Mehryar Mohri
Advanced Machine Learning 2016
Courant Institute of Mathematical Sciences
Homework assignment 1
March 7th, 2015
Due: March 21, 2015

A. Deep ensembles

In class, we presented a guarantee for the SRM method. This problem suggests you to derive a similar guarantee for deep ensembles or the principle of voted risk minimization (VRM). We will largely adopt the notation used in class. Let $\mathcal{X}$ denote the input space and let $H_1, \ldots, H_p$ families of functions mapping from $\mathcal{X}$ to $\mathbb{R}$. $\mathcal{F}$ is the convex hull of the union of these families:

$$\mathcal{F} = \operatorname{conv}\left(\bigcup_{k=1}^{p} H_k\right) = \left\{\sum_{j=1}^{N} \alpha_j h_j \colon \alpha_j \geq 0, \sum_{j=1}^{N} \alpha_j \leq 1\right\}.$$

Let F be the objective function defined for all $f = \sum_{j=1}^{N} \alpha_j h_j \in \mathcal{F}$ by

$$F(f) = \widehat{R}_{S,\rho}(f) + \frac{4}{\rho} \sum_{j=1}^{N} \alpha_j \mathfrak{R}_m(H_{k(j)}),$$

where $H_{k(j)}$ is the least complex family containing h_j. Define the VRM solution as the function f_{VRM} minimizing F:

$$f_{\mathrm{VRM}} = \operatorname*{argmin}_{f \in \mathcal{F}} F(f).$$

Use Corollary 1 and Corollary 2 (see appendix) to show the following result. For any $\delta > 0$, with probability at least $1 - \delta$ over the choice of a sample S of size m drawn i.i.d. according to $\mathcal{D}^m$, the following inequality holds for f_{VRM}:

$$R(f_{\mathrm{VRM}}) \leq \inf_{f \in \mathcal{F}} \left(R_\rho(f) + \frac{8}{\rho} \sum_{j=1}^{N} \alpha_j \mathfrak{R}_m(H_{k(j)}) + \frac{4}{\rho} \sqrt{\frac{\log p}{m}} \left[1 + \sqrt{\left\lceil \log \left[\frac{\rho^2 m}{\log p} \right] \right\rceil} \right] \right) + \sqrt{\frac{\log \frac{4}{\delta}}{2m}},$$

with $R_\rho(f) = \mathrm{E}[\Phi_\rho(yf(x))]$, where Φ_ρ is the margin loss function: $\Phi_\rho(x) = \min\left(1, \max\left(0, 1 - \frac{x}{\rho}\right)\right)$ for all $x \in \mathbb{R}$.

Note that the bound applies in particular with the right-hand side function chosen to be f^*, the one minimizing F: $f_{\mathrm{VRM}} = \operatorname{argmin}_{f \in \mathcal{F}} F(f)$.

Solution: Let $C = \frac{2}{\rho}\sqrt{\frac{\log p}{m}}\left[1 + \sqrt{\left\lceil \log\left[\frac{\rho^2 m}{\log p}\right]\right\rceil}\right]$. Then, by Corollary 1 and Corollary 2, the following holds:

$$\begin{aligned}
&\Pr\left[R(f_{\mathrm{VRM}}) - R_\rho(f^*) - \frac{8}{\rho}\sum_{j=1}^{N}\alpha_j^* \mathfrak{R}_m(H_{k(j)}) - 2C > \epsilon\right] \\
&\quad \leq \Pr\left[R(f_{\mathrm{VRM}}) - F(f_{\mathrm{VRM}}) - C > \frac{\epsilon}{2}\right] \\
&\qquad + \Pr\left[F(f_{\mathrm{VRM}}) - R_\rho(f^*) - \frac{8}{\rho}\sum_{j=1}^{N}\alpha_j^* \mathfrak{R}_m(H_{k(j)}) - C > \frac{\epsilon}{2}\right] \\
&\quad \leq 2e^{-\frac{m\epsilon^2}{2}} + \Pr\left[F(f^*) - R_\rho(f^*) - \frac{8}{\rho}\sum_{j=1}^{N}\alpha_j^* \mathfrak{R}_m(H_{k(j)}) - C > \frac{\epsilon}{2}\right] \\
&\quad = 2e^{-\frac{m\epsilon^2}{2}} + \Pr\left[\widehat{R}_S(f^*) - R_\rho(f^*) - \frac{4}{\rho}\sum_{j=1}^{N}\alpha_j^* \mathfrak{R}_m(H_{k(j)}) - C > \frac{\epsilon}{2}\right] \\
&\quad = 2e^{-\frac{m\epsilon^2}{2}} + 2e^{-\frac{m\epsilon^2}{2}} = 4e^{-\frac{m\epsilon^2}{2}}.
\end{aligned}$$

The proof is completed by setting the right-hand side to δ. □

B. Zero-sum games

In class, we gave a proof of von Neumann's minimax theorem by assuming that one of the players was using the RWM algorithm. Consider the scenario where both players use RWM at each round, which we alluded to in class. Prove von Neumann's minimax theorem using that scenario, proceeding as follows.

