

PROBABILISTIC MACHINE LEARNING  
LECTURE 13  
GAUSSIAN PROCESS CLASSIFICATION


Philipp Hennig

08 June 2020

EBERHARD KARLS  
UNIVERSITÄT  
TÜBINGEN



FACULTY OF SCIENCE  
DEPARTMENT OF COMPUTER SCIENCE  
CHAIR FOR THE METHODS OF MACHINE LEARNING




#	date	content	Ex	#	date	content	Ex
1	20.04.	Introduction	1	14	09.06.	Generalized Linear Models	
2	21.04.	Reasoning under Uncertainty		15	15.06.	Exponential Families	8
3	27.04.	Continuous Variables	2	16	16.06.	Graphical Models	
4	28.04.	Monte Carlo		17	22.06.	Factor Graphs	9
5	04.05.	Markov Chain Monte Carlo	3	18	23.06.	The Sum-Product Algorithm	
6	05.05.	Gaussian Distributions		19	29.06.	Example: Topic Models	10
7	11.05.	Parametric Regression	4	20	30.06.	Mixture Models	
8	12.05.	Learning Representations		21	06.07.	EM	11
9	18.05.	Gaussian Processes	5	22	07.07.	Variational Inference	
10	19.05.	Understanding Kernels		23	13.07.	Topics	
12	26.05.	Gauss-Markov Models		25	20.07.	Example: Kernel Topic Models	
11	25.05.	An Example for GP Regression	6	24	14.07.	Example: Inferringevision	
13	08.06.	GP Classification	7	26	21.07.	Revision	




# Classification Problems

a typography of supervised learning problems




# Classification Problems

a typography of supervised learning problems




# Classification Problems

a typography of supervised learning problems




# Classification Problems

a typography of supervised learning problems




# Classification Problems

a typography of supervised learning problems



# Classification Problems


a typography of supervised learning problems







# Classification Problems

a typography of supervised learning problems



# Classification Problems

a typography of supervised learning problems



<https://github.com/zalandoresearch/fashion-mnist>

# Classification vs. Regression

Two types of supervised learning problems

## Regression:

Given supervised *data* (special case  $d = 1$ : univariate regression)

$$(X, Y) := (x_i, y_i)_{i=1, \dots, n} \text{ with } x_i \in \mathbb{X}, y_i \in \mathbb{R}^d$$

find function  $f : \mathbb{X} \rightarrow \mathbb{R}^d$  such that  $f$  "models"  $Y \approx f(X)$ .

## Classification:

Given supervised *data* (special case  $d = 2$ : binary classification)

$$(X, Y) := (x_i, c_i)_{i=1, \dots, n} \text{ with } x_i \in \mathbb{X}, c_i \in \{1, \dots, d\}$$

find probability  $\pi : \mathbb{X} \rightarrow U^d$  ( $U^d = \{p \in [0, 1]^d : \sum_{i=1}^d p_i = 1\}$ ) such that  $\pi$  "models"  $y_i \sim \pi_{x_i}$ .

*Regression* predicts a **function**, *classification* predicts a **probability**.

Until further notice, consider only discriminative **binary** classification:

$$y \in \{-1; +1\} \quad x \mapsto \pi(x) =: \pi_x \in [0, 1]$$

$$p(y | x) = \begin{cases} \pi(x) & \text{if } y = 1 \\ 1 - \pi(x) & \text{if } y = -1 \end{cases}$$



Until further notice, consider only discriminative **binary** classification:


$$\begin{aligned}
 y \in \{-1; +1\} \quad x \mapsto \pi(x) &=: \pi_x \in [0, 1] \\
 p(y | x) &= \begin{cases} \pi(x) & \text{if } y = 1 \\ 1 - \pi(x) & \text{if } y = -1 \end{cases}
 \end{aligned}$$

Discriminative learning phrased probabilistically:

- ▶ We would like to *learn*  $\pi_x(y) = p(y | x)$
- ▶ This is *almost* like regression:  $p(y | x) = \mathcal{N}(y; f_x, \sigma^2) = \pi_x(y)$
- ▶ only the *domain* is wrong:  $y \in \{-1; 1\}$  vs.  $y \in \mathbb{R}$ .

