CS 7545: Machine Learning Theory	**Fall 2018**

Lecture 5: Martingales + Online Learning

Lecturer: Jacob Abernethy — *Scribes: Zaiwei Chen, Junghyun Kim*

5.1 Martingales

In this section, we will introduce the concept of martingales and prove Azuma's Inequality.

Definition 5.1 (Martingale) *A sequence of random variables $Z_0, Z_1, ..., Z_n, ...$ is called a **martingale** sequence if for all $n \in \mathbb{N}$*

(i) $\mathbb{E}[|Z_n|] < \infty$.

(ii) $\mathbb{E}[Z_{n+1}|Z_0, Z_1, ..., Z_n] = Z_n$.

The followings are three examples of martingale.

Example 1 (Linear martingale) Let $\{X_i\}_{i\geq 0}$ be a sequence of *i.i.d* random variables with $\mathbb{E}[X_i] = 0$, $\forall i \geq 0$, then $Z_n = \sum_{i=0}^n X_i$ is a martingale.

(i) In measure theory, $\mathbb{E}[X_i]$ is well defined only if $\mathbb{E}[|X_i|]$ is finite. So we have $\mathbb{E}[|Z_n|] \leq \mathbb{E}[\sum_{i=0}^n |X_i|] < \infty$, $\forall n \geq 0$.

(ii) $\mathbb{E}[Z_{n+1}|Z_0, Z_1, ..., Z_n] = \mathbb{E}[Z_n + X_{n+1}|Z_0, Z_1, ..., Z_n] = Z_n + \mathbb{E}[X_{n+1}] = Z_n$.

Example 2 (Quadratic martingale) Let $\{X_i\}_{i\geq 0}$ be a sequence of *i.i.d* random variables with $\mathbb{E}[X_i] = 0$ and $\sigma^2 = \mathrm{var}(X_i) < \infty$, in this case $Z_n = S_n^2 - n\sigma^2$ is a martingale ($S_n = \sum_{i=0}^n X_i$).

(i) $\mathbb{E}[|S_n^2 - n\sigma^2|] \leq \mathbb{E}[S_n^2] + n\sigma^2 = \mathrm{var}(S_n) + \mathbb{E}^2[S_n] + n\sigma^2 = n\mathrm{var}(X_1) + (n\mathbb{E}[X_i])^2 + n\sigma^2 = 2n\sigma^2 < \infty$.

(ii) Since $\mathbb{E}[S_n^2|S_1, S_2, ..., S_n] = S_n^2$ and $\mathbb{E}[S_n X_{n+1}|S_1, S_2, ..., S_n] = S_n\mathbb{E}[X_{n+1}] = 0$, we have

$$\begin{aligned}\mathbb{E}[S_{n+1}^2 - (n+1)\sigma^2|S_1, S_2, ..., S_n] &= \mathbb{E}[S_n^2 + 2S_nX_{n+1} + X_{n+1}^2|S_1, S_2, ..., S_n] - (n+1)\sigma^2 \\ &= S_n^2 + \mathbb{E}[X_{n+1}^2] - (n+1)\sigma^2 \\ &= S_n^2 - n\sigma^2\end{aligned}$$

Example 3 (Exponential martingale) Let $\{X_i\}_{i\geq 0}$ be a sequence of *i.i.d* **nonnegative** random variables with $\mathbb{E}[X_i] = 1$. Then $Z_n = \Pi_{i=0}^n X_i$ defines a martingale.

(i) $\mathbb{E}[|Z_n|] = \Pi_{i=0}^n \mathbb{E}[X_i] = 1 < +\infty$.

(ii) $\mathbb{E}[Z_{n+1}|Z_0, Z_1, ..., Z_n] = \Pi_{i=0}^n X_n \mathbb{E}[X_{n+1}] = Z_n$.

The next theorem gives a concentration result for the values of martingales that have bounded differences.

