

least-squares regression - Estimator

$$\text{OLS} \quad \hat{\theta} = (\Phi^T \Phi)^{-1} \Phi^T y$$

$$\text{Model: } y = \Phi \theta_* + \varepsilon$$

Regularization

$$\frac{\sigma^2 d}{n} \rightarrow \text{ridge}$$

dimension independent

TODAY : Convex surrogates

Estimation error for ERM

Data x_i, y_i $i = 1, \dots, n$

$\cap$ $\cap$
 X Y

loss function : $\ell(y, f(x))$
 $f: X \rightarrow Y$
 $\hookrightarrow$ prediction function

empirical risk : $\hat{R}(f) = \frac{1}{n} \sum_{i=1}^n \ell(y_i, f(x_i))$ Empirical

$R(f) = \mathbb{E}[\ell(y, f(x))]$ Expected

Goal : only $f(x) \in \mathbb{R}$.

OK for regression $\ell(y, f(x)) = (y - f(x))^2$

but if $Y = \{-1, 1\}$

Convex surrogates for classification and binary 0-1 loss

going from $g(x) \in \mathbb{R}$ to $f(x) \in \{-1, 1\}$

consider $f(x) = \text{sign}(g(x)) = \begin{cases} 1 & \text{if } g(x) > 0 \\ -1 & \text{if } g(x) < 0 \end{cases}$

(= random if $g(x) = 0$)

$R(g)$ Risk of g vs the risk $R(f)$
(= $R_{01}(g)$) $R(f) = P(f(x) \neq g) = \mathbb{E}[1_{f(x) \neq g}]$
 $= P[f(x) \neq g, g(x) > 0] + P[f(x) \neq g, g(x) < 0]$
 $+ P(f(x) \neq g, g(x) = 0)$

$R(g)$ Risk of g vs the risk $R(f)$


$$R(f) = P(f(x) \neq g) = E[1_{f(x) \neq g}]$$

$$= \underbrace{P[f(x) \neq g, g(x) > 0] + P[f(x) \neq g, g(x) < 0]}_{P(yg(x) < 0)}$$


$$+ P(f(x) \neq g, g(x) = 0)$$

$$R(g) = E \left[1_{yg(x) < 0} + \frac{1}{2} 1_{yg(x) = 0} \right]$$

$$= E[\Phi_{01}(yg(x))]$$



$$\phi_{01}(u) = \begin{cases} 1 & \text{if } u < 0 \\ \frac{1}{2} & \text{if } u = 0 \\ 0 & \text{if } u > 0 \end{cases}$$



square loss : $\phi(u) = (u-1)^2$


$$\phi(yg(x)) = (yg(x) - 1)^2 = (g(x) - y)^2$$

logistic loss : $\phi(u) = \log(1 + e^{-u}) \Rightarrow \phi(yg(x)) = \log(1 + e^{-yg(x)})$

probabilistic model : $P(y=1 | x) = \sigma(g(x))$
 $\hookrightarrow$ sigmoid

$$P(y | x) = \sigma(yg(x))$$

maximum likelihood
 $\max \log P(y|x)$




$$\sigma(u) = \frac{1}{1 + e^{-u}}$$

$$\sigma(-u) = 1 - \sigma(u)$$

Hinge loss

α loss



link with support vector machine
(SVM)

0-1 loss binary classification

$$R(f) = \mathbb{E} \phi_{01}(yg(x)) = \mathbb{E} \phi(yg(x))$$

values in $\{-1, 1\}$ values in $\mathbb{R}$

0-1 loss binary classification

$$R(g) = \mathbb{E} \phi_{01}(yg(x)) = \mathbb{E} \phi(yg(x))$$

values in $\{-1, 1\}$
values in $\mathbb{R}$

① Classification calibration = if $g^*(x)$ which minimizes the expected ϕ -risk $\mathbb{E}(\phi(yg(x)))$

Assumption:

$$\phi \text{ convex} / \phi'(c) < \infty$$

if square loss $g^*(x) = \mathbb{E}(y|x)$
 $\text{sign}(\mathbb{E}(y|x))$ is the Bayes predictor

then $\text{sign}(g^*(x))$ is a Bayes predictor

② There is a calibration function

$$R_\alpha(g) - R_\alpha^* \leq H \left[R_\phi(g) - R_\phi^* \right]$$

($H(n) = n$ hinge loss), for square / log

Risk decomposition = $\mathcal{F}$ family of functions: $X \rightarrow \mathbb{R}$