# Let's not throw out Gauss just yet

Turning Gaussian process models into discrete likelihoods



$$\pi_f = \sigma(f) = \frac{1}{1 + \exp(-f)} = \int_{-\infty}^f \frac{1}{4} \operatorname{sech}^2\left(\frac{x}{2}\right) dx$$

$$\sigma(f) = 1 - \sigma(-f) \quad f(\pi) = \ln \pi - \ln(1 - \pi) \quad \frac{d\pi}{df} = \pi(f) \cdot (1 - \pi(f))$$

# CODE

`Generative Model for Logistic Regression.ipynb`





$$p(f) = \mathcal{GP}(f; m, k)$$
$$p(y | f_x) = \sigma(yf_x) = \begin{cases} \sigma(f) & \text{if } y = 1 \\ 1 - \sigma(f) & \text{if } y = -1 \end{cases} \quad \text{using } \sigma(x) = 1 - \sigma(-x).$$


$$\begin{aligned} p(f) &= \mathcal{GP}(f; m, k) \\ p(y | f_x) &= \sigma(yf_x) = \begin{cases} \sigma(f) & \text{if } y = 1 \\ 1 - \sigma(f) & \text{if } y = -1 \end{cases} \quad \text{using } \sigma(x) = 1 - \sigma(-x). \end{aligned}$$

The problem: The posterior is not Gaussian!

$$\begin{aligned} p(f_X | Y) &= \frac{p(Y | f_X)p(f_X)}{p(Y)} = \frac{\mathcal{N}(f_X; m, k) \prod_{i=1}^n \sigma(y_i f_{x_i})}{\int \mathcal{N}(f_X; m, k) \prod_{i=1}^n \sigma(y_i f_{x_i}) df_X} \\ \log p(f_X | Y) &= -\frac{1}{2} f_X^T k_{XX}^{-1} f_X + \sum_{i=1}^n \log \sigma(y_i f_{x_i}) + \text{const.} \end{aligned}$$


# Logistic Regression is non-analytic

We'll have to break out the toolbox




# Logistic Regression is non-analytic

We'll have to break out the toolbox



# Logistic Regression is non-analytic

We'll have to break out the toolbox



We do not always care about all details of the posterior, just “key aspects”.

- ▶ Remember that the Gaussian choice was also one of convenience.
- ▶ **Moments** of  $p(f, y) = p(y | f)p(f)$  we may be interested in

$$\mathbb{E}_p(1) = \int p(y, f) df = \int 1 \cdot p(y, f) df = Z \quad \text{the evidence}$$

$$\mathbb{E}_{p(f|y)}(f) = \int f \cdot p(f | y) df = \frac{1}{Z} \int f \cdot p(f, y) df = \bar{f} \quad \text{the mean}$$

$$\mathbb{E}_{p(f|y)}(f^2) - \bar{f}^2 = \int f^2 \cdot p(f | y) df - \bar{f}^2 = \frac{1}{Z} \int f^2 \cdot p(f, y) df - \bar{f}^2 = \text{var}(f) \quad \text{the variance}$$

$Z$  for hyperparameter tuning

$\bar{f}$  as a point estimator

$\text{var}(f)$  as an error estimator

Unfortunately, all these are usually intractable. But we can aim to approximate them.

