

today: linear regression
logistic "

linear regression : derive/motivate with conditional approach to regression ($Y \in \mathbb{R}$)

$$p(y|x; w) = N(y | \underbrace{\langle w, x \rangle}_{\text{parameter}}, \sigma^2)$$

$$x \in \mathbb{R}^d$$

$$w \in \mathbb{R}^d$$

equivalently: $Y_i = w^T X_i + \epsilon_i$ where $\epsilon_i | X_i \sim \text{iid } N(0, \sigma^2)$


[aside: we use "offset" relation for x]

$$\text{i.e. } x = \begin{pmatrix} \tilde{x} \\ 1 \end{pmatrix} \quad \tilde{x} \in \mathbb{R}^{d-1}$$

"constant feature"

$$\text{thus } \langle w, x \rangle = \langle w_{1:d-1}, \tilde{x} \rangle + w_d$$

"bias" / "offset" (often b)



* dataset $(x_i, y_i)_{i=1}^n$

$X_i \sim \text{whenever (don't care)}$

$$Y_i | X_i \sim \text{iid } N(w^T X_i, \sigma^2)$$

conditional likelihood

$$p(y_{1:n} | x_{1:n}) \stackrel{\text{iid}}{=} \prod_{i=1}^n p(y_i | x_i)$$

$$\log(\quad) = \sum_{i=1}^n \left[-\frac{(y_i - w^T x_i)^2}{2\sigma^2} - \frac{1}{2} \log(2\pi\sigma^2) \right]$$

this is not a concave fct. of σ^2

$\frac{a}{x} + b \log x$

$$\frac{\partial}{\partial \sigma^2}(\quad) = 0 \Rightarrow \sum_{i=1}^n \left[-\frac{(y_i - w^T x_i)^2}{2} \left(-\frac{1}{(\sigma^2)^2} \right) - \frac{1}{2} \left(\frac{1}{\sigma^2} \right) \right] = 0$$

$$\Rightarrow \hat{\sigma}_{MLE}^2 = \frac{1}{n} \sum_{i=1}^n (y_i - w^T x_i)^2$$

obj $\rightarrow -\infty$ as $\sigma \rightarrow 0$
or $\sigma \rightarrow \infty$

So conclude that this is correct global max. in σ^2 for w fixed

(also see notes below)

"design matrix" $X \triangleq \begin{pmatrix} -x_1^T \\ -x_2^T \\ \vdots \\ -x_n^T \end{pmatrix}$ $n \times d$ matrix

vector $y \triangleq \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}$


$$Xw = \begin{pmatrix} x_1^T w \\ x_2^T w \\ \vdots \\ x_n^T w \end{pmatrix} \in \mathbb{R}^{n \times 1} \quad \sum_{i=1}^n (y_i - w^T x_i)^2 = \|y - Xw\|_2^2$$

$$Xw = \begin{pmatrix} x_1^T w \\ x_2^T w \\ \vdots \\ x_n^T w \end{pmatrix} \in \mathbb{R}^{n \times 1} \quad \sum_{i=1}^n (y_i - w^T x_i)^2 = \|y - Xw\|_2^2$$

can rewrite $-\log p(y_{1:n}|X) = \frac{\|y - Xw\|_2^2}{2\sigma^2} + \text{const.}$

design matrix

MCL $\rightarrow \min_w \|y - Xw\|_2^2 \Leftrightarrow$ projecting y on the column space of design matrix X



$$Xw = \sum_{j=1}^d x_{:j} w_j$$

\uparrow
jth column of X

$\hat{w}_{MLE} = \arg\min_{w \in \mathbb{R}^d} \|y - Xw\|_2^2$ "least square"

algebra: want $\nabla_w \xrightarrow{\text{set to}} 0$ want

$$\frac{\partial}{\partial w} [(y - Xw)^T (y - Xw)] \stackrel{!}{=} 0$$

$$\frac{\partial}{\partial w} [\|y\|^2 - 2y^T Xw + w^T X^T X w] = 0$$

$$0 - 2X^T y + 2X^T X w = 0$$

$$\Rightarrow (X^T X) w^* = X^T y$$

vector \downarrow

$$\nabla_w (w^T A w) = (A + A^T)w$$

(convex fct. of w
 \Rightarrow stat. pt. is global min)

"normal equation"

a) if $X^T X$ is invertible, then have unique solution

$$\hat{w}_{MLE} = (X^T X)^{-1} X^T y$$

X is $n \times d \Rightarrow \text{rank}(X) \leq \min\{n, d\}$
 $\text{rank}(X^T X) = \text{rank}(X) \leq \min\{n, d\}$

\Rightarrow if $n < d$ (i.e. high dimension or low data regime) then $X^T X$ is not invertible