1. Assume without loss of generality that $u_1 \leq 1$. Let p_t be the distribution defined at the tth round by the row player and q_t the one for the column player. Show that

$$\max_{\mathsf{p}} \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1 \sim \mathsf{p} \\ a_2 \sim \mathsf{q}_t}}[u_1(\mathbf{a})] - \frac{R_T}{T} \leq \min_{\mathsf{q}} \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1 \sim \mathsf{p}_t \\ a_2 \sim \mathsf{q}}}[u_1(\mathbf{a})] + \frac{R_T}{T}.$$

Solution: By the definition of the regret of the RWM algorithm played by the row player, we can write

$$\frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}_t\\ a_2\sim \mathsf{q}_t}}[u_1(\mathbf{a})] \leq \min_{\mathsf{q}} \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}_t\\ a_2\sim \mathsf{q}}}[u_1(\mathbf{a})] + \frac{R_T}{T}.$$

Similarly, for the column player,

$$\begin{aligned}
-\frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}_t\\ a_2\sim \mathsf{q}_t}}[u_1(\mathbf{a})] &\leq \min_{\mathsf{p}} -\frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \mathsf{q}_t}}[u_1(\mathbf{a})] + \frac{R_T}{T}\\
\Leftrightarrow \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}_t\\ a_2\sim \mathsf{q}_t}}[u_1(\mathbf{a})] &\geq \max_{\mathsf{p}} \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \mathsf{q}_t}}[u_1(\mathbf{a})] - \frac{R_T}{T}.
\end{aligned}$$

Combining the inequalities, we obtain

$$\max_{\mathsf{p}} \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \mathsf{q}_t}}[u_1(\mathbf{a})] - \frac{R_T}{T} \leq \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}_t\\ a_2\sim \mathsf{q}_t}}[u_1(\mathbf{a})] \leq \min_{\mathsf{q}} \frac{1}{T}\sum_{t=1}^{T} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}_t\\ a_2\sim \mathsf{q}}}[u_1(\mathbf{a})] + \frac{R_T}{T}.$$

2. Use the previous inequality and the regret bound for RWM to prove von Neumann's minimax theorem.

Solution: The inequality between the left-most and right-most terms can be rewritten as follows

$$\begin{aligned}
\max_{\mathsf{p}} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \frac{1}{T}\sum_{t=1}^{T}\mathsf{q}_t}}[u_1(\mathbf{a})] - \frac{R_T}{T} &\leq \min_{\mathsf{q}} \mathop{\mathrm{E}}_{\substack{a_1\sim \frac{1}{T}\sum_{t=1}^{T}\mathsf{p}_t\\ a_2\sim \mathsf{q}}}[u_1(\mathbf{a})] + \frac{R_T}{T}\\
\Rightarrow \max_{\mathsf{p}} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \frac{1}{T}\sum_{t=1}^{T}\mathsf{q}_t}}[u_1(\mathbf{a})] - \frac{R_T}{T} &\leq \max_{\mathsf{p}}\min_{\mathsf{q}} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \mathsf{q}}}[u_1(\mathbf{a})] + \frac{R_T}{T}\\
\Rightarrow \min_{\mathsf{q}}\max_{\mathsf{p}} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \mathsf{q}}}[u_1(\mathbf{a})] - \frac{R_T}{T} &\leq \max_{\mathsf{p}}\min_{\mathsf{q}} \mathop{\mathrm{E}}_{\substack{a_1\sim \mathsf{p}\\ a_2\sim \mathsf{q}}}[u_1(\mathbf{a})] + \frac{R_T}{T}.
\end{aligned}$$

Using $\frac{R_T}{T} = O(1/\sqrt{T})$ and taking $T \to +\infty$ proves the result. $\square$

C. Correlated Equilibria

Consider the game defined by the following matrix:

	A	B
A	(8, 8)	(1, 9)
B	(9, 1)	(0, 0)

1. Which are the pure Nash equilibria for this game?

   Solution:

   Consider first (B, B). If we are the row player and column player is playing B then we are better off playing A since that gives a payoff of 1 against 0. Therefore, (B, B) is not a Nash equilibrium.

   Next consider (A, A). If we are the row player and the column player is playing A then we are better off playing B since this gives a payoff of 9 against 8 if we are playing A. Therefore, (A, A) is not a Nash equilibrium.

   Now consider (A, B). If we are the row player again and column player is playing B, then we do not have incentive to switch from A to B since in that case we will have a smaller payoff (0 against 1). If we are a column player and row player is playing A then we do not have any incentive to switch from B to A since that leads to a smaller payoff (8 against 9). Therefore, (A, B) is a Nash equilibrium.

