$\sum_{i=1}^{n} \mathbb{E}_{Z_i^{(y)}} \xi_i < \sum_{i=1}^{n} \xi_i + \sqrt{\frac{\sum_{i=1}^{n} M_i^2 \log(1/\delta)}{2}}.$$

3.2 Freedman's inequality

Lemma 5 (Freedman's/Bernstein's inequality for martingale)**.** *Let* ξ_t' *be the martingale difference defined in Section 1. If* $|\xi_t'| \leq R$ *almost surely, then for any* $\eta \in (0, 1/R)$*, with probability at least* $1-\delta$*, for all* $T' \leq T$*,*

$$\sum_{t=1}^{T'} \xi_t' \leq \eta \sum_{t=1}^{T'} \mathbb{E}_{t-1}[(\xi_t')^2] + \frac{\log(1/\delta)}{\eta}.$$

Proof. We define $\zeta_s = \eta\xi_s' - \eta^2 \mathbb{E}_{Z_s^{(y)}}(\xi_s')^2$, which is measurable in $\mathcal{S}_s$. Then, we estimate the conditional log-moment-generating function as follows

$$\begin{aligned}
&\mathbb{E}_{Z_s^{(y)}} \exp\Big(\zeta_s\Big) \\
&= \mathbb{E}_{Z_s^{(y)}} \exp\Big(\eta\xi_s' - \eta^2 \mathbb{E}_{Z_s^{(y)}}(\xi_s')^2\Big) \\
&= \exp\big(-\eta^2 \mathbb{E}_{Z_s^{(y)}}(\xi_s')^2\big) \cdot \mathbb{E}_{Z_s^{(y)}} \exp\big(\eta\xi_s'\big) \\
&\leq \exp\big(-\eta^2 \mathbb{E}_{Z_s^{(y)}}(\xi_s')^2\big) \cdot \mathbb{E}_{Z_s^{(y)}} \big[1 + \eta\xi_s' + (e-2)(\eta\xi_s')^2\big] \\
&= \exp\big(-\eta^2 \mathbb{E}_{Z_s^{(y)}}(\xi_s')^2\big) \cdot \big[1 + (e-2)\mathbb{E}_{Z_s^{(y)}}(\eta\xi_s')^2\big] \\
&\leq \exp\big(-\eta^2 \mathbb{E}_{Z_s^{(y)}}(\xi_s')^2\big) \cdot \mathbb{E}_{Z_s^{(y)}} \exp\Big((e-2)(\eta\xi_s')^2\Big) \\
&\leq e^0 = 1,
\end{aligned}$$

where we use $e^a \leq 1 + a + (e-2)a^2$ for $a \leq 1$ in the first inequality (the constraint on η), and $\mathbb{E}_{Z_s^{(y)}} \xi_s' = 0$ because ξ_s' is a martingale difference sequence. The second inequality uses $1 + a \leq e^a$ for all $a \in \mathbb{R}$. Then, we can invoke Lemma 2 with $\{\zeta_t\}$ and $\lambda = 1$ to obtain that

$$\forall n > 0, \sum_{i=1}^{n} \zeta_i < \frac{\log(1/\delta)}{\eta} + \sum_{i=1}^{n} \log \mathbb{E}_{Z_i^{(y)}} \exp(-\zeta_i) \leq \log(1/\delta).$$

Plugging $\zeta_i = \eta\xi'_i - \eta^2\mathbb{E}_{Z_s^{(y)}}(\xi'_s)^2$ finishes the proof. □

Example 1. *We consider the contextual bandit problem with general function approximation. For $f \in \mathcal{F}$, we define*

$$U^t(f) = \left(f(x^t, a^t) - r^t\right)^2 - \left(f^*(x^t, a^t) - r^t\right)^2.$$

We have

$$\mathbb{E}_{t-1}U^t(f) = \mathbb{E}_{t-1}(f(x^t, a^t) - f^*(x^t, a^t))^2,$$

where on the RHS, there is still randomness for a^t and $\mathbb{E}_{t-1}[\cdot] = \mathbb{E}[\cdot|\mathcal{H}^{t-1}, x^t]$. Then $Z^t(f) = \mathbb{E}_{t-1}U^t(f) - U^t(f)$ is a martingale difference and $\sum_{t=1}^{\tau} Z^t(f)$ is a martingale sequence. Since the increment $|Z^t(f)| \leq 1$, we can apply the Freedman's inequality to get that w.p. at least $1-\delta$, for all $\tau \leq T$,

$$\sum_{t=1}^{\tau} Z^t(f) \leq \frac{1}{8}\mathbb{E}_{t-1}[Z^t(f)^2] + 8\log(1/\delta).$$

