		Contents – lecture 3	2(21)
Machine LearningLecture 3 – Linear classificationVision of Systems and Control Department of Information Technology Uppsala University.Email: thomas.schon@it.uu.se, :ww: user.it.uu.se/~thosc112		 Summary of lecture 2 Problem setup Discriminant functions (mainly least squares) Probabilistic generative models Logistic regression (discriminative model) Bayesian logistic regression 	
Machine Learning, Lecture 3 – Linear classification T. Schön, 2014		Machine Learning, Lecture 3 – Linear classification T. Schön, 2014	
Summary of lecture 2 (I	3(21)	Summary of lecture 2 (II/III)	4(21)
Linear regression models the relationship between a continuous target variable <i>t</i> and a possibly nonlinear function $\phi(x)$ of the input variable <i>x</i> , $t_n = \underbrace{w^T \phi(x_n)}_{y(x_n,w)} + \epsilon_n.$ Solved this problem using 1. Maximum Likelihood (ML) 2. Bayesian approach ML with a Gaussian noise model is equivalent to least squares (LS).		Theorem (Gauss-Markov)In a linear regression model $T = \Phi w + E$,where E is white noise with zero mean and covariance R, the bestlinear unbiased estimate (BLUE) of w is $\hat{w} = (\Phi^T R^{-1} \Phi)^{-1} \Phi^T R^{-1} T$, $Cov(\hat{w}) = (\Phi^T R^{-1} \Phi)^{-1}$.Interpretation: The least squares estimator has the smallest meansquare error (MSE) of all linear estimators with no bias, BUT theremay exist a biased estimator with lower MSE.	
Machine Learning, Lecture 3 – Linear classification T. Schön, 2014		Machine Learning, Lecture 3 – Linear classification T. Schön, 2014	

Summary of lecture 2 (III/III) 5(21)	Generative and discriminative models 6(21)
Two potentially biased estimators are ridge regression $(p = 2)$ and the Lasso $(p = 1)$ $\begin{array}{l} \min_{w} \sum_{n=1}^{N} \left(t_n - w^T \phi(x_n)\right)^2 \\ \text{s.t.} \sum_{j=0}^{M-1} w_j ^p \leq \eta \end{array}$ which using a Lagrange multiplier λ can be stated $\begin{array}{l} \min_{w} \sum_{n=1}^{N} \left(t_n - w^T \phi(x_n)\right)^2 + \lambda \sum_{j=0}^{M-1} w_j ^p \end{array}$ Alternative interpretation: The MAP estimate with the likelihood $\prod_{n=1}^{N} (t_n - w^T \phi(x_n))^2 \text{ together with a Gaussian prior leads to ridge regression and together with a Laplacian prior it leads to the LASSO. \end{array}$	Approaches that model the distributions of both the inputs and the outputs are known as generative models . The reason for the name is the fact that using these models we can generate new samples in the input space. Approaches that models the posterior probability directly are referred to as discriminative models .
Machine Learning, Lecture 3 – Linear classification T. Schön, 2014	Machine Learning, Lecture 3 – Linear classification T. Schön, 2014
ML for probabilistic generative models (I/IV) 7(21)	ML for probabilistic generative models (II/IV) 8(21)
ML for probabilistic generative models (I/IV) Consider the two class case, where the class-conditional densities $p(x \mid C_k)$ are Gaussian and the training data is given by $\{x_n, t_n\}_{n=1}^N$. Furthermore, assume that $p(C_1) = \alpha$. The task is now to find the parameters α , μ_1 , μ_2 , Σ by maximizing the likelihood function (i.i.d. Bernoulli), $p(T, X \mid \alpha, \mu_1, \mu_2, \Sigma) = \prod_{n=1}^N (p(x_n, C_1))^{t_n} (p(x_n, C_2))^{1-t_n}$, where $p(x_n, C_1) = p(C_1)p(x_n \mid C_1) = \alpha \mathcal{N}(x_n \mid \mu_1, \Sigma),$ $p(x_n, C_2) = p(C_2)p(x_n \mid C_2) = (1 - \alpha)\mathcal{N}(x_n \mid \mu_2, \Sigma).$	ML for probabilistic generative models (II/IV) Let us now maximize the logarithm of the likelihood function, $L(\alpha, \mu_1, \mu_2, \Sigma) = \ln \left(\prod_{n=1}^N (\alpha \mathcal{N}(x_n \mid \mu_1, \Sigma))^{t_n} ((1 - \alpha) \mathcal{N}(x_n \mid \mu_2, \Sigma))^{1 - t_n} \right)$ The terms that depends on α are $\sum_{n=1}^N (t_n \ln \alpha + (1 - t_n) \ln(1 - \alpha))$ which is maximized by $\widehat{\alpha} = \frac{1}{N} \sum_{n=1}^N t_n = \frac{N_1}{N_1 + N_2}$ (as expected). N_k denotes the number of data in class C_k . Straightforwardly we get $\widehat{\mu}_1 = \frac{1}{N_1} \sum_{n=1}^N t_n x_n, \qquad \widehat{\mu}_2 = \frac{1}{N_2} \sum_{n=1}^N (1 - t_n) x_n.$