Theorem 5.2 (Azuma's Inequality) *Let $Z_0 = 0, Z_1, \cdots, Z_n$ be a martingale sequence with $-c_i \leq Z_i - Z_{i-1} \leq c_i$, $\forall i \geq 1$. Then*

$$\mathbb{P}(Z_n \geq t) \leq \exp\left(-\frac{t^2}{2\sum_{i=1}^n c_i^2}\right)$$

Proof: Let $\lambda > 0$, the following derivation based on Markov's inequality and the identity $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X|Y]]$.

$$\begin{aligned}\mathbb{P}(Z_n \geq t) &= \mathbb{P}(\exp{(\lambda Z_n)} \geq \exp{(\lambda t)}) \\ &\leq \exp{(-\lambda t)}\mathbb{E}[\exp{(\lambda Z_n)}] \\ &= \exp{(-\lambda t)}\mathbb{E}[\exp{(\lambda Z_{n-1} + \lambda(Z_n - Z_{n-1}))}] \\ &= \exp{(-\lambda t)}\mathbb{E}[\mathbb{E}[\exp{(\lambda Z_{n-1} + \lambda(Z_n - Z_{n-1}))}|Z_1, Z_2, ..., Z_{n-1}]] \\ &= \exp{(-\lambda t)}\mathbb{E}[\exp{(\lambda Z_{n-1})}\mathbb{E}[\exp{(\lambda(Z_n - Z_{n-1}))}|Z_1, Z_2, ..., Z_{n-1}]]\end{aligned}$$

Notice that $\mathbb{E}[\lambda(Z_n - Z_{n-1})|Z_1, Z_2, ..., Z_{n-1}] = 0$ by definition of martingale, and $-c_n \leq Z_n - Z_{n-1} \leq c_n$. We are ready to use **Conditional Hoeffding's Lemma**[1]

$$\begin{aligned}\mathbb{E}[\exp{(\lambda(Z_n - Z_{n-1}))}|Z_1, Z_2, ..., Z_{n-1}] &\leq \exp{(\frac{\lambda^2(2c_n)^2}{8})} \\ &= \exp{(\frac{\lambda^2 c_n^2}{2})}\end{aligned}$$

Now we can compute the upper bound recursively as

$$\begin{aligned}\mathbb{P}(Z_n \geq t) &\leq \exp{(-\lambda t)}\exp{(\frac{\lambda^2 c_n^2}{2})}\mathbb{E}[\exp{(\lambda Z_{n-1})}] \leq \cdots \\ &\leq \exp{(-\lambda t)}\exp{(\frac{\lambda^2\sum_{i=1}^n c_i^2}{2})}\mathbb{E}[\exp{(\lambda Z_0)}] \\ &= \exp{(\frac{\lambda^2\sum_{i=1}^n c_i^2}{2} - \lambda t)}\end{aligned}$$

Since the above inequality holds for all $\lambda > 0$, we have

$$\begin{aligned}\mathbb{P}(Z_n \geq t) &\leq \inf_{\lambda > 0}\exp{(\frac{\lambda^2\sum_{i=1}^n c_i^2}{2} - \lambda t)} \\ &= \exp{(-\frac{t^2}{2\sum_{i=1}^n c_i^2})}\end{aligned}$$

■

5.2 Online Learning

In this section we will introduce prediction with expert advice.

Example 1 (Weather report) On each round $t \in \mathbb{N}^+$

(i) N experts predict weather, $x_{it} \in \{0, 1\}$, $i = 1, 2, \cdots, N$.

(ii) Algorithm listens, and predicts $\hat{y}_t \in \{0, 1\}$.

(iii) Nature reveals $y_t \in \{0, 1\}$.

