

MA4270 Data Modelling and Computation

Preliminaries

- Variance** $= \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$
- Markov ineq.:** Given $a > 0$ and X nonnegative, $\Pr(X \geq a) \leq \frac{\mathbb{E}[X]}{a}$
- Chebyshev ineq.:** $\Pr(|Y - \mathbb{E}[Y]| \geq b) \leq \frac{\sigma^2}{b^2}$

Binary Classification

- Training data:** $\mathcal{D} = \{(\mathbf{x}_t, y_t)\}_{t=1}^n$
 - where $y_t \in \{-1, 1\}$ and $\mathbf{x}_t \in \mathbb{R}^d$
 - where we have n training data
- Classifier:** $f_\theta : \mathbb{R}^d \rightarrow \{-1, 1\}$
- Training error:** $\hat{E}(\theta) = \frac{1}{n} \sum_{t=1}^n \text{Loss}(y_t, f_\theta(\mathbf{x}_t))$
 - i.e. the proportion of failed training data
- $\text{Loss}(y, \hat{y}) = \mathbf{1}\{\hat{y} \neq y\}$

Linear classifier

- Formula:** $\hat{y} := \text{sgn}(\theta, \mathbf{x}) := \text{sgn} \sum_{i=1}^d \theta_i x_i$ (for some fixed θ
 - i.e. partitioning space with plane passing through origin
 - \mathcal{D} is linearly separable: $\exists \theta$ such that $\hat{E}(\theta) = 0$
 - Perceptron update algorithm:**
 - current estimate θ_{curr}
 - look at one (\mathbf{x}_t, y_t)
 - if $\text{sgn}(\theta_{curr}, \mathbf{x}_t) = y_t$ then $\theta_{next} := \theta_{curr}$
 - otherwise set $\theta_{next} := \theta_{curr} + y_t \mathbf{x}_t$ (intuitively, we are increasing the value of $y_t \theta_{next}^T \mathbf{x}_t$)
 - Perceptron algorithm:**
 - init $\theta^{(0)} := \mathbf{0}, k := 0$
 - For each data point, run the update, and increment k only when we made a mistake (repeat after reaching n)
 - stop when we do full pass of all data without mistakes
 - Perceptron algorithm assumptions:**
 - $\exists R > 0$ such that $\forall t, \|\mathbf{x}_t\| \leq R$ (i.e. \mathbf{x}_t is uniformly bdd)
 - $\exists \theta^*$ and $\exists \gamma > 0$ such that $\min_t y_t (\theta^*)^T \mathbf{x}_t \geq \gamma$ (i.e. stronger lin. separable cond. - bounded away from zero)
- Thm:** under above assumptions, the algorithm finds $\theta^{(k)}$ such that $\hat{E}(\theta^{(k)}) = 0$ after at most $k_{max} = \frac{R^2 \|\theta^*\|^2}{\gamma^2}$ steps (we make at most k_{max} mistakes)
- Cauchy-Schwarz inequality:** $|\langle \mathbf{u}, \mathbf{v} \rangle| \leq \|\mathbf{u}\| \cdot \|\mathbf{v}\|$ (with equality when \mathbf{u} and \mathbf{v} are in the same direction)

Perceptron proof:

- show that $\langle \theta^*, \theta^{(k+1)} \rangle \geq \langle \theta^*, \theta^{(k)} \rangle + \gamma$, so $\langle \theta^*, \theta^{(k)} \rangle \geq k\gamma$
 - show that $\|\theta^{(k+1)}\|^2 \leq \|\theta^{(k)}\|^2 + R^2$, so $\|\theta^{(k)}\|^2 \leq kR^2$
 - use Cauchy-Schwartz ineq. to conclude that $k \leq \frac{R^2 \|\theta^*\|^2}{\gamma^2}$
- Details:
- $(\theta^*)^T \theta^{(k+1)} = (\theta^*)^T (\theta^{(k)} + y_t \mathbf{x}_t) = (\theta^*)^T \theta^{(k)} + y_t (\theta^*)^T \mathbf{x}_t \geq (\theta^*)^T \theta^{(k)} + \gamma$
 - $\|\theta^{(k+1)}\|^2 = \|\theta^{(k)} + y_t \mathbf{x}_t\|^2 = \|\theta^{(k)}\|^2 + 2\langle \theta^{(k)}, y_t \mathbf{x}_t \rangle + \|\mathbf{x}_t\|^2 \leq \|\theta^{(k)}\|^2 + \|\mathbf{x}_t\|^2$ (since when a mistake occurs we must have $\langle \theta^{(k)}, y_t \mathbf{x}_t \rangle \leq 0$)

$$3. 1 \geq \frac{\langle \theta^*, \theta^{(k)} \rangle}{\|\theta^*\| \cdot \|\theta^{(k)}\|} \geq \frac{k\gamma}{\|\theta^*\| \cdot \sqrt{kR^2}} = \frac{\sqrt{k}\gamma}{\|\theta^*\| \cdot R}, \text{ so } k \leq \frac{R^2 \|\theta^*\|^2}{\gamma^2}$$