ERM $\hat{f} \in \arg \min_{f \in \mathcal{F}} \hat{R}(f) = \frac{1}{n} \sum_{i=1}^n \ell(y_i, f(x_i))$

$\hookrightarrow$ square
 $\hookrightarrow$ logistic
 $\hookrightarrow$ hinge

in most cases: $\mathcal{F} = \{f_\theta, \theta \in \Theta\}$

$$R(\hat{f}) - R^* = \underbrace{R(\hat{f}) - \inf_{f' \in \mathcal{F}} R(f')}_{\text{estimation error}} + \underbrace{\inf_{f' \in \mathcal{F}} R(f') - R^*}_{\text{approximation error}}$$

$$|c^h|_5 \rightarrow |c^9|_5$$

Approximation error = $\mathcal{E} = \{f_{\theta}; \theta \in \Theta \subset \mathbb{R}^d\}$
 ex: $f_{\theta}(x) = \theta^T \varphi(x)$

$$\inf_{\theta \in \Theta} R(f_{\theta}) - R^* = \underbrace{\inf_{\theta \in \Theta} R(f_{\theta}) - \inf_{\theta \in \mathbb{R}^d} R(f_{\theta})}_{\text{effect of restricting } \theta \in \Theta} + \underbrace{\inf_{\theta \in \mathbb{R}^d} R(f_{\theta}) - R^*}_{\text{incompressible \& / max error}}$$

effect of
restricting $\theta \in \Theta$

incompressible
& / max error


ex: $\Theta = \{\theta, \|\theta\|_2 \leq D\}$

Example: linear model with G-lipschitz-continuous loss
 Model $f_{\theta}(x) = \theta^T \varphi(x)$, with optimal $\theta_* = \theta^*$

Bounding $\inf_{\|\theta\|_2 \leq D} R(f_{\theta}) - R(f_{\theta_*}) = \inf_{\|\theta\|_2 \leq D} \mathbb{E}[\ell(y, f_{\theta}(x)) - \ell(y, f_{\theta_*}(x))]$

Cauchy
Schwarz

$$\leq \inf_{\|\theta\|_2 \leq D} \mathbb{E}[\|\varphi(x)\|_2] \cdot \|\theta - \theta_*\|_2$$



$$= G \mathbb{E}[\|\varphi(x)\|_2] (\|\theta_*\|_2 - D)_+ + G \mathbb{E}[\|\varphi(x)\|_2] (D - \|\theta_*\|_2)_+$$

$$R(\hat{g}) - \min_{f \in \mathcal{F}} R(f) = R(\hat{g}) - R(g_F) \quad \text{, minimizer of } R \text{ on } \mathcal{F}.$$

$$= \underbrace{R(\hat{g}) - \hat{R}(\hat{g})}_1 + \underbrace{\hat{R}(\hat{g}) - \hat{R}(g_F)}_{\leq \varepsilon \text{ optimization error}} + \underbrace{\hat{R}(g_F) - R(g_F)}_{\downarrow}$$

$$\frac{1}{n} \sum_{i=1}^n \ell(g_i, \hat{g}(x_i)) \rightarrow \mathbb{E} \ell \quad \begin{matrix} \text{iid} \\ \text{ } \end{matrix} \quad \stackrel{?}{=} \underbrace{\sum_n \ell(g_i, \hat{g}(x_i)) - R(\hat{g})}_1$$

$$\stackrel{?}{=} \sum_n \ell(g_i, g_F(x_i)) - \mathbb{E} \ell$$

$$\leq \min_{f \in \mathcal{F}} R(f) - \hat{R}(\hat{g}) + \varepsilon + \sup_b \hat{R}(b) - R(\hat{g}),$$

uniform deviation

$$\leq 2 \sup_b |R(b) - \hat{R}(b)| + \varepsilon$$

goal: control uniform deviations: $\sup_{f \in \mathcal{F}} R(f) - \hat{R}(f)$

Two types of guarantees: (1) $\mathbb{E} \left[\sup_{f \in \mathcal{F}} R(f) - \hat{R}(f) \right] \leq \dots$

(2) with probability $1 - \delta$, $\sup_{f \in \mathcal{F}} R(f) - \hat{R}(f) \leq O(\delta)$.