# The Toolbox

---

## Framework:

$$\int p(x_1, x_2) dx_2 = p(x_1)$$

$$p(x_1, x_2) = p(x_1 | x_2)p(x_2)$$

$$p(x | y) = \frac{p(y | x)p(x)}{p(y)}$$


---

## Modelling:

- ▶ Directed Graphical Models
- ▶ Gaussian Distributions
- ▶ Kernels
- ▶ Markov Chains
- ▶
- ▶

## Computation:

- ▶ Monte Carlo
  - ▶ Linear algebra / Gaussian inference
  - ▶ maximum likelihood / MAP
  - ▶ Laplace approximations
  - ▶
- 




# The Laplace Approximation

A local Gaussian approximation



Pierre Simon M. de Laplace, 1814





# An idea as old as Probabilistic Inference

Théorie analytique des probabilités, 1814



Pierre-Simon, marquis de Laplace  
(1749-1827)

364

## THÉORIE ANALYTIQUE

L'intégrale du numérateur étant prise depuis  $x = \theta$  jusqu'à  $x = \theta'$ , et celle du dénominateur étant prise depuis  $x = 0$  jusqu'à  $x = 1$

La valeur de  $x$  la plus probable, est celle qui rend  $y$  un *maximum*. Nous la désignerons par  $a$ . Si aux limites de  $x$ ,  $y$  est nul, alors chaque valeur de  $y$  a une valeur égale correspondante de l'autre côté du *maximum*.

Quand les valeurs de  $x$ , considérées indépendamment du résultat observé, ne sont pas également possibles; en nommant  $x$  la fonction de  $x$  qui exprime leur probabilité; il est facile de voir, par ce qui a été dit dans le premier chapitre de ce Livre, qu'en changeant dans la formule (1),  $y$  dans  $yz$ , on aura la probabilité que la valeur de  $x$  est comprise dans les limites  $x = \theta$  et  $x = \theta'$ . Cela revient à supposer toutes les valeurs de  $x$  également possibles *a priori*, et à considérer le résultat observé, comme étant formé de deux résultats indépendants, dont les probabilités sont  $y$  et  $x$ . On peut donc ramener ainsi tous les cas à celui où l'on suppose *a priori*, avant l'événement, une égale possibilité aux différentes valeurs de  $x$ , et par cette raison, nous adopterons cette hypothèse dans ce qui va suivre.

Nous avons donné dans les n° 22 et suivans du premier Livre, les formules nécessaires pour déterminer par des approximations convergentes, les intégrales du numérateur et du dénominateur de la formule (1), lorsque les événements simples dont se compose l'événement observé, sont répétés un très-grand nombre de fois; car alors  $y$  a pour facteurs, des fonctions de  $x$  élevées à de grandes puissances. Nous allons, au moyen de ces formules, déterminer la loi de probabilité des valeurs de  $x$ , à mesure qu'elles s'éloignent de la valeur  $a$ , la plus probable, ou qui rend  $y$  un *maximum*. Pour cela, reprenons la formule (c) du n° 27 du premier Livre,

$$fydx = Y \cdot \left\{ U + \frac{1}{2} \cdot \frac{d^2 U^2}{1.2. dx^2} + \frac{1.5}{2^2} \cdot \frac{d^4 U^4}{1.2.3.4. dx^4} + \text{etc.} \right\} \cdot \int dx \cdot c^{-x} \\ + \frac{y}{a} \cdot c^{-T} \cdot \left\{ \frac{dU}{dx} - T \cdot \frac{d^2 U^2}{1.2. dx^2} + (T^2 + 1) \cdot \frac{d^4 U^4}{1.2.3.4. dx^4} - \text{etc.} \right\}; \quad (2) \\ - \frac{y}{a} \cdot c^{-T'} \cdot \left\{ \frac{dU}{dx} + T' \cdot \frac{d^2 U^2}{1.2. dx^2} + (T'^2 + 1) \cdot \frac{d^4 U^4}{1.2.3.4. dx^4} + \text{etc.} \right\}$$