Perceptron with offset:

$$\text{Set } \hat{\theta} := \begin{bmatrix} \theta \\ \theta_0 \end{bmatrix} \text{ and } \hat{\mathbf{x}} := \begin{bmatrix} \mathbf{x} \\ 1 \end{bmatrix},$$

$$\text{so that } \hat{\theta}^T \hat{\mathbf{x}} = \theta^T \mathbf{x} + \theta_0 \text{ and } \hat{R}^2 = R^2 + c^2$$

- (Non-geometric) margin** (w.r.t. θ^*): $\gamma := \min_t y_t (\theta^*)^T \mathbf{x}_t$ (note: usually ‘margin’ means the geometric margin)
 - Geometric margin** (normalised by $\|\theta^*\|$): $\gamma_{geom} := \frac{\gamma}{\|\theta^*\|}$ (= the shortest distance from a point to the hyperplane)
 - Maximum margin classifier** (a kind of support vector machine): find the classifier that makes the margin largest
 - maximise (over θ, γ) $\frac{\gamma}{\|\theta\|}$ subject to $\forall t, y_t \theta^T \mathbf{x}_t \geq \gamma$
 - equiv. minimise (over new θ) $\frac{1}{2} \|\theta\|^2$ subject to $\forall t, y_t \theta^T \mathbf{x}_t \geq 1$
 - which is a convex optimisation problem solvable efficiently
 - then the classifier $f_\theta(\mathbf{x}) = \text{sgn}(\theta^T \mathbf{x})$ is max margin
 - and $\gamma_{geom} = \frac{1}{\|\theta^*\|}$ (new θ)
- Proof of uniqueness:** can show that if we have distinct θ_1, θ_2 both optimal, then their average most also be optimal and hence $\|\theta_1 - \theta_2\|^2 = 0$

- Support vector:** the \mathbf{x} s that touch the margin

Linear classifier with offset

- Formula:** $\hat{y} := \text{sgn}(\langle \theta, \mathbf{x} \rangle + \theta_0)$ where $\theta \in \mathbb{R}^d$ and $\theta_0 \in \mathbb{R}$
- SVM with offset:** minimise (over θ, θ_0) $\frac{1}{2} \|\theta\|^2$ subject to $\forall t, y_t (\theta^T \mathbf{x}_t + \theta_0) \geq 1$ (note: we don’t penalise for θ_0)
- Soft-margin SVM:** allows misclassified points by introducing slack variables (ζ):
 - minimise (over θ, θ_0, ζ) $\frac{1}{2} \|\theta\|^2 + C \sum_{t=1}^n \zeta_t$ subject to $\forall t, y_t (\theta^T \mathbf{x}_t + \theta_0) \geq 1 - \zeta_t$ and $\forall t, \zeta_t \geq 0$ (allows us to pay a measured penalty for points too close to the margin or misclassified)
 - equiv. minimise (over θ, θ_0) $\frac{1}{2} \|\theta\|^2 + C \sum_{t=1}^n \left[1 - y_t (\theta^T \mathbf{x}_t + \theta_0) \right]_+$ where $[z]_+ := \max\{0, z\}$
 - higher $C \implies$ favour fewer violations
 - $C \rightarrow \infty$ and \mathcal{D} is lin. separable \implies hard-margin SVM
- Support vector:** the \mathbf{x} s that touch the margin, are inside the margin, or misclassified
 - moving these points will change the line
 - if we remove all points except those from support vector, the classifier will remain the same

- Loss: (where $z = y_t (\theta^T \mathbf{x}_t + \theta_0)$)
 - 0-1 loss: $\text{Loss}(z) := \mathbf{1}\{z \leq 0\}$
 - Hinge loss: $\text{Loss}_h(z) := [1 - z]_+$ (it is convex)

Logistic regression

- Logistic function:**

$$g(z) := \frac{1}{1 + e^{-z}}$$



- Soft decision:** We answer using a probability space in $[0, 1]$ (logistic regression is a soft-decision algorithm)

- Probability** (given θ, θ_0):

$$P(y = 1 | \mathbf{x}) = g(\theta^T \mathbf{x} + \theta_0) = \frac{1}{1 + \exp(-(\theta^T \mathbf{x} + \theta_0))}$$

$$P(y = -1 | \mathbf{x}) = 1 - P(y = 1 | \mathbf{x}) = \frac{1}{1 + \exp(\theta^T \mathbf{x} + \theta_0)}$$

$$\text{Hence } P(y | \mathbf{x}) = g(y \cdot (\theta^T \mathbf{x} + \theta_0))$$

$$\text{Note: } \text{Log} \left(\frac{P(y=1|\mathbf{x})}{P(y=-1|\mathbf{x})} \right) = \theta^T \mathbf{x} + \theta_0$$