We can control the second-order bound by

$$\mathbb{E}_{t-1}[Z^t(f)^2] \leq 4\mathbb{E}_{t-1}[(f(x^t, a^t) - f^*(x^t, a^t))^2] = 4\mathbb{E}_{t-1}U^t(f),$$

where it follows that

$$\frac{1}{2}\sum_{t=1}^{\tau}\mathbb{E}_{t-1}U^t(f) \leq \sum_{t=1}^{\tau} U^t(f) + 8\log(1/\delta).$$

Since $0 \leq \frac{1}{2}\sum_{t=1}^{\tau}\mathbb{E}_{t-1}U^t(f) \leq \sum_{t=1}^{\tau} U^t(f) + 8\log(1/\delta)$, we have

$$\sum_{t=1}^{\tau}(f^*(x^t, a^t) - r^t)^2 \leq \sum_{t=1}^{\tau}(f(x^t, a^t) - r^t)^2 + 8\log(1/\delta).$$

Taking a union bound over f, we have

$$\sum_{t=1}^{\tau}(f^*(x^t, a^t) - r^t)^2 \leq \sum_{t=1}^{\tau}(f(x^t, a^t) - r^t)^2 + 8\log(|\mathcal{F}|/\delta),$$

holds with probability at least $1-\delta$ for any $f \in [F], \tau \in [T]$.

A distinct feature is that since we do not tune $\eta > 0$ for different times step, we can directly invoke the inequality that holds for all $n > 0$. Therefore, we do not pay for an additional $\log T$ here.

If one wishes to work with the random variables directly, we have the following result.

Lemma 6. *Let $(\xi_t)_{t \leq T}$ be a sequence of random variables adapted to a filtration $(\mathcal{F}_t)_{t \leq T}$. If $0 \leq \xi_t \leq R$ almost surely, then with probability at least $1-\delta$,*

$$\sum_{t=1}^{T}\xi_t \leq \frac{3}{2}\sum_{t=1}^{T}\mathbb{E}_{Z_t^{(y)}}[\xi_t] + 4R\log(2\delta^{-1}),$$

and

$$\sum_{t=1}^{T}\mathbb{E}_{Z_t^{(y)}}[\xi_t] \leq 2\sum_{t=1}^{T}\xi_t + 8R\log(2\delta^{-1}).$$

3.3 MLE Analysis

As an additional example, we consider the model-based case, where we need to estimate the Hellinger distance between the model and the true model via the likelihood estimator.

Suppose that for each iteration t, we will choose a model $M^t \in \mathcal{M}$ (e.g. by UCB or posterior sampling) and collect a new trajectory by executing $x_h \sim \pi_{M^t}, a_h \sim \widetilde{\pi}_t$, where $\widetilde{\pi}_t$ is either π_{M^t} (Q-type) or a uniform exploration over $\mathcal{A}$ (V-type). Then, suppose that we set $\beta_h = 2\log(H|\mathcal{M}_h|/\delta)$, with probability at least $1-\delta$, for all $t \in [T]$, we have

$$\sum_{h=1}^{H}\sum_{s=1}^{t-1} \mathbb{E}_{x_h \sim \pi_{M^s}, a_h \sim \widetilde{\pi}_s} D_{\mathrm{H}}^2\big(\mathbb{P}_{h,M}(\cdot \mid x_h, a_h), \mathbb{P}_{h,M^*}(\cdot \mid x_h, a_h)\big) \leq \beta,$$

where $\beta := \sum_{h=1}^{H} \beta_h = 2\log(H|\mathcal{M}|/\delta)$ and $\widetilde{\pi}_s$ is either π_{M^s} or $\mathrm{Unif}(\mathcal{A})$.