   By symmetry, we conclude that (B, A) is also a Nash equilibrium.

2. Find a mixed Nash equilibrium. Which is the expected payoff for the row player?

   Solution: Suppose $(\mathsf{p}_{\text{row}}, \mathsf{p}_{\text{col}})$ is a mixed Nash equilibrium. Then either $\mathsf{p}_{\text{row}}(A) > 0$ and $\mathsf{p}_{\text{row}}(B) > 0$ or $\mathsf{p}_{\text{col}}(A) > 0$ and $\mathsf{p}_{\text{col}}(B) > 0$. We consider the case that $\mathsf{p}_{\text{row}}(A) > 0$ and $\mathsf{p}_{\text{row}}(B) > 0$. Then the expected payoff of playing A and playing B must be the same for the row player assuming $(\mathsf{p}_{\text{row}}, \mathsf{p}_{\text{col}})$ is a mixed Nash equilibrium. Otherwise, if A has a higher payoff then row player has an incentive to switch to pure strategy only playing A which would contradict the fact that $(\mathsf{p}_{\text{row}}, \mathsf{p}_{\text{col}})$ is a mixed Nash equilibrium.

   We use this observation to figure out p_{col}. Note that if $\mathsf{p}_{\text{col}}(A) = p$ this implies that

$$p \cdot 8 + (1 - p) \cdot 1 = p \cdot 9 + (1 - p) \cdot 0$$

which leads to $p = 1/2$. A symmetric argument leads us to conclude that $((1/2, 1/2), (1/2, 1/2))$ is a mixed Nash equilibrium.

The expected payoff of the row player is $9/2$.

3. Suppose now that a correlation device draws each of (A, A), (A, B), (B, A) with equal probability. Prove that this defines a correlated equilibrium. What is the expected payoff? How does it compare to the expected payoff for a mixed Nash equilibrium found in the previous question?

   Solution:

   We first show that this is a correlated equilibrium. Consider row first player. We have the following:

$$\frac{1}{3} \cdot u(A, A) + \frac{1}{3} \cdot u(A, B) = \frac{10}{3} \geq \frac{1}{3} = \frac{1}{3} \cdot u(A, A) + \frac{1}{3} \cdot u(B, A)$$
$$\frac{1}{3} \cdot u(B, A) + 0 \cdot u(B, B) = 3 \geq \frac{8}{3} = \frac{1}{3} \cdot u(A, A) + 0 \cdot u(B, A),$$

   where u denotes a payoff of the row player. Since the same set of linear inequalities holds for column player as well by symmetry, we conclude that this is indeed a correlated equilibrium.

   The expected payoff of the row player in this case is

$$\frac{1}{3} \cdot u(A, A) + \frac{1}{3} \cdot u(A, B) + \frac{1}{3} \cdot u(B, A) = 6$$

   which is higher than for a mixed Nash equilibrium found in the previous question.

   This a version of the so-called Chicken Game where each player either dares or chickens out. This game is famous for his relevance to the Cuban missile crisis.

A Appendix

Corollary 1 *Assume $p > 1$. Fix $\rho > 0$. Then, for any $\delta > 0$, with probability at least $1 - \delta$ over the choice of a sample S of size m drawn i.i.d. according to $\mathcal{D}^m$, the following inequality holds for all $f = \sum_{j=1}^{N} \alpha_j h_j \in \mathcal{F}$:*

$$R(f) \leq \widehat{R}_{S,\rho}(f) + \frac{4}{\rho} \sum_{j=1}^{N} \alpha_j \mathfrak{R}_m(H_{k(j)}) + \frac{2}{\rho}\sqrt{\frac{\log p}{m}} \left[1 + \sqrt{\left\lceil \log \left[\frac{\rho^2 m}{\log p}\right] \right\rceil}\right] + \sqrt{\frac{\log \frac{2}{\delta}}{2m}}.$$

Corollary 2 *Assume $p > 1$. Fix $\rho > 0$. Then, for any $\delta > 0$, with probability at least $1 - \delta$ over the choice of a sample S of size m drawn i.i.d. according to $\mathcal{D}^m$, the following inequality holds for all $f = \sum_{j=1}^{N} \alpha_j h_j \in \mathcal{F}$:*

$$\widehat{R}_{S,\rho}(f) \leq R_\rho(f) + \frac{4}{\rho} \sum_{j=1}^{N} \alpha_j \mathfrak{R}_m(H_{k(j)}) + \frac{2}{\rho}\sqrt{\frac{\log p}{m}} \left[1 + \sqrt{\left\lceil \log \left[\frac{\rho^2 m}{\log p}\right] \right\rceil}\right] + \sqrt{\frac{\log \frac{2}{\delta}}{2m}}.$$