- Multiplying θ and θ_0 by a large constant makes the probability further away from 0.5 (i.e. the graph gets steeper, but the decision boundary remains the same)
- To determine θ and θ_0 ,** we generally use ‘maximum likelihood’:
 - $L(\theta, \theta_0 | \mathcal{D}) = \prod_{t=1}^n P(y_t | \mathbf{x}_t; \theta, \theta_0)$ (generally good enough if n is large)
 - Equiv. $L(\theta, \theta_0 | \mathcal{D}) = P(y_1, \dots, y_n | \mathbf{x}_1, \dots, \mathbf{x}_n; \theta, \theta_0)$
 - max likelihood: $(\theta, \theta_0) = \arg \max_{\theta, \theta_0} L(\theta, \theta_0 | \mathcal{D})$
 - To solve: $\arg \max_{\theta, \theta_0} L(\theta, \theta_0 | \mathcal{D}) = \arg \max_{\theta, \theta_0} \prod_{t=1}^n P(y_t | \mathbf{x}_t; \theta, \theta_0) = \arg \max_{\theta, \theta_0} \sum_{t=1}^n \text{Log} P(y_t | \mathbf{x}_t; \theta, \theta_0) = \arg \min_{\theta, \theta_0} \sum_{t=1}^n \text{Log}(1 + \exp(-y_t (\theta^T \mathbf{x}_t + \theta_0)))$
 - ... which can only be solved numerically using gradient descent, because it is convex
 - We can call it a ‘loss’:
 - Logistic loss: $\widetilde{\text{Loss}}(z) = \text{Log}(1 + e^{-z})$

- Gradient descent for max likelihood logistic regression:**
 - $\frac{\partial}{\partial \theta_0} \text{Log}(1 + \exp(-y_t (\theta^T \mathbf{x}_t + \theta_0))) = -y_t (1 - P(y_t | \mathbf{x}_t; \theta, \theta_0))$
 - $\frac{\partial}{\partial \theta} \text{Log}(1 + \exp(-y_t (\theta^T \mathbf{x}_t + \theta_0))) = -y_t \mathbf{x}_t (1 - P(y_t | \mathbf{x}_t; \theta, \theta_0))$
 - Thus: $\theta_0^{(i+1)} \leftarrow \theta_0^{(i)} + \eta \sum_{t=1}^n y_t (1 - P(y_t | \mathbf{x}_t; \theta^{(i)}, \theta_0^{(i)}))$
 - $\theta^{(i+1)} \leftarrow \theta^{(i)} + \eta \sum_{t=1}^n y_t \mathbf{x}_t (1 - P(y_t | \mathbf{x}_t; \theta^{(i)}, \theta_0^{(i)}))$

- Stochastic gradient descent:**

- For large data sets, iterating all points at every step is too costly, so we instead pick one (or a few) data points only (at random, or cycle, etc):
 - $\theta_0^{(i+1)} \leftarrow \theta_0^{(i)} + \eta y_t (1 - P(y_t | \mathbf{x}_t; \theta^{(i)}, \theta_0^{(i)}))$
 - $\theta^{(i+1)} \leftarrow \theta^{(i)} + \eta y_t \mathbf{x}_t (1 - P(y_t | \mathbf{x}_t; \theta^{(i)}, \theta_0^{(i)}))$
- For separable data, max likelihood will lead to an arbitrarily steep logistic regression function (i.e. $L(\theta, \theta_0 | \mathcal{D}) \rightarrow 1$ as $\theta, \theta_0 \rightarrow \infty$ is the max likelihood), so we want to prevent this using regularisation

- Regularisation:** Change the expression to $\arg \min_{\theta, \theta_0} \sum_{t=1}^n \text{Log}(1 + \exp(-y_t (\theta^T \mathbf{x}_t + \theta_0))) + \frac{\lambda}{2} \|\theta\|^2$ (where λ is the regularisation parameter) (prevents overly large θ)

- Logistic regression can be extended to multi-class classification

Gradient descent

- Def:** Want to find $\hat{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^d} f(\mathbf{z})$ for some evaluation function f
- Algorithm:** At the current point, evaluate the current gradient, and move a bit in the steepest downhill direction
 - update operation: $\mathbf{z}^{(i+1)} \leftarrow \mathbf{z}^{(i)} - \eta \nabla f(\mathbf{z}^{(i)})$

$$\text{where } \nabla f(\mathbf{z}^{(i)}) := \begin{bmatrix} \frac{\partial f}{\partial z_1} \\ \vdots \\ \frac{\partial f}{\partial z_n} \end{bmatrix} \text{ and } \eta \text{ is the } \underline{\text{step size}}$$

- guaranteed to find the minimum if the function is convex

Linear regression

- The output is now over \mathbb{R}

- Training data:** $\mathcal{D} = \{(\mathbf{x}_t, y_t)\}_{t=1}^n$
 - where $y_t \in \mathbb{R}$ and $\mathbf{x}_t \in \mathbb{R}^d$

- Objective:** To learn a prediction rule (line): $\hat{y} = \theta^T \mathbf{x} + \theta_0$

- Gaussian model:** Add some noise, because the relationship may not be perfect
 - $y = (\theta^*)^T \mathbf{x} + \theta_0^* + z$, where $z \sim N(0, \sigma^2)$
 - PDF of a normal distribution: $\mathcal{N}(z; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(z-\mu)^2}{2\sigma^2}\right)$
 - So $P(y | \mathbf{x}) = \mathcal{N}(y; (\theta^*)^T \mathbf{x} + \theta_0^*, \sigma^2)$