ML for probabilistic generative models (III/IV) ML for probabilistic generative models (IV/IV) 9(21) 10(21)Lemma (Useful matrix derivatives) $L(\Sigma) = -\frac{1}{2} \sum_{n=1}^{N} t_n \ln \det \Sigma - \frac{1}{2} \sum_{n=1}^{N} t_n (x_n - \mu_1)^T \Sigma^{-1} (x_n - \mu_1)$ $\frac{\partial}{\partial M} \ln \det M = M^{-T},$ $\frac{\partial}{\partial M} \operatorname{Tr} \left(M^{-1} N \right) = -M^{-T} N^T M^{-T}.$ $-\frac{1}{2}\sum_{i=1}^{N}(1-t_{n})\ln\det\Sigma-\frac{1}{2}\sum_{i=1}^{N}(1-t_{n})(x_{n}-\mu_{2})^{T}\Sigma^{-1}(x_{n}-\mu_{2})$ Differentiating $L(\Sigma) = -\frac{N}{2} \ln \det \Sigma - \frac{N}{2} \operatorname{Tr} (\Sigma^{-1}S)$ results in Using the fact that $x^{T}Ax = \text{Tr}(Axx^{T})$ we have $\frac{\partial L}{\partial \Sigma} = -\frac{N}{2}\Sigma^{-T} + \frac{N}{2}\Sigma^{-T}S\Sigma^{-T}$ $L(\Sigma) = -\frac{N}{2} \ln \det \Sigma - \frac{N}{2} \operatorname{Tr} \left(\Sigma^{-1} S \right),$ Hence, $\Sigma = S$ $\frac{\partial L}{\partial \Sigma} = 0$ where $S = \frac{1}{N} \sum_{n=1}^{N} \left(t_n (x_n - \mu_1) (x_n - \mu_1)^T + (1 - t_n) (x_n - \mu_2) (x_n - \mu_2)^T \right)$ More results on matrix derivatives are available in Magnus, J. R., & Neudecker, H. (1999). Matrix differential calculus with applications in statistics and econometrics. 2nd Edition, UK: Wiley Machine Learning, Lecture 3 - Linear classification Machine Learning, Lecture 3 - Linear classification T. Schön, 2014 T. Schön, 2014 Generalized linear models for classification Gradient of L(w) for logistic regression (I/II) 11(21)12(21)The negative log-likelihood is In linear regression we made use of a linear model $L(w) = -\sum_{n=1}^{N} (t_n \ln y_n + (1 - t_n) \ln(1 - y_n)),$ $t_n = y(x, w) = w^T \phi(x_n) + \epsilon_n.$ where For classification problems the target variables are discrete, or $y_n = \sigma(a_n) = \frac{1}{1 + \exp(-a_n)}$, and $a_n = w^T \phi_n$. slightly more general, posterior probabilities in the range (0, 1). This is achieved using a so called *activation* function $f(f^{-1} \text{ must exist})$, Using the chain rule we have, $y(x) = f(w^T x + w_0).$ (1) $g = \frac{\partial L}{\partial w} = \sum_{n=1}^{N} \frac{\partial L}{\partial y_n} \frac{\partial y_n}{\partial a_n} \frac{\partial a_n}{\partial w}$ Note that the decision surface corresponds to y(x) = constant,