Review of concentration inequality

(1) Markov inequality
if $Z \geq 0$

$$P(Z > t) \leq$$

$\frac{\mathbb{E}[Z]}{t} = \delta$

Solve

$$\Rightarrow \text{with probability at least } 1 - \delta, \quad Z \leq \frac{\mathbb{E}[Z]}{\delta}$$

(2) Hoeffding's inequality. if $Z_1, \dots, Z_n \in [0, 1]$ and independent
then $P\left(\frac{1}{n} \sum_{i=1}^n Z_i - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_i] > t\right) \leq e^{-2nt^2} = \delta$

$$\text{with proba at least } 1 - \delta, \quad \frac{1}{n} \sum_{i=1}^n Z_i - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_i] \leq \sqrt{\log \frac{1}{\delta} \frac{1}{2n}}$$

Concentration inequality for $R(\ell) - \bar{R}(\ell)$ for fixed ℓ
 if loss function $\ell \in [c, \ell_\infty)$, $Z_i = -\ell(y_i, \ell(x_i))$

Hoeffding

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} Z_i - \frac{1}{n} \sum_{i=1}^n Z_i \leq \sqrt{\frac{\ell_\infty}{2n}} \sqrt{\ln \frac{1}{\delta}}$$

$$R(\ell) - \bar{R}(\ell)$$

~~" $\Rightarrow$ "~~ with probability $1 - \delta$ $\max_{\ell \in F} R(\ell) - \bar{R}(\ell) \leq \dots$

$\Rightarrow$ see at Section 4.4.1 on MacDiarmid inequality

with probability at least $1 - \delta$, $Z \leq \mathbb{E} Z + \frac{\ell_\infty}{\sqrt{2n}} \sqrt{\ln \frac{1}{\delta}}$

How to bound uniform deviations?

- finite class of models
- general case

① Finite class of models

$$\mathbb{P}\left(\sup_{f \in F} R(f) - \hat{R}(f) \geq t\right) = \mathbb{P}\left(\bigcup_{f \in F} \{R(f) - \hat{R}(f) \geq t\}\right)$$

$$\leq \sum_{f \in F} \mathbb{P}(R(f) - \hat{R}(f) \geq t)$$

union
bound

$$\leq \sum_{f \in F} e^{-2nt^2/\ell^2} = |F| e^{-2nt^2/\ell^2}$$

" δ

$$t = \frac{\ell}{\sqrt{2n}} \left(\sqrt{\lg \frac{1}{\delta}} + \lg |F| \right)$$

$$\text{with prob} \geq 1 - \delta, \sup_{f \in F} R(f) - \hat{R}(f) \leq \rightarrow$$

Rademacher complexity

$\mathcal{F}$ family of functions: $\mathcal{X} \rightarrow \mathbb{R}$
 $\mathcal{H} = \{(\underline{z}, y) \mapsto \ell(y, f(\underline{z})), f \in \mathcal{F}\}$

family ^{$\mathcal{F}$} of functions from
 $\mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}$

$$\sup_f R(f) - \widehat{R}(f) = \sup_{h \in \mathcal{H}} \mathbb{E}[h(\underline{z})] - \frac{1}{n} \sum_{i=1}^n h(\underline{z}_i)$$

$\ell(y_i, f(\underline{z}_i))$

goal: bound the expectation of $\uparrow$

$$R_n(\mathcal{H}) = \mathbb{E}_{\text{data}, \varepsilon} \left[\sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \varepsilon_i h(\underline{z}_i) \right]$$

Rademacher
complexity of
 $\mathcal{H}$

$\varepsilon \in \{-1, 1\}^n$ with independent
uniform values

$\mathbb{P}(\varepsilon_i = \pm 1) = \frac{1}{2}$ Rademacher r.v

Symmetrization: $z_1, \dots, z_n$ iid. = data D
 independent copy $z'_1, \dots, z'_n$ iid $\Rightarrow$ data D'

$$\mathbb{E} \left[\sup_{h \in H} \underbrace{\mathbb{E} \left[\frac{1}{n} \sum_i h(z_i) \right]}_{\text{data } D} - \frac{1}{n} \sum_i h(z_i) \right] = \mathbb{E} \left[\sup_{h \in H} \underbrace{\frac{1}{n} \sum_{i=1}^n \mathbb{E} [h(z'_i) | D]}_{\text{data } D'} - \frac{1}{n} \sum_i h(z_i) \right]$$

because D & D'
are indep.

