Machine Learning for Phenotype Classification

Diagnosing breast cancer using Logistic Regression

Christoph.Lippert@mdc-berlin.de

Fine needle aspartate biopsy images

Diagnosing breast cancer using Logistic Regression

Given:

- **Training Data** with known diagnosis
 - 249 benign •
 - 149 malignant •
- 2 features from pre-processed images •

Task:

- **classify new** biopsies from features
 - c_1 : Malignant (**M**) ٠
 - c_2 : Benign (**B**) ٠

Pre-processing (given)





 c_1 : Malignant (**M**)

 c_2 : Benign (**B**)

Features:

*x*¹ (concavity_mean):

Fraction of chords outside nucleus



*x*² (texture_mean):

Variance in gray-scale intensities •

Linear Classification

Use the training data to find a linear decision function

 $x_1w_1 + x_2w_2 + b = 0$ to **separate** the **two classes** $c_1 = M$ and $c_2 = B$.

or equiv.
$$\mathbf{x}\mathbf{w} = 0$$
 where $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix}$ feature vector (given)
and $\mathbf{w} = \begin{bmatrix} w_1 & w_2 & b \end{bmatrix}$ weight vector (unknown)



Linear Classification

Use the **training data** to find a **linear decision function**

 $x_1w_1 + x_2w_2 + b = 0$ to **separate** the **two classes** $c_1 = M$ and $c_2 = B$.

or equiv.
$$\mathbf{x}\mathbf{w} = 0$$
 where $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix}$ feature vector
(given)
and $\mathbf{w} = \begin{bmatrix} w_1 & w_2 & b \end{bmatrix}$ weight vector
(unknown)

Questions we will answer:

- **Q1.** How to deal with **samples at the boundary**?
 - **A1:** Predict **probabilities** $0 \le p(y = c_1 | \mathbf{x}) \le 1$
- Q2. How to compare different functions? A2: log-loss
- Q3. How to determine the best function? A3: Use optimization
- Q4. How to assess the classifier performance? A4: Quality metrics



Q1. How to deal with uncertain predictions?

Predict probabilities.

The logistic sigmoid:

$$p(y = c_1 | \mathbf{x}) = \pi(\mathbf{x}\mathbf{w})$$
$$= \frac{1}{1 + \exp(-\mathbf{x}\mathbf{w})}$$

By symmetry:

$$p(y = c_2 | \mathbf{x}) = 1 - \pi(\mathbf{x}\mathbf{w})$$



Q2. How to **compare different** functions?

The log-loss function

• **log-error** function for a **single sample**

 $-\ln p(y = c_{true} | \mathbf{x})$

- Large, when assigning low probability
- Small, when assigning high probability
- log-loss function for training data set
 - Sum error functions for all data points

$$loss = -\sum_{n \in c_1} \ln(\pi(\mathbf{x}_n \mathbf{w})) - \sum_{n' \in c_2} \ln(1 - \pi(\mathbf{x}_{n'} \mathbf{w}))$$

Q3. How to **determine** the **best** function?

direction of largest increase in $L(\mathbf{w}^t)$

Update rule:

 $\mathbf{w}^{t+1} = \mathbf{w}^t - \alpha \nabla L(\mathbf{w}^t)$ for a small learning rate α (here, 10^{-4})

direction of largest increase in $L(\mathbf{w}^t)$

Update rule:

 $\mathbf{w}^{t+1} = \mathbf{w}^t - \alpha \nabla L(\mathbf{w}^t)$ for a small learning rate α (here, 10⁻⁴)

Account for the **curvature**!

$$L(\mathbf{w}) = -\sum_{n \in c_{1}} \ln(\pi(\mathbf{x}_{n}\mathbf{w})) - \sum_{n' \in c_{2}} \ln(1 - \pi(\mathbf{x}_{n'}\mathbf{w})) + \lambda \cdot 0.5 \cdot \sum_{d=1}^{D} w_{d}^{2}$$