where

$$\frac{\partial L}{\partial y_n} = \frac{1-t_n}{1-y_n} - \frac{t_n}{y_n} = \frac{y_n - t_n}{y_n(1-y_n)}$$

Machine Learning, Lecture 3 – Linear classification T. Schön, 2014

Machine Learning, Lecture 3 – Linear classification T. Schön, 2014

name generalized linear model for (1).

implying that $w^T x + w_0 = \text{constant}$. This means that the decision

surface is a linear function of x, even if f is nonlinear. Hence, the

Gradient of L(w) for logistic regression (II/II)

Furthermore,

$$\frac{\partial y_n}{\partial a_n} = \frac{\partial \sigma(a_n)}{\partial a_n} = \dots = \sigma(a_n)(1 - \sigma(a_n)) = y_n(1 - y_n),$$
$$\frac{\partial a_n}{\partial w} = \phi_n.$$

which results in the following expression for the gradient

$$g = \sum_{n=1}^{N} (y_n - t_n)\phi_n = \Phi^T (Y - T),$$

where

$$\Phi = \begin{pmatrix} \phi_0(x_1) & \phi_1(x_1) & \dots & \phi_{M-1}(x_1) \\ \phi_0(x_2) & \phi_1(x_2) & \dots & \phi_{M-1}(x_2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_0(x_N) & \phi_1(x_N) & \dots & \phi_{M-1}(x_N) \end{pmatrix} \quad Y = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{pmatrix} \quad T = \begin{pmatrix} t_1 \\ t_2 \\ \vdots \\ t_N \end{pmatrix}$$

Machine Learning, Lecture 3 – Linear classification T. Schön, 2014

Bayesian logistic regression

15(21)

13(21)

Recall that

$$p(T \mid w) = \prod_{n=1}^{N} \sigma(w^{T} \phi_{n})^{t_{n}} \left(1 - \sigma(w^{T} \phi_{n})\right)^{1-t_{n}}$$

Hence, computing the posterior density

$$p(w \mid T) = \frac{p(T \mid w)p(w)}{p(T)}$$

is intractable. We are forced to an approximation. Three alternatives

- 1. Laplace approximation (this lecture)
- 2. VB & EP (lecture 7)
- 3. Sampling methods, e.g., MCMC (lecture 10)

Hessian of L(w) for logistic regression

14(21)

$H = \frac{\partial^2 L}{\partial w \partial w^T} = \cdots = \sum_{n=1}^N (y_n - t_n) \phi_n \phi_n^T = \Phi^T R \Phi$

where

$$R = \begin{pmatrix} y_1(1-y_1) & 0 & \dots & 0 \\ 0 & y_2(1-y_2) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & y_N(1-y_N) \end{pmatrix}$$

Machine Learning, Lecture 3 – Linear classification T. Schön, 2014

Laplace approximation (I/III)

16(21)

The Laplace approximation is a simple approximation that is obtained by fitting a Gaussian centered around the (MAP) mode of the distribution.

Consider the density function $p(\boldsymbol{z})$ of a scalar stochastic variable $\boldsymbol{z},$ given by

$$p(z) = \frac{1}{Z}f(z),$$

where $Z = \int f(z) dz$ is the normalization coefficient.