$X^T X$ is invertible $\Rightarrow n \geq d$

prediction on training set

$$\hat{y} = Xw^* = X \underbrace{(X^T X)^{-1} X^T y}_{\hat{w}_{MLE}}$$

projection on column space of X

(recall geometric perspective)

b) if $X^T X$ is not invertible \rightarrow there is no unique sol'n

any \hat{w} s.t. $(X^T X)\hat{w} = X^T y$ is a MCL solution

could choose $\hat{w} = \arg\min \|w\|_2 = X^+ y$ Moore-Penrose pseudo-inverse

only w s.t. $Xw = y$ is a math solution
 could choose $\hat{w} = \arg \min_w \|w\|_2 = X^+ y$ Moore-Penrose pseudo-inverse

SVD
 $X = U \Sigma V^T$
 $n \times d$ $n \times n$ $n \times d$ $d \times d$

s.t. $(X^T X) w = X^T y$
 $U U^T = I_n$ $V V^T = I_d$

$X^+ = (X^T X)^+ X^T$ when X is full-rank

$\Sigma = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_d \\ 0 & \end{pmatrix}$
 $n \times d$

$X^+ = V \Sigma^+ U^T$

$\Sigma^+ = \begin{pmatrix} \sigma_1^+ & 0 & \\ 0 & \sigma_d^+ & \\ & & 0 \end{pmatrix}$

$\sigma_i^+ = \begin{cases} 1/\sigma_i & \text{if } \sigma_i \neq 0 \\ 0 & \text{otherwise} \end{cases}$

problem: pseudo-inverse is not numerically stable

instead it is better to regularize to get similar effect

$\hat{w}_{\text{MAP}}(\lambda) \xrightarrow{\lambda \rightarrow 0} \hat{w}_{\text{pseudo-inverse}}$

16h04

regularization: (can be motivated from MAP point of view)

suppose we put a prior $p(w) = N(w | 0, \frac{1}{\lambda} I)$
 λ is "precision parameter"
 I is $d \times d$ identity matrix

log posterior: $\log p(w | \text{data}) = \log p(y, w | X, w) + \log p(w) + \text{cst.}$
 $= -\frac{1}{2\sigma^2} \|y - Xw\|^2 - \frac{\lambda}{2\sigma^2} \|w\|^2 + \text{cst.}$

MAP here

$\hat{w}_{\text{MAP}} = \arg \min_{w \in \mathbb{R}^d} \frac{1}{2} \|y - Xw\|^2 + \frac{\lambda}{2} \|w\|^2$

"ridge regression"

same as "regularized ERM"

(empirical risk minimization) with squared loss

$\ell(y_i, w^T x_i) = \frac{1}{2} (y_i - w^T x_i)^2$

$\frac{1}{n} \sum_{i=1}^n \ell(y_i, h_w(x_i)) + \frac{\lambda}{2n} \|w\|^2$

Squared loss
 empirical error

regularization

this obj.

is strongly convex in w
 \Rightarrow a unique sol'n

$f(\cdot)$ is λ -strongly convex $\Leftrightarrow f(\cdot) - \frac{\lambda \|\cdot\|^2}{2}$ is convex in (\cdot)

$\nabla w = 0 \Rightarrow$

$(X^T X + \lambda I) w = X^T y$
 always invertible $\lambda > 0$

$g(x) \triangleq f(x) - \frac{\lambda \|x\|^2}{2} \forall x$

\hat{w}_{MAP} or
 ridge regression
 $\hat{w}_{\text{MAP}} = (X^T X + \lambda I)^{-1} X^T y$

no problem for $d > n$

good practice: to either standardize features i.e. make each feature with zero mean & unit empirical variance

or
normalize \leftarrow make x_i unit norm $\|x_i\|_2 = 1$
or
scale features to $[0,1]$ or $[-1,1]$

logistic regression

Setup: binary classification $\mathcal{Y} = \{0,1\}$ $X \in \mathbb{R}^d$

generative model motivation:

suppose only assumption is there exists a pdf (densities) in \mathbb{R}^d for each class conditionals:

$$p(x|Y=1) \text{ \& } p(x|Y=0)$$

$$\begin{aligned} p(Y=1|X=x) &= \frac{p(Y=1, X=x)}{p(Y=1, X=x) + p(Y=0, X=x)} \cdot p(X=x) \\ &= \frac{1}{1 + \frac{p(Y=0, X=x)}{p(Y=1, X=x)}} = \frac{1}{1 + \exp(-f(x))} \end{aligned}$$


where $f(x) \triangleq \log \frac{p(X=x|Y=1)}{p(X=x|Y=0)} + \log \frac{p(Y=1)}{p(Y=0)}$

"log odds"
class-conditional ratio
prior odds ratio

in general

$$p(Y=1|X=x) = \sigma(f(x)) \quad \text{where } \sigma(z) \triangleq \frac{1}{1 + \exp(-z)}$$

"sigmoid function"



