

Logistic Regression

Classification

Aarti Singh & Geoff Gordon

Machine Learning 10-701

Feb 22, 2021



MACHINE LEARNING DEPARTMENT

Carnegie Mellon.
School of Computer Science

Logistic Regression

Not really regression

Assumes the following functional form for $P(Y|X)$:


$$P(Y = 1|X) = \frac{1}{1 + \exp(-w_0 - \sum_i w_i X_i)} = \sigma(w_0 + \sum_i w_i X_i)$$

$X = \begin{bmatrix} X_1 \\ \vdots \\ X_p \end{bmatrix}$

Logistic function applied to a linear function of the data

Logistic
function

(or Sigmoid), $\sigma(z) = \frac{1}{1 + \exp(-z)}$



Features can be discrete or continuous!

Logistic Regression is a Linear Classifier!

Assumes the following functional form for $P(Y|X)$:


$$\rightarrow P(Y = 1|X) = \frac{1}{1 + \exp(-w_0 - \sum_i w_i X_i)} \quad \Rightarrow P(Y=0|X) = \frac{e^{-c}}{1+e^{-c}}$$

Decision boundary:

$$\Rightarrow \frac{P(Y = 1|X)}{P(Y = 0|X)} = \exp(w_0 + \sum_i w_i X_i) \stackrel{\textcircled{1}}{\geq} \textcircled{0} 1$$

$$\Rightarrow w_0 + \sum_i w_i X_i \stackrel{1}{\geq} \underset{0}{} 0$$

(Linear Decision Boundary)



Training Logistic Regression

How to learn the parameters $w_0, w_1, \dots, w_d$? (d features)

Training Data $\{(X^{(j)}, Y^{(j)})\}_{j=1}^n$ $X^{(j)} = (X_1^{(j)}, \dots, X_d^{(j)})$

Maximum Likelihood Estimates

$$\hat{\mathbf{w}}_{MLE} = \arg \max_{\mathbf{w}} \prod_{j=1}^n P(\underbrace{X^{(j)}, Y^{(j)}}_{\mathcal{D}} \mid \mathbf{w})$$

$w_2 (w_0 \dots w_d)$

But there is a problem ...

Don't have a model for $P(X)$ or $P(X|Y)$ – only for $P(Y|X)$

Training Logistic Regression

How to learn the parameters $w_0, w_1, \dots, w_d$? (d features)

Training Data $\{(X^{(j)}, Y^{(j)})\}_{j=1}^n$ $X^{(j)} = (X_1^{(j)}, \dots, X_d^{(j)})$

Maximum (Conditional) Likelihood Estimates

$$\hat{w}_{MCLE} = \arg \max_w \prod_{j=1}^n P(\underline{Y^{(j)}} \mid \underline{X^{(j)}}; w)$$

Discriminative philosophy – Don't waste effort learning $P(X)$, focus on $P(Y|X)$ – that's all that matters for classification!

$$\arg \max_y P(Y|X) = \arg \max_y \underbrace{P(X|Y)}_{\sim \mathcal{N}(\cdot)} \underbrace{P(Y)}_{\leftarrow}$$

Expressing Conditional log Likelihood

$$P(Y = 1|X, w) = \frac{1}{1 + \exp(-w_0 - \sum_i w_i X_i)} \quad \leftarrow$$

$$\frac{e^{-l}}{1+e^{-l}}$$

$$P(Y = 0|X, w) = \frac{1}{1 + \exp(w_0 + \sum_i w_i X_i)} \quad \leftarrow$$

$$l(\mathbf{w}) \equiv \log \prod_j P(y^j | \mathbf{x}^j, \mathbf{w})$$

Expressing Conditional log Likelihood

$$P(Y = 1|X, w) = \frac{1}{1 + \exp(-w_0 - \sum_i w_i X_i)}$$

$$P(Y = 0|X, w) = \frac{1}{1 + \exp(w_0 + \sum_i w_i X_i)}$$

$$l(\mathbf{w}) \equiv \ln \prod_j P(\underline{y}^j | \mathbf{x}^j, \mathbf{w})$$

$$= \sum_j \left[y^j (w_0 + \sum_i w_i x_i^j) - \ln(1 + \exp(w_0 + \sum_i w_i x_i^j)) \right]$$

Good news: $l(\mathbf{w})$ is concave function of $\mathbf{w}$!