We have the following estimation of the moment-generating function. For any $M \in \mathcal{M}$, we have

$$\begin{aligned}
&\mathbb{E}\Big[\exp\Big(\frac{1}{2}\sum_{s=1}^{t}\log\frac{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^s, a_h^s)}{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^s, a_h^s)}\Big)\Big] \\
&= \mathbb{E}\Big[\exp\Big(\frac{1}{2}\sum_{s=1}^{t-1}\log\frac{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^s, a_h^s)}{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^s, a_h^s)}\Big)\Big]\mathbb{E}_t\sqrt{\frac{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^t, a_h^t)}{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^t, a_h^t)}} \\
&= \mathbb{E}\Big[\exp\Big(\frac{1}{2}\sum_{s=1}^{t-1}\log\frac{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^s, a_h^s)}{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^s, a_h^s)}\Big)\Big]\mathbb{E}_{x_h \sim \pi_{M^t}, a_h \sim \widetilde{\pi}_t}\int_{x \in \mathcal{S}}\sqrt{\mathbb{P}_{h,M}(x \mid x_h, a_h)\cdot\mathbb{P}_{h,M^*}(x \mid x_h, a_h)} \\
&= \mathbb{E}\Big[\exp\Big(\frac{1}{2}\sum_{s=1}^{t-1}\log\frac{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^s, a_h^s)}{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^s, a_h^s)}\Big)\Big]\Big(1 - \mathbb{E}_{x_h \sim \pi_{M^t}, a_h \sim \widetilde{\pi}_t} D_{\mathrm{H}}^2\big(\mathbb{P}_{h,M}(\cdot \mid x_h, a_h), \mathbb{P}_{h,M^*}(\cdot \mid x_h, a_h)\big)\Big) \\
&= \cdots \\
&= \prod_{s=1}^{t}\Big(1 - \mathbb{E}_{x_h \sim \pi_{M^s}, a_h \sim \widetilde{\pi}_s} D_{\mathrm{H}}^2\big(\mathbb{P}_{h,M}(\cdot \mid x_h, a_h), \mathbb{P}_{h,M^*}(\cdot \mid x_h, a_h)\big)\Big)
\end{aligned}$$

We now invoke Lemma 1 to obtain that for any fixed M_h, we have

$$\begin{aligned}
1 - \frac{\delta}{H|\mathcal{M}_h|} &\leq \mathbb{P}\Big[\forall t > 0 : \frac{1}{2}\sum_{s=1}^{t}\log\frac{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^s, a_h^s)}{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^s, a_h^s)} \leq \log(H|\mathcal{M}_h|/\delta) \\
&\qquad + \sum_{s=1}^{t}\log\Big(1 - \mathbb{E}_{x_h \sim \pi_{M^s}, a_h \sim \widetilde{\pi}_s} D_{\mathrm{H}}^2\big(\mathbb{P}_{h,M}(\cdot \mid x_h, a_h), \mathbb{P}_{h,M^*}(\cdot \mid x_h, a_h)\big)\Big]\\
&\leq \mathbb{P}\Big[\forall t > 0 : \frac{1}{2}\sum_{s=1}^{t}\log\frac{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^s, a_h^s)}{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^s, a_h^s)} \leq \log(H|\mathcal{M}_h|/\delta) \\
&\qquad - \sum_{s=1}^{t}\mathbb{E}_{x_h \sim \pi_{M^s}, a_h \sim \widetilde{\pi}_s} D_{\mathrm{H}}^2\big(\mathbb{P}_{h,M}(\cdot \mid x_h, a_h), \mathbb{P}_{h,M^*}(\cdot \mid x_h, a_h)\big)\Big]
\end{aligned}$$

where we use $\log(1-x) \leq -x$ for $x \leq 1$. With a union bound over $\mathcal{M}_h$ and then $[H]$, we conclude that with probability at least $1-\delta$, we have for all $t \in [T]$, and for all $h \in [H]$,

$$\sum_{s=1}^{t}\mathbb{E}_{x_h \sim \pi_{M^s}, a_h \sim \widetilde{\pi}_s} D_{\mathrm{H}}^2\big(\mathbb{P}_{h,M}(\cdot \mid x_h, a_h), \mathbb{P}_{h,M^*}(\cdot \mid x_h, a_h)\big) \leq \sum_{s=1}^{t}\frac{1}{2}\log\frac{\mathbb{P}_{h,M^*}(x_{h+1}^s \mid x_h^s, a_h^s)}{\mathbb{P}_{h,M}(x_{h+1}^s \mid x_h^s, a_h^s)} + \log(H|\mathcal{M}_h|/\delta).$$

We also save a $\log T$ because we do not need to tune the parameter λ in Lemma 2.

References

Tong Zhang. Mathematical analysis of machine learning algorithms. 2023. doi: 10.1017/9781009093057.