- Using maximum likelihood to derive least squares formula** (works for Gaussian noise model only):

$$L(\theta, \theta_0, \sigma^2 | \mathcal{D}) = P(y_1, \dots, y_n | \mathbf{x}_1, \dots, \mathbf{x}_n; \theta, \theta_0, \sigma^2) = \prod_{t=1}^n P(y_t | \mathbf{x}_t; \theta, \theta_0, \sigma^2) = \prod_{t=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(y_t - \theta^T \mathbf{x}_t - \theta_0)^2}{2\sigma^2}\right)$$

$$\text{Log}(L(\theta, \theta_0, \sigma^2 | \mathcal{D})) = -\frac{n}{2} \text{Log}(2\pi\sigma^2) - \frac{1}{2\sigma^2} \sum_{t=1}^n (y_t - \theta^T \mathbf{x}_t - \theta_0)^2$$

$$\dots \text{ so } \theta \text{ and } \theta_0 \text{ do not depend on } \sigma^2$$

... so equiv. to finding

$$(\hat{\theta}, \hat{\theta}_0) = \arg \min_{\theta, \theta_0} \sum_{t=1}^n (y_t - (\theta^T \mathbf{x}_t + \theta_0))^2$$

(note that $(\theta^T \mathbf{x}_t + \theta_0)$ is simply the predicted value of y

using the line defined by $(\hat{\theta}, \hat{\theta}_0)$, so it is the formula for least squares regression)

- Closed form solution** (for the prediction/estimate):

$$\text{Solve } J(\Theta) := \sum_{t=1}^n (y_t - (\theta^T \mathbf{x}_t + \theta_0))^2 = \|\mathbf{y} - \mathbf{X}\Theta\|^2, \text{ where:}$$

$$-\mathbf{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} \in \mathbb{R}^n; \mathbf{X} = \begin{bmatrix} \mathbf{x}_1^T & 1 \\ \vdots & \vdots \\ \mathbf{x}_n^T & 1 \end{bmatrix} \in \mathbb{R}^{n \times (d+1)}; \Theta = \begin{bmatrix} \theta \\ \theta_0 \end{bmatrix} \in \mathbb{R}^{d+1}$$

then to solve for min point, we make the gradient equal 0:

$$\nabla \left(\|\mathbf{y} - \mathbf{X}\Theta\|^2 \right) = -2\mathbf{X}^T (\mathbf{y} - \mathbf{X}\Theta) = \mathbf{0}$$

$$\implies \mathbf{X}^T \mathbf{y} = \mathbf{X}^T \mathbf{X} \Theta$$

$$\implies \Theta = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y} \text{ (so } \Theta \text{ is linear in } \mathbf{y} \text{ but not } \mathbf{X})$$

(($\mathbf{X}^T \mathbf{X}$)⁻¹ \mathbf{X}^T is called the pseudo-inverse of \mathbf{X})

$$\text{Prediction rule: } \hat{y} = \hat{\theta}^T \mathbf{x} + \hat{\theta}_0$$


- Maximum likelihood with high n : low bias and low variance
- Maximum likelihood with low n : low bias and high variance

Bias & Variance

- Bias:** systematic error
- Variance:** random error

- we want to determine good $\theta_1, \dots, \theta_M$

- **Adaboost algorithm** for learning $\theta_1, \dots, \theta_M$ and $\alpha_1, \dots, \alpha_M$
 - **Input:** $\mathcal{D} = \{(\mathbf{x}_t, y_t)\}_{t=1}^n$, number of iterations M
 - 1. Init weights $w_0(t) := \frac{1}{n}$ for $t = 1, \dots, n$
 - 2. For $m = 1, \dots, M$:
 - a. Choose $h(\cdot; \hat{\theta}_m)$ as $\hat{\theta}_m := \arg \min_{\theta} \sum_{y_t \neq h(\mathbf{x}_t; \theta)}^{t=1} w_{m-1}(t)$ (i.e. minimise the weight of all misclassified points) (note: $\arg \min_{\theta} \sum_{y_t \neq h(\mathbf{x}_t; \theta)}^{t=1} w_{m-1}(t) = \arg \min_{\theta} \sum_{t=1}^n (-y_t h(\mathbf{x}_t; \theta)) w_{m-1}(t)$ since $-y_t h(\mathbf{x}_t; \theta) = 2 \cdot \mathbf{1}\{y_t \neq h(\mathbf{x}_t; \theta)\} - 1$)
 - b. $\hat{\alpha}_m := \frac{1}{2} \log \frac{1 - \hat{\varepsilon}_m}{\hat{\varepsilon}_m}$ where $\hat{\varepsilon}_m := \sum_{y_t \neq h(\mathbf{x}_t; \hat{\theta}_m)}^{t=1} w_{m-1}(t)$



(plot of $\hat{\alpha}_m$ against $\hat{\varepsilon}_m$.)

- c. Update weights: $w_m(t) := \frac{1}{Z_m} w_{m-1}(t) e^{-y_t h(\mathbf{x}_t; \hat{\theta}_m) \hat{\alpha}_m}$ where $Z_m := \sum_{t=1}^n w_{m-1}(t) e^{-y_t h(\mathbf{x}_t; \hat{\theta}_m) \hat{\alpha}_m}$ is the normalization factor (i.e. we want $\sum_{t=1}^n w_m(t) = 1$) (note: $e^{\hat{\alpha}_m} > 1$ and $e^{-\hat{\alpha}_m} < 1$ since $\hat{\alpha}_m > 0$, so it increases relative weights for misclassified points w.r.t. current decision stump)