## DES PROBABILITÉS.

365

$\psi$  est égal à  $\frac{x-a}{\sqrt{\log Y - \log y}}$ , et  $U, \frac{dU}{dx}, \frac{d^2 U^2}{dx^2}$ , etc. sont ce que deviennent  $\psi, \frac{d\psi}{dx}, \frac{d^2 \psi^2}{dx^2}$ , etc., lorsqu'on y change après les différentiations,  $x$  en  $a$ ,  $a$  étant la valeur de  $x$  qui rend  $y$  un *maximum*:  $T$  est égal à ce que devient la fonction  $\sqrt{\log Y - \log y}$ , lorsqu'on change  $x$  en  $a - \theta$  dans  $y$ , et  $T'$  est ce que devient la même fonction, lorsqu'on y change  $x$  dans  $a + \theta'$ . L'expression précédente de  $fydx$  donne la valeur de cette intégrale, dans les limites  $x = a - \theta$  et  $x = a + \theta'$ ; l'intégrale  $\int dx \cdot c^{-x}$  étant prise depuis  $t = -T$  jusqu'à  $t = T'$ .

Le plus souvent, aux limites de l'intégrale  $fydx$ , étendue depuis  $x = 0$  jusqu'à  $x = 1$ ,  $y$  est nul; ou lorsque  $y$  n'est pas nul, il devient si petit à ces limites, qu'on peut le supposer nul. Alors, on peut faire à ces limites  $T$  et  $T'$  infinis, ce qui donne pour l'intégrale  $fydx$ , étendue depuis  $x = 0$  jusqu'à  $x = 1$ ,

$$fydx = Y \cdot \left\{ U + \frac{1}{2} \cdot \frac{d^2 U^2}{1.2. dx^2} + \frac{1.5}{2^2} \cdot \frac{d^4 U^4}{1.2.3.4. dx^4} + \text{etc.} \right\} \cdot \sqrt{\pi};$$

ainsi la probabilité que la valeur de  $x$  est comprise dans les limites  $x = a - \theta$  et  $x = a + \theta'$ , est égale à

$$\frac{\int dx \cdot c^{-x}}{\sqrt{\pi}} + \left\{ \frac{1}{2} \cdot c^{-T} \cdot \left\{ \frac{dU}{dx} - T \cdot \frac{d^2 U^2}{1.2. dx^2} + (T^2 + 1) \cdot \frac{d^4 U^4}{1.2.3.4. dx^4} - \text{etc.} \right\} \right. \\ \left. - \frac{1}{2} \cdot c^{-T'} \cdot \left\{ \frac{dU}{dx} + T' \cdot \frac{d^2 U^2}{1.2. dx^2} + (T'^2 + 1) \cdot \frac{d^4 U^4}{1.2.3.4. dx^4} + \text{etc.} \right\} \right\}; \quad (3) \\ \left\{ U + \frac{1}{2} \cdot \frac{d^2 U^2}{1.2. dx^2} + \frac{1.5}{2^2} \cdot \frac{d^4 U^4}{1.2.3.4. dx^4} + \text{etc.} \right\} \cdot \sqrt{\pi}.$$

On voit par le n° 23 du premier Livre, que dans le cas où  $y$  a pour facteurs, des fonctions de  $x$  élevées à de grandes puissances de l'ordre  $\frac{1}{a}$ ,  $a$  étant une fraction extrêmement petite, alors  $U$  est le plus souvent de l'ordre  $\sqrt{a}$ , ainsi que ses différences successives;  $U, \frac{dU}{dx}, \frac{d^2 U^2}{dx^2}$ , etc. sont respectivement des ordres  $\sqrt{a}, a, a^2$ , etc.; d'où il suit que la convergence de séries de la formule (3), exige que  $T$  et  $T'$  ne soient pas d'un ordre supérieur à  $\frac{1}{\sqrt{a}}$ .