Let M_t (L_t^i) be the number of mistakes made by the algorithm (the *ith* expert) after t rounds (i.e. $M_t = \sum_{s=1}^t \mathbb{1}[\hat{y_s} \neq y_s]$, $L_t^i = \sum_{s=1}^t \mathbb{1}[\hat{x_{is}} \neq y_s]$). Assume there exists a perfect expert, namely $\exists i^* \in \{1, 2, \cdots, N\}$ such that $L_t^{i^*} = 0$, $\forall t \in \mathbb{N}^+$. The question is: What is the best algorithm to minimize M_t? What is the upper bound on M_t for a particular algorithm? The second question can be rephrased as: **For a particular algorithm, at most how many mistakes the algorithm needs to make until it finally finds the perfect expert?**

[1]The proof for Conditional Hoeffding's Lemma is similar to the proof for Hoeffding's Lemma. Use $\mathbb{E}[\cdot|Z]$ instead of $\mathbb{E}[\cdot]$ and use $a = f(Z)$, $b = f(Z) + c$ for the lower and upper bound.

Theorem 5.3 (Trivial Bound) $M_t \leq N - 1$

Proof: Let the algorithm be following the expert whose index is the smallest among those who have not yet made mistakes. The worst case should be $x_{it} \neq y_t$ when $i \leq t$ and $x_{it} = y_t$ when $i > t$, for all $1 \leq t \leq N-1$. In this case the algorithm will follow the wrong prediction for $N-1$ times until it finally finds the perfect expert and stops making mistakes since then. ■

Theorem 5.4 (Existence of an algorithm with Log Bound) *There exists an algorithm such that* $M_t \leq \log_2 N$

Proof: The idea is to follow the majority vote of "perfect to now" experts (Halving Algorithm). Let C_{t+1} be the index set of experts who have not yet made mistakes in the first t rounds. We have

$$C_1 = \{1, 2, 3, \cdots, N\}$$
$$C_{t+1} = C_t \backslash \{i : x_{it} \neq y_t\}$$
$$\hat{y_t} = [\frac{1}{|C_t|} \sum_{i \in C_t} x_{it}]$$

Here $[\ \cdot\ ]$ means rounding to the nearest integer. If $\frac{1}{|C_t|} \sum_{i \in C_t} x_{it} = \frac{1}{2}$, we break ties arbitrarily. Observe that once the algorithm made a mistake, the majority (no less than half) of "perfect to now" experts must made a wrong prediction. Then the number of the remaining "perfect to now" experts in the next round would be reduced by at least half. Mathematically, it means

$$\hat{y_t} \neq y_t \quad \text{(or say } M_t = M_{t-1} + 1) \quad \Longrightarrow \quad |C_{t+1}| \leq \frac{|C_t|}{2}$$
$$\Longrightarrow \quad |C_{t+1}| \leq |C_1| (\frac{1}{2})^{M_t}$$

Suppose the algorithm finds the perfect expert after round T, that means $|C_{T+1}| \leq 1$. Solve $|C_1|(\frac{1}{2})^{M_T} = 1$ and we get $M_T = \log_2 N$. After round T the algorithm stops making mistakes and we have $M_t \leq \log_2 N$, $\forall t \in \mathbb{N}^+$. ■

Example 2 (Team match). For each round $t = 1, 2, 3, \cdots$

(i) Team i, j arrive.

(ii) The algorithm predicts $i_t < j_t$ (team j wins the match) or $i_t > j_t$ (team i wins the match).

(iii) Nature reveals winner i_t or j_t.

Assume there exists a correct ranking among all n teams, namely there exists a permutation $\Pi \in S_n$ (all permutations on $[n]$) such that $i < j \Leftrightarrow \Pi(i) < \Pi(j)$. The questions is: **What is the best algorithm? What is its corresponding M_T?**

Theorem 5.5 (Existence of an algorithm with Log factorial Bound) *There exists an algorithm such that* $M_t \leq \log_2(n!)$.

Proof: Consider each permutation as an expert prediction. Then there are total $n!$ experts. The correct ranking corresponds to the perfect expert. Now the result follows from using the majority vote algorithm mentioned in the previous example. ■

Remarks

(i) To check whether an expert makes a right prediction, we need to check whether the outcome of the match among n teams is consistent with the permutation provided by the expert. This decision problem requires computational efforts.

(ii) Notice that $n! \leq n^n$, so the upper bound on M_t in this case is of order $n \log_2 n$.