3. The output classifier is $f_M(\mathbf{x}) := \sum_{m=1}^M \hat{\alpha}_m h(\mathbf{x}; \hat{\theta}_m)$

• Note: Adaboost never overfits points for some reason

- **Training error:** fraction of misclassified points: $\frac{1}{n} \sum_{t=1}^n \mathbf{1}\{y_t f_m(\mathbf{x}) \leq 0\}$
- **Theorem:** After M iterations, training error $\leq \exp(-2 \sum_{m=1}^M (\frac{1}{2} - \hat{\varepsilon}_m)^2)$ (in particular, if $\hat{\varepsilon}_m \leq \frac{1}{2} - \gamma \forall m$, then training error $\leq \exp(-2M\gamma^2)$; and since training error must be a multiple of $\frac{1}{n}$, if $\exp(-2M\gamma^2) < \frac{1}{n}$ then training error must be zero (i.e. everything classified correctly))

• **Adaboost proof:**

1. Bound 0-1 loss with the exponential loss: $\frac{1}{n} \sum_{t=1}^n \mathbf{1}\{y_t f_M(\mathbf{x}) \leq 0\} \leq \frac{1}{n} \sum_{t=1}^n \exp(-y_t f_M(\mathbf{x}))$ (proof: it is obvious)
2. By weight update formula, $w_M(t) = \frac{1}{n} \prod_{m=1}^M \frac{\exp(-y_t h(\mathbf{x}_t; \hat{\theta}_m) \hat{\alpha}_m)}{Z_m} = \frac{1}{n} \frac{\exp(-\sum_{m=1}^M y_t h(\mathbf{x}_t; \hat{\theta}_m) \hat{\alpha}_m)}{\prod_{m=1}^M Z_m} = \frac{1}{n} \frac{\exp(-y_t f_M(\mathbf{x}))}{\prod_{m=1}^M Z_m}$ Hence $1 = \sum_{t=1}^n w_M(t) = \frac{1}{n} \frac{\sum_{t=1}^n \exp(-y_t f_M(\mathbf{x}))}{\prod_{m=1}^M Z_m}$, and thus $\frac{1}{n} \sum_{t=1}^n \exp(-y_t f_M(\mathbf{x})) = \prod_{m=1}^M Z_m$ Combining with step 1, $\frac{1}{n} \sum_{t=1}^n \mathbf{1}\{y_t f_M(\mathbf{x}) \leq 0\} \leq \prod_{m=1}^M Z_m$
3. $Z_m = \sum_{t=1}^n w_{m-1}(t) e^{-y_t h(\mathbf{x}_t; \hat{\theta}_m) \hat{\alpha}_m} = \sum_{t: y_t \neq h(\mathbf{x}_t; \hat{\theta}_m)} w_{m-1}(t) e^{\hat{\alpha}_m} + \sum_{t: y_t = h(\mathbf{x}_t; \hat{\theta}_m)} w_{m-1}(t) e^{-\hat{\alpha}_m} = e^{\hat{\alpha}_m} \sum_{t: y_t \neq h(\mathbf{x}_t; \hat{\theta}_m)} w_{m-1}(t) + e^{-\hat{\alpha}_m} \sum_{t: y_t = h(\mathbf{x}_t; \hat{\theta}_m)} w_{m-1}(t) = \hat{\varepsilon}_m e^{\hat{\alpha}_m} + (1 - \hat{\varepsilon}_m) e^{-\hat{\alpha}_m}$ So we want to pick $\hat{\alpha}_m$ to minimise Z_m (proof skipped... the optimum $\hat{\alpha}_m = \frac{1}{2} \log \frac{1 - \hat{\varepsilon}_m}{\hat{\varepsilon}_m}$)
4. Substitute $\hat{\alpha}_m$ to step 3, and get $Z_m = \hat{\varepsilon}_m \sqrt{\frac{1 - \hat{\varepsilon}_m}{\hat{\varepsilon}_m}} + (1 - \hat{\varepsilon}_m) \sqrt{\frac{\hat{\varepsilon}_m}{1 - \hat{\varepsilon}_m}} = 2\sqrt{\hat{\varepsilon}_m(1 - \hat{\varepsilon}_m)}$

$$\sqrt{1 - (1 - 2\hat{\varepsilon}_m)^2} = \exp\left(\frac{1}{2} \log\left(1 - (1 - 2\hat{\varepsilon}_m)^2\right)\right) \leq \exp\left(-\frac{1}{2}(1 - 2\hat{\varepsilon}_m)^2\right) = \exp\left(-2\left(\frac{1}{2} - \hat{\varepsilon}_m\right)^2\right)$$

(note: to remove log, use formula $\log(1 + a) \leq a \forall a \in \mathbb{R}$)

