EECS598: Prediction and Learning: It's Only a Game Fall 2013

Lecture 24: Generalized Calibration and Correlated Equilibria

Prof. Jacob Abernethy *Scribe: Yuqing Kong*

24.1 Generalized Calibration

In previous section, we make predictions in [0,1]. [0,1] interval can be generalized to convex set. As before we divide [0,1] into small sections, now we devide the convex set into n small pieces and pick one point q_i in each piece. Now the calibration setting will be generalized to:
For t=1,...,T

1. Forecaster "guesses" $\hat{y}_t$ with q_{i_t}
2. Outcome is y_t
   .

In the end, we want to guarantee that:
$\exists T_0, \forall i, \forall T > T_0, \|\frac{\sum_{t=1}^{T} y_t \mathbb{1}[q_{i_t}=q_i]}{\sum_{t=1}^{T} \mathbb{1}[q_{i_t}=q_i]} - q_i\| < c\epsilon$

With this generalized calibration you can:

1. Get lower regret
2. Get minmax duality
3. Show Approachability Theorem.

24.2 Two players zero-sum game

Consider a repeated zero-sum game between two players.
Given matrix M, two players chooses $(x,y) \in \Delta_n \times \Delta_n$ to get value $x^T My$. Player 1 chooses $x \in \Delta_n$ and wants to minimize $x^T My$ while Player 2 chooses $y \in \Delta_n$ and wants to maximinze $x^T My$. They play this game repeatedly. Consider the following setting:
For t=1,...,T

1. Player 1 chooses $x_t \in \Delta_n$
2. Player 2 chooses $y_t \in \Delta_n$

Let V^* denote $\min_x \max_y (xMy)$
Given any ϵ, we want to find an algorithm such that in the end $\frac{1}{T}\sum_{t=1}^{T} x_t M y_t \leq V^* + O(\epsilon)$,
The idea is to reduce this problem to generalized calibration and use ϵ calibration algorithm. Consider the following algorithm:

Reduction to Calibration:
For t=1,2,...,T

1. Player 1 guesses $q_{i_t} \in \Delta_n$

2. Player 1 computes the best response
$x_t = x(q_{i_t}) = \underset{x \in \Delta_n}{\arg\min} x^T M q_{i_t}$

3. Player 2 reveals y_t

We assume that this algorithm is calibrated and now let's analyze the value $\frac{1}{T}\sum_{t=1}^{T} x_t M y_t$ to see whether it exceeds V^* much:
For the sake of analysis, let n_T^i denote $\sum_{t=1}^{T} \mathbb{1}[q_{i_t} = q_i]$, we can see $\sum_i n_T^i = T$

$$\frac{1}{T}\sum_{t=1}^{T} x_t M y_t = \sum_{i=1}^{N}\left(\frac{1}{T}\sum_{t=1}^{T} x_t M y_t \mathbb{1}[q_{i_t} = q_i]\right) \tag{24.1}$$

$$= \sum_{i=1}^{N}\left(\frac{1}{T}\sum_{t=1}^{T} x(q_i) M y_t \mathbb{1}[q_{i_t} = q_i]\right) \tag{24.2}$$

$$= \sum_{i=1}^{N}\left(\sum_{t=1}^{T} \frac{n_T^i}{T} x(q_i) M \left(\frac{y_t \mathbb{1}[q_{i_t} = q_i]}{n_T^i}\right)\right) \tag{24.3}$$

$$= \sum_{i=1}^{N} \frac{n_T^i}{T} x(q_i) M (q_i + \epsilon U) \tag{24.4}$$

$$= \sum_{i=1}^{N} \frac{n_T^i}{T} x(q_i) M q_i + o(\epsilon) \le V^* + o(\epsilon) \tag{24.5}$$

From line 3 to line 4, we are assuming forecast is calibrated. In line 4, U is a vector and $\|U\| \le 1$. In line 5, $\sum_{i=1}^{N} \frac{n_T^i}{T} x(q_i) M q_i \le V^*$, V^* is the value of game.
So we can see:

Theorem 24.1. *Existence of $\epsilon-$ Nash Equilibrium is reducible to ϵ calibration algorithm.*

24.3 Correlated Equilibrium

Now let's consider a game among k players.
For all i, player i has M_i strategies. Let $[M_i]$ denote the set of the M_i strategies player i can use. Each time k players play $(j_1, j_2, ..., j_k) \in [M_1] \times [M_2] \times ... \times [M_k]$ and then player i would get loss: $C_i(j_1, ..., j_k)$
We assign a joint distribution $\mu \in \Delta([M_1] \times [M_2] \times ... \times [M_k])$ to the actions of k players. Then we can see the expected loss to Player i with distribution μ would be:

$$C_i(\mu) = \sum_{(j_1,...,j_k)} \mu(j_1, j_2, ..., j_k) C_i(j_1, ..., j_k)$$

A *strategy modification* is a function $\phi[M_i] \to [M_i]$ such that $\phi(j) = j$ for all j but one j_o. $\phi(j_o)$ is arbitrary. Then after this modification, the expected loss would change to:

$$C_i^{\phi}(\mu) = \sum_{(j_1,\dots,j_k)} \mu(j_1, j_2, \dots, j_k) C_i(j_1, \dots, j_{i-1}, \phi(j_i), j_{i+1}, \dots j_k)$$

Now we can give the definition of *Correlated Equilibrium*(CE):
Distribution μ is a CE if for all i, $C_i(\mu) \le C_i^{\phi}$ for all modifications ϕ.
Distribution μ is an $\epsilon-$CE if for all i, $C_i(\mu) \le C_i^{\phi}(\mu) + \epsilon$ for all modifications ϕ.

