The following homework will be discussed on June 5th and June 8th:

H7.1. (Ada-Boost) Let $\mathcal{F} \subset \{-1,1\}^{\mathbb{R}^d}$ be the class of functions that correspond to single Perceptron-like artificial neurons with activation function $\sigma(z) = \text{sgn}(z)$ and d inputs. Let AdaBoost run on this set as set of base hypotheses for T rounds. Can you represent the resulting function in terms of a neural network?

H7.2. (**Multiplicative Weights and Adaptive Game Playing**) In this exercise we study the problem of learning to play a repeated Matrix game by using an ensemble method. Let M be a matrix. On each of a series of rounds, one player, which we call *learner*, chooses a row i and an adversary chooses a column j of M. The entry $M_{ij} \in [0,1]$ is the loss suffered by the row player. The *expected loss* is defined as

$$M(p,q) := \langle p, Mq \rangle = \sum_{ij} p(i) M_{ij} q(j)$$

where p an q are probability vectors representing mixed strategies of the two players. The game matrix M is fixed, but unknown to the learner. On each round $t = 1, \dots, T$:

1. the learner chooses a mixed strategy p_t.
2. the adversary chooses mixed strategy q_t. This strategy may be chosen with knowledge of p_t
3. the learner is permitted to observe the loss $M(e_i, q_t)$ for all unit standard vectors e_i
4. the learner suffers loss $M(p_t, q_t)$

The main goal of the learner is to minimize its total loss $\sum_{t=1}^T M(p_t, q_t)$. We describe now the following ensemble method known as *multiplicative weights algorithm (MWA)*. The learner starts with some initial mixed strategy p_1. After each round t, he updates his mixed strategy by computing

$$p_{t+1}(i) = p_t(i) \frac{\beta^{M(e_i, q_t)}}{Z_t},$$

where Z_t is a normalization factor and $\beta \in [0,1)$ is a parameter of the algorithm.

a) Prove that for any iteration t where the MWA is used with parameter β and for any mixed strategy p

$$KL(p||p_{t+1}) - KL(p||p_t) \leq \left(\log \frac{1}{\beta} \right) M(p, q_t) + \log \left(1 - (1-\beta) M(p_t, q_t) \right).$$

The inequality $\beta^x \leq 1 - (1-\beta)x$ for $\beta \geq 0$ and $x \in [0,1]$ might be useful.

b) Show that for any matrix M with n rows and entries in $[0,1]$ and for any sequence of mixed strategies $q_1, \dots, q_T$ played by the adversary, the sequence of mixed strategies $p_1, \dots, p_T$ produced by the MWA satisfies

$$\sum_{t=1}^T M(p_t, q_t) \leq \min_p \left[a_\beta \sum_{t=1}^T M(p, q_t) + c_\beta KL(p||p_1) \right],$$

where $a_\beta = \frac{\log(1/\beta)}{1-\beta}$ and $c_\beta = \frac{1}{1-\beta}$.
Hint: $\log(1-x) \leq -x$ for $x < 1$.
Remark: this means that the total loss from the MWA is "not much worse" than the total loss of the best strategy in hindsight, namely $\min_p \sum_{t=1}^{T} M(p, q_t)$.

c) Show that if the MWA is used with p_1 equal to the uniform distribution and

$$\beta = \left(1 + \sqrt{\frac{2 \log n}{T}}\right)^{-1},$$

then

$$\frac{1}{T}\sum_{t=1}^{T} M(p_t, q_t) \leq \min_p \frac{1}{T}\sum_{t=1}^{T} M(p, q_t) + \sqrt{\frac{2\log n}{T}} + \frac{\log n}{T}.$$

Hint: $-\beta \log \beta \leq (1-\beta^2)/2$ for $\beta \in (0, 1]$.

d) What are the main differences of the described MWA and the Ada-boost?