5. Substitute into step 2 end, we get $\frac{1}{n} \sum_{t=1}^n \mathbf{1}\{y_t f_M(\mathbf{x}) \leq 0\} \leq \exp\left(-2 \sum_{m=1}^M \left(\frac{1}{2} - \hat{\varepsilon}_m\right)^2\right)$
- **Claim:** $\sum_{t: y_t \neq h(\mathbf{x}_t; \hat{\theta}_m)} w_m(t) = \frac{1}{2}$ (hence the same $\hat{\theta}$ will not be chosen twice in a row)
Proof: LHS $= \frac{1}{2} \iff \sum_{t=1}^n (-y_t h(\mathbf{x}_t; \theta)) w_m(t) = 0$, and we have $\sum_{t=1}^n (-y_t h(\mathbf{x}_t; \theta)) w_m(t) = \sum_{t: y_t = h(\mathbf{x}_t; \hat{\theta}_m)} w_m(t) + \sum_{t: y_t \neq h(\mathbf{x}_t; \hat{\theta}_m)} (-w_m(t)) = \sum_{t: y_t = h(\mathbf{x}_t; \hat{\theta}_m)} \frac{1}{Z_m} w_{m-1}(t) e^{-\hat{\alpha}_m} - \sum_{t: y_t \neq h(\mathbf{x}_t; \hat{\theta}_m)} \frac{1}{Z_m} w_{m-1}(t) e^{\hat{\alpha}_m} = \frac{1}{Z_m} (e^{-\hat{\alpha}_m} (1 - \hat{\varepsilon}_m) - e^{\hat{\alpha}_m} \hat{\varepsilon}_m) = 0$ (last step is by definition of $\hat{\alpha}_m$ being chosen to minimise the expression obtained in step 3 of Adaboost proof)

Theory

Concentration

- **Concentration:** General idea: to show how well things concentrate around the mean: $\mathbb{P}[|Y - m| > t] \leq \text{TailBound}(t)$
- Consider $Y_n = \frac{1}{n} \sum_{i=1}^n X_i$ where $\mathbb{E}[X_i] = \mu$ and $\text{Var}[X_i] = \sigma^2$ and X_i are i.i.d.
Law of large numbers: $\mathbb{P}[|Y_n - \mu| > \varepsilon] \rightarrow 0$ as $n \rightarrow \infty$ for any $\varepsilon > 0$
Central limit theorem: $\mathbb{P}[|Y_n - \mu| > \frac{\sigma}{\sqrt{n}}] \rightarrow 2\Phi\left(-\frac{\sigma}{\sigma}\right)$ as $n \rightarrow \infty$ where Φ is the standard normal c.d.f. (inaccurate for very small probabilities)
Large deviations theory (important): $\mathbb{P}[|Y_n - \mu| > \varepsilon] \leq e^{-n\psi(\varepsilon)}$ (this is true for any ε and any n , not just for large n)
- **Basic inequalities:**
Markov's ineq.: If Z is a nonnegative rand. var., then $\mathbb{P}[Z \geq t] \leq \frac{\mathbb{E}[Z]}{t}$
Markov's ineq. for functions: If Z is a rand. var. and ϕ is a nonnegative increasing function, then $\mathbb{P}[Z \geq t] \leq \mathbb{P}[\phi(Z) \geq \phi(t)] \leq \frac{\mathbb{E}[\phi(Z)]}{\phi(t)}$
Chebyshev's ineq.: If Z is a rand. var., then $\mathbb{P}[|Z - \mathbb{E}[Z]| \geq t] \leq \frac{\text{Var}[Z]}{t^2}$ (proof by letting $\phi(t) = t^2$ and replacing Z by $|Z - \mathbb{E}[Z]|$)
Chernoff bound: If Z is a rand. var. and $\lambda \geq 0$, then $\mathbb{P}[Z \geq t] \leq e^{-\lambda t} \mathbb{E}[e^{\lambda Z}]$ (proof by letting $\phi(t) = e^{\lambda t}$)
- **Sum of independent random variables:**
 $Z = X_1 + \dots + X_n$ and $Y_n = \frac{1}{n} Z$ where X_i are i.i.d.:
Chebyshev's ineq.: $\mathbb{P}[|Y_n - \mathbb{E}[Y_n]| \geq \varepsilon] \leq \frac{\text{Var}[X]}{n\varepsilon^2}$
Chernoff bound: $\mathbb{P}[Z \geq n\varepsilon] \leq \exp(-n \cdot \psi_X^*(\varepsilon))$ where $\psi_X^*(t) = \max_{\lambda} (\lambda t - \psi_X(\lambda))$ and $\psi_X(\lambda) = \text{Log } \mathbb{E}[e^{\lambda X}]$ (e.g. for Gaussian ($X \sim N(0, \sigma^2)$), $\psi_X(\lambda) = \frac{\lambda^2 \sigma^2}{2}$, so $\psi_X^*(t) = \frac{t^2}{2\sigma^2}$, so $\mathbb{P}[Z \geq n\varepsilon] \leq \exp\left(-\frac{n\varepsilon^2}{2\sigma^2}\right)$, so $\mathbb{P}[|Z| \geq n\varepsilon] \leq 2 \exp\left(-\frac{n\varepsilon^2}{2\sigma^2}\right)$) (note that for σ^2 -sub-Gaussian, we instead assume that $\psi_X(\lambda) \leq \frac{\lambda^2 \sigma^2}{2}$, and the same result holds)
- **Hoeffding's ineq.:** $Z = X_1 + \dots + X_n$ where X_i are i.i.d. and $X_i \in [a_i, b_i]$:
 $\mathbb{P}\left[\frac{1}{n} |Z - \mathbb{E}[Z]| \geq \varepsilon\right] \leq 2 \exp\left(\frac{-2n\varepsilon^2}{\sum_{i=1}^n (b_i - a_i)^2}\right)$