In the past, the loss we analyze is compared to a constant sequence. But now, we can generalize the definition and discuss a loss which is compared to a "class" of sequences. Let's see the definitions of *external regret* and *internal regret*.

- An algorithm(Alg) has no *external regret* if $\mathbb{E}[\frac{1}{T}(\sum l_{I_t} - l_i)] \le \epsilon$ for large T. Here $(i, i, \dots, i)$ is the best constant sequence we can choose in hindsight.
- An algorithm(Alg) has no *internal regret* if for all ϕ, $\mathbb{E}[\frac{1}{T}(\sum l_{I_t} - l_{\phi(I_t)})] \le \epsilon$ for large T. Here $\{(\phi(I_1), \phi(I_2), \dots, \phi(I_T))\}_{\phi}$ are a "class" of sequences compared to our actions.

We know that no-external-regret algorithm can give us an algorithm to get an $\epsilon-$ Nash Equilibrium. Now let's see whether no-internal-regret algorithm can give us an algorithm to get an $\epsilon-$ Correlated Equilibrium and discuss the relation among B.A.T, no-internal-regret algorithm and calibration algorithm.

Theorem 24.2. *Existence of No-Internal Alg is reducible to Black Well Approachibility*

Proof. If we want to use B.A.T, firstly we need to define a vector game. Let's define a biaffine $r : \Delta_n \times [0,1]^n \to \mathbb{R}^{n^2}$
$r(\underline{w}, \underline{l}) =< (l_i - l_j) w_i >_{(i,j) \in [n]^2}$
Then we need to define the set: $S = \mathbb{R}_{-}^{n^2}$
So we need to know whether the assumption of B.A.T is satisfied. In other words, we need to know $\forall \underline{l} \in [0,1]^n$ whether there exist $w \in \Delta_n$ such that $r(w, \underline{l}) \in S$.
The answer is yes, since we can find $w = e_i$ where $i = \arg\min_{i'} l_{i'}$. Now we can use the result of B.A.T, which means given any ϵ we can find an adaptive strategy such that $\exists T_0, \forall T > T_0, d(\frac{1}{T}\sum_{t=1}^{T} < (l_i^t - l_j^t) w_i^t >, S) < \epsilon$. No-internal-regret algorithm requires that $\frac{1}{T}\sum_T \sum_I (l_{I_t} - l_{\phi(I_t)}) w_{I_t} \le \epsilon$, which can be satisfied by the result B.A.T gives us. So we can see we find a no-internal-regret algorithm through Black Well Approachibility.

□

Theorem 24.3. *If all players use a no internal regret algorithm to play then $\bar{\mu}_t$, the empirical distribution of*

$$\{(j_1^1, \dots, j_k^1), (j_1^2, \dots, j_k^2), \dots, (j_1^T, \dots, j_k^T)\}$$

is an $\epsilon-$CE.

Proof. The definition of $\epsilon-$ CE is for all i, for all ϕ
$C_i(\mu) \le C_i^{\phi}(\mu) + \epsilon = \sum_{(j_1,\dots,j_k)} \mu(j_1, j_2, \dots, j_k) C_i(j_1, \dots, j_{i-1}, \phi(j_i), j_{i+1}, \dots j_k) + \epsilon$

If all players use a no-internal-regret algorithm, then for all i, for all ϕ, $\frac{1}{T}\sum_t(C_i(\mu_t) - C_i^{\phi}(\mu_t)) \leq \epsilon$ $\Rightarrow C_i(\bar{\mu}_t) \leq C_i^{\phi}(\bar{\mu}_t) + \epsilon$, which means $\bar{\mu}_t$ is an ϵ−CE

□

Theorem 24.4. *We can reduce calibration to no-internal-regret.*

Proof. The definition of calibration is: $\forall i \ \|\frac{\sum_{t=1}^{T} y_t \mathbb{1}[q_{i_t}=q_i]}{\sum_{t=1}^{T}\mathbb{1}[q_{i_t}=q_i]} - q_i\| < c\epsilon$ for large T. So if the algorithm is not calibrated, then$\exists\epsilon\ \forall T_0$, $\exists T > T_0$ such that $\exists$ a set I for all $i \in I$ $\|\frac{\sum_{t=1}^{T} y_t \mathbb{1}[q_{i_t}=q_i]}{\sum_{t=1}^{T}\mathbb{1}[q_{i_t}=q_i]} - q_i\| > c\epsilon$ but $\|\frac{\sum_{t=1}^{T} y_t \mathbb{1}[q_{i_t}=q_i]}{\sum_{t=1}^{T}\mathbb{1}[q_{i_t}=q_i]} - q_j\| < c\epsilon (j \neq i)$. At this time, if we define a modification ϕ to change strategy from q_i to q_j at time $\{t : q_{i_t} = q_i\}$ for all $i \in I$, then $\sum_i |\frac{1}{T}\sum_t (q_{i_t} - y_t)\mathbb{1}(q_{i_t} = q_i)| - \sum_i |\frac{1}{T}\sum_t(\phi(q_{i_t}) - y_t)\mathbb{1}(\phi(q_{i_t}) = q_i)| > O(\epsilon)$, which means the algorithm has internal regret. By this contradiction, we can reduce calibration to no-internal-regret.

□

So we can see B.A.T $\Rightarrow$ Existence of no internal algorithm $\Rightarrow$ Existence of an ϵ−CE;
No-internal-regret algorithm $\Rightarrow$ Calibration algorithm $\Rightarrow$ an ϵ−NE.
$\Rightarrow$ means "gives".