Bad news: no closed-form solution to maximize $l(\mathbf{w})$

HW2!

Good news: can use iterative optimization methods (gradient ascent)

That's M(C)LE. How about M(C)AP?

$$p(\mathbf{w} \mid Y, \mathbf{X}) \propto \underbrace{P(Y \mid \mathbf{X}, \mathbf{w})}_{\text{likelihood}} \underbrace{p(\mathbf{w})}_{\text{prior}}$$

- Define priors on $\mathbf{w}$
 - Common assumption: Normal distribution, zero mean, identity covariance
 - “Pushes” parameters towards zero


$$p(\mathbf{w}) = \prod_i \frac{1}{\kappa \sqrt{2\pi}} e^{-\frac{w_i^2}{2\kappa^2}}$$

Zero-mean Gaussian prior
 $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \kappa^2 \mathbf{I})$

Logistic
function

(or Sigmoid):
$$\frac{1}{1 + \exp(-z)}$$

$$z = \sum_i w_i x_i + w_0$$




➤ What happens if we scale z by a large constant?

$$z = w_0 + \sum_i w_i x_i$$

**Logistic
function**

(or Sigmoid):

$$\frac{1}{1 + \exp(-z)}$$



➤ Poll: What happens if we scale z by a large constant?

- A) The logistic classifier decision boundary shifts towards class 1
- B) The logistic classifier decision boundary remains same
- C) The logistic classifier tries to separate the data perfectly ←
- D) The logistic classifier allows more mixing of labels on each side of decision boundary

That's M(C)LE. How about M(C)AP?

$$p(\mathbf{w} \mid Y, \mathbf{X}) \propto P(Y \mid \mathbf{X}, \mathbf{w})p(\mathbf{w})$$

$$p(\mathbf{w}) = \prod_i \frac{1}{\kappa \sqrt{2\pi}} e^{-\frac{w_i^2}{2\kappa^2}}$$

- M(C)AP estimate

Zero-mean Gaussian prior

$$\mathbf{w}^* = \arg \max_{\mathbf{w}} \ln \left[\underbrace{p(\mathbf{w})}_{\text{prior}} \prod_{j=1}^n \underbrace{P(y^j \mid \mathbf{x}^j, \mathbf{w})}_{\text{likelihood}} \right]$$

$$\ln p(\mathbf{w}) + \sum_{j=1}^n \ln P(y^j \mid \mathbf{x}^j, \mathbf{w})$$

$$\mathbf{w}^* = \arg \max_{\mathbf{w}} \sum_{j=1}^n \ln P(y^j \mid \mathbf{x}^j, \mathbf{w}) - \underbrace{\sum_{i=1}^d \frac{w_i^2}{2\kappa^2}}_{\text{penalty}}$$

$\propto -\|\mathbf{w}\|^2$

Still concave objective!

Penalizes large weights

Logistic Regression for more than 2 classes

- Logistic regression in more general case, where $Y \in \{y_1, \dots, y_K\}$

for $k < K$

$$\rightarrow P(Y = y_k | X) = \frac{\exp(w_{k0} + \sum_{i=1}^d w_{ki} X_i)}{1 + \sum_{j=1}^{K-1} \exp(w_{j0} + \sum_{i=1}^d w_{ji} X_i)} \quad \leftarrow$$


for $\underline{k=K}$ (normalization, so no weights for this class)

$$P(Y = y_K | X) = \frac{1}{1 + \sum_{j=1}^{K-1} \exp(w_{j0} + \sum_{i=1}^d w_{ji} X_i)}$$

Predict $f^*(x) = \arg \max_{Y=y} P(Y = y | X = x)$

Is the decision boundary still linear? $\leftarrow$

Gaussian Naïve Bayes vs. Logistic Regression



- Representation equivalence (both can yield linear decision boundaries)
 - **But only in a special case!!!** (GNB with class-independent variances)
 - **LR makes no assumptions about $P(X|Y)$ in learning!!!**
 - **Optimize different functions (MLE/MCLE) or (MAP/MCAP)! Obtain different solutions**

Experimental Comparison (Ng-Jordan'01)

UCI Machine Learning Repository 15 datasets, 8 continuous features, 7 discrete features