If $a_i = a$ and $b_i = b$ for all i then:

$$\mathbb{P}\left[\frac{1}{n} |Z - \mathbb{E}[Z]| \geq \varepsilon\right] \leq 2 \exp\left(\frac{-2n\varepsilon^2}{(b-a)^2}\right)$$

If $a_i = 0$ and $b_i = 1$ for all i then:

$$\mathbb{P}\left[\frac{1}{n} |Z - \mathbb{E}[Z]| \geq \varepsilon\right] \leq 2 \exp(-2n\varepsilon^2)$$

Statistical learning theory

- **Underfitting:** High training error (and hence high test error)
- **Overfitting:** Low training error but high test error
- **Setup:**
 - Data drawn from distribution $P_{\mathbf{X}Y}$ (unknown), i.e. $\mathcal{D} = \{(\mathbf{x}_t, y_t)\}_{t=1}^n$, where each $(\mathbf{x}_t, y_t) \sim P_{\mathbf{X}Y}$ and i.i.d.
 - Loss function $\ell(y, \hat{y})$ where $\hat{y} = f(x)$ (where f is our classifier) (ℓ could be any loss function)
 - Training error (empirical risk): $R_n(f) = \frac{1}{n} \sum_{t=1}^n \ell(y_t, f(\mathbf{x}_t))$
 - Test error (true risk): $R(f) = \mathbb{E}[\ell(y, f(\mathbf{x}))]$ where $(x, y) \sim P_{\mathbf{X}Y}$
 - $\underbrace{R(f)}_{\text{Test error}} = \underbrace{R_n f}_{\text{Training error}} + \underbrace{(R(f) - R_n f)}_{\text{Generalisation error}}$ (Generalisation error is large if we overfit)
 - Consider algorithm that outputs $f \in \mathcal{F}$ (if \mathcal{F} is too small then we underfit; if \mathcal{F} is too large then we overfit)
- **Task:** Let $f_{\text{erm}} = \arg \min_{f \in \mathcal{F}} R_n(f)$, want to find out if f_{erm} has small $R(f)$ too? Will show that $R(f_{\text{erm}}) \leq R(f^*) + \varepsilon$ with probability $\geq 1 - \delta$ provided that $n \geq \bar{n}(\varepsilon, \delta)$ (where $f^* = \arg \min_{f \in \mathcal{F}} R(f)$, and \bar{n} is known as the sample complexity) (called the “probably approximately correct” (PAC) guarantee) (the analysis is true even for the worst case $P_{\mathbf{X}Y}$, so for good distributions, it may be possible to use less data points)
 \mathcal{F} is PAC-learnable iff this is attainable for all $\varepsilon > 0, \delta > 0$ with $\bar{n} < \infty$ (regardless of $P_{\mathbf{X}Y}$) (when $|\mathcal{F}|$ is small enough, this is usually possible)
Realisable: $y_i = f(\mathbf{x}_i)$ for all i , for some $f \in \mathcal{F}$
Agnostic: Just compare with the (ideal) $f^* \in \mathcal{F}$ (we usually use this)
- **PAC guarantee for finite \mathcal{F} :** If $|\mathcal{F}| < \infty$ and $\ell(y, \hat{y}) \in [0, 1]$, then \mathcal{F} is PAC-learnable (i.e. $R(f_{\text{erm}}) \leq R(f^*) + \varepsilon$ with probability $\geq 1 - \delta$) with $\bar{n}_{\mathcal{F}}(\varepsilon, \delta) = \frac{2}{\varepsilon^2} \log \frac{2|\mathcal{F}|}{\delta}$ (equivalently, $R(f_{\text{erm}}) - R(f^*) \leq \varepsilon \leq \sqrt{\frac{2}{n} \log \frac{2|\mathcal{F}|}{\delta}}$)
- **Proof of PAC guarantee for finite \mathcal{F} :**
 1. Fix $f \in \mathcal{F}$. Let $z_i = \ell(y_i, f(\mathbf{x}_i)) \forall i$. Then since (\mathbf{x}_i, y_i) are i.i.d., z_i is also i.i.d.. Since $z_i \in [0, 1]$, and $\mathbb{E}[z_i] = \mathbb{E}[\ell(y_i, f(\mathbf{x}_i))] = R(f)$, we apply Hoeffding's ineq to get $\mathbb{P}\left[\left|\frac{1}{n} \sum_{i=1}^n z_i - R(f)\right| > \varepsilon_0\right] \leq 2 \exp(-2n\varepsilon_0^2)$. Then observe that $\frac{1}{n} \sum_{i=1}^n z_i = \frac{1}{n} \sum_{i=1}^n \ell(y_i, f(\mathbf{x}_i)) = R_n(f)$, so $\mathbb{P}[|R_n(f) - R(f)| > \varepsilon_0] \leq 2 \exp(-2n\varepsilon_0^2)$.
 2. Take the union bound: $\mathbb{P}\left[\bigcup_{f \in \mathcal{F}} \{|R_n(f) - R(f)| > \varepsilon_0\}\right] \leq 2|\mathcal{F}| \exp(-2n\varepsilon_0^2)$
 3. So we pick $\delta = 2|\mathcal{F}| \exp(-2n\varepsilon_0^2)$, so $n = \frac{1}{2\varepsilon_0^2} \log \frac{2|\mathcal{F}|}{\delta}$.
 4. Suppose $|R_n(f) - R(f)| \leq \varepsilon_0 \forall f \in \mathcal{F}$ (holds with probability $\geq 1 - \delta$), then: $R(f_{\text{erm}}) - R(f^*) = \underbrace{R(f_{\text{erm}}) - R_n(f_{\text{erm}})}_{\leq \varepsilon_0} + \underbrace{R_n(f_{\text{erm}}) - R_n(f^*)}_{\leq 0} + \underbrace{R_n(f^*) - R(f^*)}_{\leq \varepsilon_0} \leq 2\varepsilon_0$. Hence we let $\varepsilon := 2\varepsilon_0$, so $n = \frac{2}{\varepsilon^2} \log \frac{2|\mathcal{F}|}{\delta}$.
- **Shattering:** A set of points $\mathbf{x}_1, \dots, \mathbf{x}_k$ is said to be shattered by \mathcal{F} if $|\{(f(\mathbf{x}_1), \dots, f(\mathbf{x}_k)) \mid f \in \mathcal{F}\}| = 2^k$