Some properties of $\sigma(z)$: $\sigma(-z) = 1 - \sigma(z)$ $[\sigma(z) + \sigma(-z) = 1]$

$$\frac{d}{dz} \sigma(z) = \sigma(z) \sigma(-z) = \sigma(z) (1 - \sigma(z))$$

(*) to motivate linear logistic regression,

consider class conditionals in the exponential family

$$p(x|m) \triangleq \underbrace{h(x)}_{\text{"canonical parameter"}} \exp \underbrace{m^T T(x)}_{\text{"sufficient statistics"}} - A(m)$$

scalar prod. \rightarrow log partition fct.
normalize

"canonical parameter"

these specify the "flat" exponential family that we are considering

Gaussian:

$$\log p(x|\mu, \sigma^2) = -\frac{1}{2} \log 2\pi\sigma^2 - \frac{(x-\mu)^2}{2\sigma^2}$$

$$\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

$-\frac{1}{2} \log 2\pi\sigma^2 \sim -\frac{1}{2} \log \sigma^2$

$$\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \quad \frac{1}{2} \log \frac{1}{2\sigma^2} - \left[\frac{x^2}{2\sigma^2} - x\frac{\mu}{\sigma^2} + \frac{\mu^2}{2\sigma^2}\right]$$

$$\text{let } T(x) = \begin{bmatrix} -x^2/2 \\ x \\ 1/\sigma^2 \end{bmatrix}$$

$$n(\mu, \sigma^2) = \begin{bmatrix} 1/\sigma^2 \\ \mu/\sigma^2 \end{bmatrix}$$

$$A(n) = \frac{1}{2} \log(2\pi\sigma^2) + \frac{\mu^2}{2\sigma^2}$$

$$p(x|y=1) = p(x|\eta_1)$$

$$p(x|y=0) = p(x|\eta_0)$$

$$\begin{aligned} \text{log odds} \quad f(x) &= \log \frac{p(x|\eta_1)}{p(x|\eta_0)} + \log \frac{p(y=1)}{p(y=0)} \\ &= (\eta_1 - \eta_0)^T T(x) + A(\eta_0) - A(\eta_1) + \log \frac{\pi}{1-\pi} \\ &\triangleq w^T \ell(x) \end{aligned}$$

$$\text{where } w = \begin{pmatrix} \eta_1 - \eta_0 \\ A(\eta_0) - A(\eta_1) + \log \frac{\pi}{1-\pi} \end{pmatrix} \quad \ell(x) = \begin{pmatrix} T(x) \\ 1 \end{pmatrix}$$

get logistic regression model

$$p(y=1|x=x) = \sigma(w^T \ell(x))$$

(feature map)

decision boundary

$$\mathbb{1}\{w^T \ell(x) > 0\}$$

exercise to reader: (hw 2) try argument above with $p(x|y) = N(x|\mu_y, \Sigma_y)$

$$\bullet \text{ if } \Sigma_0 = \Sigma_1, \text{ then } \ell(x) = \begin{pmatrix} x \\ 1 \end{pmatrix}$$

$$\bullet \text{ otherwise } \ell(x) = \begin{pmatrix} -xx^T \\ x \\ 1 \end{pmatrix}$$

quadratic

quadratic decision boundary

• note about sigma^2 being a global max

(aside: showing that the sigma^2 above is the **global max** is subtle because the objective is not concave in sigma^2. I give more info here for your curiosity, but it is not required for the assignment.)

- Formally, to find a global max of a *differentiable objective*, you need to check all **stationary points** (zero gradient points), **as well as the values at the boundary of the domain**.

Thus here, you would need to show that the objective cannot take higher value anywhere at the boundary of the domain (which is the case here (exercise!), as the objective goes to -infinity at the boundary), so you are done (this is the only possible global optimum -- a maximum here, as it should be, given that there are no other stationary points and all values are lower at the boundary, but one could also explicitly check the Hessian to see that it is strictly negative definite at the stationary point, i.e. it looks like a local maximum).

Note that we will see later in the class that the Gaussian is in the exponential family, with a log-concave likelihood in the right ("natural") parameterization, and thus using the invariance principle of the MLE, we could also easily deduce the MLE in the "moment" parameterization which is the usual (μ, σ^2) one, without having to worry about local optima...

- for a cute counter-example illustrating that a differentiable function could have only one stationary point which is a local min but *not a global min* (and thus why one needs to look at the values at the boundary), see:

- https://en.wikipedia.org/wiki/Maxima_and_minima#Functions_of_more_than_one_variable

- i.e.

$$f(x, y) = x^2 + y^2(1 - x)^3, \quad x, y \in \mathbb{R},$$

shows. Its only critical point is at $(0,0)$, which is a local minimum with $f(0,0) = 0$. However, it cannot be a global one, because $f(2,3) = -5$.

(see picture of function [here](#))

(and note that the "[Mountain pass theorem](#)" which basically says that if you have a strict local optimum with another point somewhere with the same value, then there must be a saddle point somewhere (a "mountain pass") i.e. another stationary point, **does not hold for this counter-example** as one of the required regularity conditions, the "Palais-Smale compactness condition" fails. Here, the saddle point (which should intuitively exist) "happens at infinity", which is why it only has one stationary point despite $(0,0)$ not being a global minimum)

- the moral of the story: intuitions for multivariate optimization are often misleading! (this counter-example would not work in 1d because of [Rolle's theorem](#))