- ▶ Consider a probability distribution  $p(\theta)$  (may be a posterior  $p(\theta | D)$  or something else)
- ▶ find a (local) **maximum** of  $p(\theta)$  or (equivalently)  $\log p(\theta)$

$$\hat{\theta} = \arg \max \log p(\theta) \quad \Rightarrow \quad \nabla \log p(\hat{\theta}) = 0$$

- ▶ perform **second order Taylor expansion** around  $\theta = \hat{\theta} + \delta$  in log space

$$\log p(\delta) = \log p(\hat{\theta}) + \frac{1}{2} \delta^\top \left( \underbrace{\nabla \nabla^\top \log p(\hat{\theta})}_{=: \Psi} \right) \delta + \mathcal{O}(\delta^3)$$

- ▶ define the **Laplace approximation**  $q$  to  $p$

$$q(\theta) = \mathcal{N}(\theta; \hat{\theta}, -\Psi^{-1})$$

- ▶ Note that, if  $p(\theta) = \mathcal{N}(\theta; m, \Sigma)$ , then  $p(\theta) = q(\theta)$

- ▶ Find maximum posterior probability for **latent  $f$**  at **training points**

$$\hat{\mathbf{f}} = \arg \max \log p(\mathbf{f}_X | y)$$

- ▶ Assign approximate Gaussian posterior at training points

$$q(\mathbf{f}_X) = \mathcal{N}(\mathbf{f}_X; \hat{\mathbf{f}}, -(\nabla \nabla^\top \log p(\mathbf{f}_X | y)|_{\mathbf{f}_X = \hat{\mathbf{f}}})^{-1}) =: \mathcal{N}(\mathbf{f}_X; \hat{\mathbf{f}}, \hat{\Sigma})$$

- ▶ approximate posterior **predictions** at  $f_x$  for **latent function**

$$\begin{aligned} q(f_x | y) &= \int p(f_x | \mathbf{f}_X) q(\mathbf{f}_X) d\mathbf{f}_X = \int \mathcal{N}(f_x; m_x + k_{xX} K_{XX}^{-1} (\mathbf{f}_X - m_X), k_{xx} - k_{xX} K_{XX}^{-1} k_{Xx}) q(\mathbf{f}_X) d\mathbf{f}_X \\ &= \mathcal{N}(f_x; m_x + k_{xX} K_{XX}^{-1} (\hat{\mathbf{f}} - m_X), k_{xx} - k_{xX} K_{XX}^{-1} k_{Xx} + k_{xX} K_{XX}^{-1} \hat{\Sigma} K_{XX}^{-1} k_{Xx}) \end{aligned}$$

Compare with exact predictions

$$\mathbb{E}_{p(f_x, \mathbf{f}_X | y)}(f_x) = \int (\mathbb{E}_{p(\mathbf{f}_X | \mathbf{f}_X)}(f_x)) p(\mathbf{f}_X | y) d\mathbf{f}_X = m_x + k_{xX} K_{XX}^{-1} (\mathbb{E}_{p(\mathbf{f}_X | y)}(\mathbf{f}_X) - m_X) =: \bar{f}_x$$

Recall:  $p(x) = \mathcal{N}(x; m, V)$ ,  $p(z | x) = \mathcal{N}(z; Ax, B) \Rightarrow p(z) = \int p(z | x) p(x) dx = \mathcal{N}(z; Am, AVA^\top + B)$ .

- Find maximum posterior probability for **latent  $f$**  at **training points**

$$\hat{\mathbf{f}} = \arg \max \log p(\mathbf{f}_X | y)$$

- Assign approximate Gaussian posterior at training points

$$q(\mathbf{f}_X) = \mathcal{N}(\mathbf{f}_X; \hat{\mathbf{f}}, -(\nabla \nabla^\top \log p(\mathbf{f}_X | y)|_{\mathbf{f}_X = \hat{\mathbf{f}}})^{-1}) =: \mathcal{N}(\mathbf{f}_X; \hat{\mathbf{f}}, \hat{\Sigma})$$