- **VC Dimension:** $d_{\text{VC}} = d_{\text{VC}}(\mathcal{F}) =$ largest k such that $\exists \mathbf{x}_1, \dots, \mathbf{x}_k$ that \mathcal{F} shatters (i.e. largest k such that $\exists \mathbf{x}_1, \dots, \mathbf{x}_k$ for which all combinations of labels (i.e. $f(\mathbf{x}_1), \dots, f(\mathbf{x}_k)$) can be produced by classifiers in \mathcal{F}) (to determine the VC dimension, must show that k works but $k + 1$ does not work)

- **Corollary for Finite \mathcal{F} :** If \mathcal{F} is finite, $d_{\text{VC}} \leq \log_2 |\mathcal{F}|$

- **Sauer's lemma:** Effective $\#f$ in $\mathcal{F} \leq \sum_{i=0}^{d_{\text{VC}}} \binom{n}{i}$
Slightly weaker version: Effective $\#f$ in $\mathcal{F} \leq (n + 1)^{d_{\text{VC}}}$

- **PAC guarantee for infinite \mathcal{F} :** Assuming 0-1 loss, \mathcal{F} is PAC-learnable with $\bar{n}_{\mathcal{F}}(\varepsilon, \delta) = C \cdot \frac{d_{\text{VC}}(\mathcal{F}) + \log \frac{1}{\delta}}{\varepsilon^2}$

- **Converse guarantee:** If $d_{\text{VC}} = \infty$, then \mathcal{F} is not PAC-learnable (though there may be certain choices of $P_{\mathbf{X}Y}$ where learning might still work)

- **Examples**

- **Linear classifier (without offset):** $d_{\text{VC}} = d$ (the dimension of the space)

- **Linear classifier (with offset):** $d_{\text{VC}} = d + 1$

Unsupervised Learning

- **Training data:** $\mathcal{D} = \{\mathbf{x}_t\}_{t=1}^n$

- **K-means clustering:**

Goal: Partition \mathcal{D} into clusters $\mathcal{D}_1, \dots, \mathcal{D}_k$ such that the associated cluster centres $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k \in \mathbb{R}^d$ minimise $J(\{\mathcal{D}_j\}_{j=1}^k, \{\boldsymbol{\mu}_j\}_{j=1}^k) = \sum_{j=1}^k \sum_{\mathbf{x} \in \mathcal{D}_j} \|\mathbf{x} - \boldsymbol{\mu}_j\|^2$ (i.e. minimise the sum of squared distance from the respective cluster centres)

- **K-means algorithm:**

1. Minimise $J(\{\mathcal{D}_j\}_{j=1}^k, \{\boldsymbol{\mu}_j\}_{j=1}^k)$ w.r.t. $\{\mathcal{D}_j\}_{j=1}^k$ for fixed $\{\boldsymbol{\mu}_j\}_{j=1}^k$ (i.e. associate each data point with the cluster centre closest to it)

2. Minimise $J(\{\mathcal{D}_j\}_{j=1}^k, \{\boldsymbol{\mu}_j\}_{j=1}^k)$ w.r.t. $\{\boldsymbol{\mu}_j\}_{j=1}^k$ for fixed $\{\mathcal{D}_j\}_{j=1}^k$ (i.e. set each cluster centre to the mean of the data points associated with it, i.e. $\boldsymbol{\mu}_j = \frac{1}{|\mathcal{D}_j|} \sum_{\mathbf{x} \in \mathcal{D}_j} \mathbf{x}$)

3. Repeat until no change

Notes:

- initially, choose $\{\boldsymbol{\mu}_j\}_{j=1}^k$ at random

- K-means finds a local minimum for J

- **Distribution learning:**

Goal: Estimate a distribution $\hat{p}(\mathbf{x})$ that models the data well (either p.m.f. (discrete) or p.d.f. (continuous))

- assume $p(\mathbf{x})$ is determined by θ , i.e. $p(\mathbf{x}; \theta)$

- to find $\hat{\theta}$, we use maximum likelihood, i.e.

$$\hat{\theta} = \arg \max_{\theta} \prod_{t=1}^n p(\mathbf{x}_t; \theta) = \arg \max_{\theta} \sum_{t=1}^n \log p(\mathbf{x}_t; \theta)$$