$$= \mathbb{E} \left[\sup_{h \in H} \frac{1}{n} \sum_i \mathbb{E} [h(z'_i) - h(z_i) | D] \right]$$

because

$$\mathbb{E} [h(z_i) | D] = h(z_i)$$

$$= \mathbb{E} \left[\sup_h \mathbb{E} \left(\frac{1}{n} \sum_i (h(z'_i) - h(z_i)) \mid D \right) \right]$$

$$\leq \mathbb{E} \left[\mathbb{E} \left(\sup_h \frac{1}{n} \sum_i (h(z'_i) - h(z_i)) \mid D \right) \right]$$

$$= \mathbb{E}_{\text{data } D} \left[\sup_h \frac{1}{n} \sum_i \varepsilon_i (h(z'_i) - h(z_i)) \right]$$

because

$$h(z'_i) - h(z_i)$$

"symmetric"

$$\leq \mathbb{E}_{\text{data } D, \varepsilon} \left[\sup_h \frac{1}{n} \sum_i \varepsilon_i h(z'_i) + \sup_h \frac{1}{n} \sum_i (-\varepsilon_i) h(z_i) \right]$$

$$= 2 P_n(H) = 2 \mathbb{E}_{\text{data } D, \varepsilon} \left(\sup_h \frac{1}{n} \sum_i \varepsilon_i h(z_i) \right)$$

$$\sup_h \mathbb{E} \leq \mathbb{E} \sup_h$$

$$\mathbb{E}(\text{estimation error}) \leq 4 R_n(H) = 4 \mathbb{E}_{\text{data}, \varepsilon} \left[\max_h \frac{1}{n} \sum_i \varepsilon_i h(x_i) \right]$$

$$= 4 \mathbb{E}_{\text{data}, \varepsilon} \left[\max_{f \in F} \frac{1}{n} \sum_i \varepsilon_i \ell(x_i, f(x_i)) \right]$$

Contracta principle:

if the loss ℓ is G -Lipschitz continuous

then

$$R_n(H) \leq G \cdot R_n(F)$$

$$= G \mathbb{E}_{\text{data}, \varepsilon} \left[\max_{f \in F} \frac{1}{n} \sum_i \varepsilon_i \ell(x_i, f(x_i)) \right]$$

Example: $\ell(\theta, x) = \theta^T \varphi(x)$, $\|\theta\|_2 \leq D$

$$R_n(F) = \mathbb{E}_{\text{data}, \varepsilon} \max_{\|\theta\|_2 \leq D} \left(\frac{1}{n} \sum_i \varepsilon_i \varphi(x_i) \right)^T \theta = \mathbb{E}_{\text{data}, \varepsilon} D \left\| \frac{1}{n} \sum_i \varepsilon_i \varphi(x_i) \right\|_2$$

Cauchy-Schwarz linear in θ

$$\mathbb{E} \|\theta\|_2 \leq \sqrt{\mathbb{E} \|\theta\|_2^2}$$

Tensor's inequality

$$\leq D \sqrt{\mathbb{E}_{\text{data}, \varepsilon} \left\| \frac{1}{n} \sum_i \varepsilon_i \varphi(x_i) \right\|_2^2} = D \sqrt{\frac{1}{n^2} \sum_i \mathbb{E} \varepsilon_i^2 \|\varphi(x_i)\|_2^2}$$

$$= \frac{D}{\sqrt{n}} \sqrt{\mathbb{E} \|\varphi(x)\|_2^2}$$

Summary for linear models

Empirical risk
minimizer

$$f_{\Theta}(x) = \varphi(x)^T \Theta$$

$$\Theta = \{ \Theta, \|\Theta\|_2 \leq D \}$$

$$\Theta_{\delta} = \arg \min_{\Theta} \mathcal{Q}(\Theta)$$

$$\|\varphi(x)\|_2 \leq R$$

$$\text{Approximation error} \leq GR(\|\Theta_{\delta}\|_2 - D)_+ +$$

$$\text{Estimation error} \leq 4G \frac{RD}{\sqrt{n}}$$

$$D = \|\Theta_{\delta}\|_2$$

$$\Rightarrow \frac{4GR\|\Theta_{\delta}\|_2}{\sqrt{n}}$$

MAGIC 3 : t iterations
of GD

$$\text{opt. error} : GRD/\sqrt{t}$$

excess risk compared
to $\mathcal{R}(\Theta_{\delta})$