- approximate posterior **predictions** at  $f_x$  for **latent function**

$$\begin{aligned} q(f_x | y) &= \int p(f_x | \mathbf{f}_X) q(\mathbf{f}_X) d\mathbf{f}_X = \int \mathcal{N}(f_x; m_x + k_{xx} K_{XX}^{-1} (\mathbf{f}_X - m_X), k_{xx} - k_{xx} K_{XX}^{-1} k_{XX}) q(\mathbf{f}_X) d\mathbf{f}_X \\ &= \mathcal{N}(f_x; m_x + k_{xx} K_{XX}^{-1} (\hat{\mathbf{f}} - m_X), k_{xx} - k_{xx} K_{XX}^{-1} k_{XX} + k_{xx} K_{XX}^{-1} \hat{\Sigma} K_{XX}^{-1} k_{XX}) \end{aligned}$$

Compare with exact predictions

$$\text{var}_{p(f_x, \mathbf{f}_X | y)}(f_x) = \int (f_x - \bar{f}_x)^2 dp(f_x | \mathbf{f}_X) dp(\mathbf{f}_X) = k_{xx} - k_{xx} K_{XX}^{-1} k_{XX} + k_{xx} K_{XX}^{-1} \text{var}_{p(\mathbf{f}_X | y)}(\mathbf{f}_X) K_{XX}^{-1} k_{XX}$$

Recall:  $p(x) = \mathcal{N}(x; m, V)$ ,  $p(z | x) = \mathcal{N}(z; Ax, B) \Rightarrow p(z) = \int p(z | x) p(x) dx = \mathcal{N}(z; Am, AVA^\top + B)$ .

- ▶ Find maximum posterior probability for latent  $f$  at training points

$$\hat{\mathbf{f}} = \arg \max \log p(\mathbf{f}_X | y)$$

- ▶ Assign approximate Gaussian posterior at training points

$$q(\mathbf{f}_X) = \mathcal{N}(\mathbf{f}_X; \hat{\mathbf{f}}, -(\nabla \nabla^\top \log p(\mathbf{f}_X | y)|_{\mathbf{f}_X=\hat{\mathbf{f}}})^{-1}) =: \mathcal{N}(\mathbf{f}_X; \hat{\mathbf{f}}, \hat{\Sigma})$$

- ▶ approximate posterior **predictions** at  $f_x$  for **latent function**

$$\begin{aligned} q(f_x | y) &= \int p(f_x | \mathbf{f}_X) q(\mathbf{f}_X) d\mathbf{f}_X = \int \mathcal{N}(f_x; m_x + k_{xx} K_{XX}^{-1} (\mathbf{f}_X - m_X), k_{xx} - k_{xx} K_{XX}^{-1} k_{xx}) q(\mathbf{f}_X) d\mathbf{f}_X \\ &= \mathcal{N}(f_x; m_x + k_{xx} K_{XX}^{-1} (\hat{\mathbf{f}} - m_X), k_{xx} - k_{xx} K_{XX}^{-1} k_{xx} + k_{xx} K_{XX}^{-1} \hat{\Sigma} K_{XX}^{-1} k_{xx}) \end{aligned}$$

- ▶ compute predictions for **label probabilities**:

$$\mathbb{E}_{p(f|y)}[\pi_x] \approx \mathbb{E}_q[\pi_x] = \int \sigma(f_x) q(f_x | y) df_x \quad \text{or (not the same!) } \hat{\pi}_x = \sigma(\mathbb{E}_q(f_x))$$