1. Find the MAP mode: We start by finding a mode z_0 of the density function,

$$\left.\frac{df(z)}{dz}\right|_{z=z_0} = 0$$

Machine Learning, Lecture 3 – Linear classification T. Schön, 2014

Laplace approximation (II/III)17(21)	Laplace approximation (III/III) 18(21)
2. Compute a Gaussian approximation around the MAP mode: Start by considering a Taylor expansion of $\ln f(z)$ around z_0 , $\ln f(z) \approx \ln f(z_0) + \frac{d}{dz} \ln f(z) \Big _{z=z_0} (z-z_0) + \frac{1}{2} \frac{d^2}{dz^2} \ln f(z) \Big _{z=z_0} (z-z_0)^2$ $= \ln f(z_0) - \frac{A}{2} (z-z_0)^2$, (2) where $A = -\frac{d^2}{dz^2} \ln f(z) \Big _{z=z_0}$ Taking the exponential of both sides in (2) results in $f(z) \approx f(z_0) \exp\left(-\frac{A}{2} (z-z_0)^2\right)$	By normalizing this expression we have now obtained a Gaussian approximation $q(z) = \left(\frac{A}{2\pi}\right)^{1/2} \exp\left(-\frac{A}{2}(z-z_0)^2\right)$ where $A = -\frac{d^2}{dz^2} \ln f(z)\Big _{z=z_0}$ The main limitation of the Laplace approximation is that it is a local method that only captures aspects of the true density around a specific value z_0 .
Machine Learning, Lecture 3 – Linear classification T. Schön, 2014	Machine Learning, Lecture 3 – Linear classification T. Schön, 2014
Bayesian logistic regression (I/II) 19(21)	Bayesian logistic regression (II/II) 20(21)
The posterior is $p(w \mid T) \propto p(T \mid w)p(w), \qquad (3)$ where we assume a Gaussian prior $p(w) = \mathcal{N}(w \mid m_0, S_0)$ and make use of the Laplace approximation. Taking the logarithm of both sides of (3) gives $\ln p(w \mid t) = -\frac{1}{2}(w - m_0)^T S_0^{-1}(w - m_0)$ $+ \sum_{n=1}^N (t_n \ln y_n + (1 - t_n) \ln(1 - y_n)) + \text{const.}$	Using the Laplace approximation we can now obtain a Gaussian approximation $q(w) = \mathcal{N}(w \mid w_{\text{MAP}}, S_N)$ where w_{MAP} is the MAP estimate of $p(w \mid T)$ and the covariance S_N is the Hessian of $\ln p(w \mid T)$, $S_N = \frac{\partial^2}{\partial w \partial w^T} \ln p(w \mid T) = S_0^{-1} + \sum_{n=1}^N y_n (1 - y_n) \phi_n \phi_n^T$ Based on this distribution we can now start making predictions for
where $y_n = \sigma(w^T \phi_n)$.	Recall that prediction corresponds to marginalization w.r.t. w .

A few concepts to summarize lecture 3

Classification: The goal of classification is to assign an input vector x to one of K classes, C_{k} , $k = 1, \ldots, K$.

Discriminant: A discriminant is a function that takes an input *x* and assigns it to one of *K* classes.

Generative models: Approaches that model the distributions of both the inputs and the outputs are known as generative models. In classification this amounts to modelling the class-conditional densities $p(x \mid C_k)$, as well as the prior densities $p(C_k)$. The reason for the name is the fact that using these models we can generate new samples in the input space.

Discriminative models: Approaches that models the posterior probability directly are referred to as discriminative models.

Logistic Regression: Discriminative model that makes direct use of a generalized linear model in the form of a logistic sigmoid to solve the classification problem.

Laplace approximation: A local approximation method that finds the mode of the posterior distribution and then fits a Gaussian centered at that mode.

Machine Learning, Lecture 3 – Linear classification T. Schön, 2014