Hessian:
$$\mathbf{H}_{\mathbf{w}^{t}} = \begin{bmatrix} \partial^{2}L/\partial^{2}w_{1} & \partial^{2}L/\partialw_{1}\partialw_{2} & \dots & \partial^{2}L/\partialw_{1}\partialw_{D} \\ \vdots & \ddots & \vdots \\ \partial^{2}L/\partialw_{D}\partialw_{1} & \partial^{2}L/\partialw_{D}\partialw_{2} & \dots & \partial^{2}L/\partial^{2}w_{D} \end{bmatrix}$$
$$= \underbrace{\left(\mathbf{X}\text{diag}\left(\pi\left(\mathbf{X}\mathbf{w}^{t}\right) \cdot \left(\mathbf{1} - \pi(\mathbf{X}\mathbf{w}^{t})\right)\right)^{T}\mathbf{X} + \underbrace{\lambda \cdot \mathbf{I}_{D \times D}}_{\mathbf{H}_{\mathbf{w}^{t}}(\text{regularizer})}\right)$$

Jpdate rule:

 $\mathbf{w}^{t+1} = \mathbf{w}^{t} - \mathbf{H}_{\mathbf{w}^{t}}^{-1} \nabla L(\mathbf{w}^{t})$ Newton-Raphson algorithm

Account for the **curvature**!

$$L(\mathbf{w}) = -\sum_{n \in c_{1}} \ln(\pi(\mathbf{x}_{n}\mathbf{w})) - \sum_{n' \in c_{2}} \ln(1 - \pi(\mathbf{x}_{n'}\mathbf{w})) + \lambda \cdot 0.5 \cdot \sum_{d=1}^{D} w_{d}^{2}$$

Hessian:

$$\mathbf{H}_{\mathbf{w}^{t}} = \begin{bmatrix} \partial^{2}L/\partial^{2}w_{1} & \partial^{2}L/\partial w_{1}\partial w_{2} & \dots & \partial^{2}L/\partial w_{1}\partial w_{D} \\ \vdots & \ddots & \vdots \\ \partial^{2}L/\partial w_{D}\partial w_{1} & \partial^{2}L/\partial w_{D}\partial w_{2} & \dots & \partial^{2}L/\partial^{2}w_{D} \end{bmatrix}$$

$$= \underbrace{\left(\mathbf{X}\text{diag}\left(\pi\left(\mathbf{X}\mathbf{w}^{t}\right) \cdot \left(\mathbf{1} - \pi(\mathbf{X}\mathbf{w}^{t})\right)\right)^{T}\mathbf{X} + \underbrace{\lambda \cdot \mathbf{I}_{D \times D}}_{\mathbf{H}_{\mathbf{w}^{t}}(\text{regularizer})} \right)$$

Jpdate rule:

 $\mathbf{w}^{t+1} = \mathbf{w}^{t} - \mathbf{H}_{\mathbf{w}^{t}}^{-1} \nabla L(\mathbf{w}^{t})$ Newton-Raphson algorithm

Q4. How to evaluate the model?

Quality measures

- Accept $p(y_1|\mathbf{x}) < 0.5$ to increase sensitivity
- Require $p(y_1|\mathbf{x}) > 0.5$ to increase specificity

Q4. How to evaluate the model?

train vs. test accuracy $= \frac{1}{N} \#$ correct TPsensitivity = $\overline{TP + FN}$ $\frac{TP}{TP + FP}$ specificity = 398 training samples 200 true diagnosis M B 236 13 160 120 80 123 26 40 В Μ predicted diagnosis Accuracy 90% Sensitivity 83% Specificity 90%

	Predicted Positive	Predicted Negative
Positive	True Positive (TP)	False Negative (FN)
Negative	False Positive (FP)	True Negative (TN)

171 test samples

*c*₁ (M)

Observation:

- Lower performance on test data
- Overfitting to training data
- Test errors yield **unbiased** performance estimates

Recap: Diagnosing breast cancer using Logistic Regression

Fine needle aspartate biopsy images

 c_1 : Malignant (**M**)

c₂: Benign (**B**)

Logistic Regression model

- Linear classification
- Predicting probabilities

Model fitting

- log-loss
- Optimizing the log-loss

Model evaluation

- Precision vs. recall
- train vs. test

	Predicted Positive	Predicted Negative
Positive	True Positive (TP)	False Negative (FN)
Negative	False Positive (FP)	True Negative (TN)