- ▶ the Laplace approximation is only very roughly motivated (see above)
- ▶ it can be **arbitrarily wrong**, since it is a **local** approximation
- ▶ but it is often among the most computationally efficient things to try
- ▶ for logistic regression, it tends to work relatively well, because
  - ▶ the log posterior is concave (see below)
  - ▶ the algebraic structure of the link function yields “almost” a Gaussian posterior (cf. picture above)



$$p(f) = \mathcal{GP}(f, m, k) \quad p(\mathbf{y} | \mathbf{f}_X) = \prod_{i=1}^n \sigma(y_i f_{X_i}) \quad \sigma(z) = \frac{1}{1 + e^{-z}}$$

$$\begin{aligned} \log p(\mathbf{f}_X | \mathbf{y}) &= \log p(\mathbf{y} | \mathbf{f}_X) + \log p(\mathbf{f}_X) - \log p(\mathbf{y}) \quad \text{with} \quad \log \sigma(y_i f_{X_i}) = -\log(1 + e^{-y_i f_{X_i}}) \\ &= \sum_{i=1}^n \log \sigma(y_i f_{X_i}) - \frac{1}{2} (\mathbf{f}_X - \mathbf{m}_X)^\top K_{XX}^{-1} (\mathbf{f}_X - \mathbf{m}_X) + \text{const.} \end{aligned}$$

$$\nabla \log p(\mathbf{f}_X | \mathbf{y}) = \sum_{i=1}^n \nabla \log \sigma(y_i f_{X_i}) - K_{XX}^{-1} (\mathbf{f}_X - \mathbf{m}_X) \quad \text{with} \quad \frac{\partial \log \sigma(y_i f_{X_i})}{\partial f_{X_j}} = \delta_{ij} \left( \frac{y_i + 1}{2} - \sigma(f_{X_i}) \right)$$

$$\begin{aligned} \nabla \nabla^\top \log p(\mathbf{f}_X | \mathbf{y}) &= \sum_{i=1}^n \nabla \nabla^\top \log \sigma(y_i f_{X_i}) - K_{XX}^{-1} \quad \text{with} \quad \frac{\partial^2 \log \sigma(y_i f_{X_i})}{\partial f_{X_a} \partial f_{X_b}} = -\delta_{ia} \delta_{ib} \underbrace{\sigma(f_{X_i})(1 - \sigma(f_{X_i}))}_{=: w_i \text{ with } 0 < w_i < 1} \end{aligned}$$

$$=: -\text{diag } \mathbf{w} - K^{-1} = -(W + K^{-1}) \quad \leftarrow \text{convex minimization / concave maximization!}$$

# Implementing the Laplace Approximation I

Newton Optimization

[Rasmussen & Williams, 2006, §3.4]

```
1 procedure OPTIMIZE( $L(\cdot), f_0$ )
2    $f \leftarrow f_0$  // initialize
3   while not converged do
4      $g \leftarrow \nabla L(f)$  // compute gradient
5      $H \leftarrow (\nabla \nabla^\top L(f))^{-1}$  // compute inverse Hessian
6      $\Delta \leftarrow Hg$  // Newton step
7      $f \leftarrow f - \Delta$  // perform step
8     converged  $\leftarrow \|\Delta\| < \epsilon$  // check for convergence
9   end while
10  return  $f$ 
11 end procedure
```



```
1 procedure GP-LOGISTIC-TRAIN( $K_{XX}, m_X, y$ )
2    $f \leftarrow m_X$  // initialize
3   while not converged do
4      $r \leftarrow \frac{y+1}{2} - \sigma(f)$  //  $= \nabla \log p(y | f_x)$ , gradient of log likelihood
5      $W \leftarrow \text{diag}(\sigma(f) \odot (1 - \sigma(f)))$  //  $= -\nabla \nabla \log p(y | f_x)$ , Hessian of log likelihood
6      $g \leftarrow r - K_{XX}^{-1}(f - m_X)$  // compute gradient
7      $H \leftarrow -(W + K^{-1})^{-1}$  // compute inverse Hessian
8      $\Delta \leftarrow Hg$  // Newton step
9      $f \leftarrow f - \Delta$  // perform step
10    converged  $\leftarrow \|\Delta\| < \epsilon$  // check for convergence
11  end while
12  return  $f$ 
13 end procedure
```

This can be numerically unstable as it (repeatedly) requires  $(W + K^{-1})^{-1}$ . For a numerically stable alternative, use  $B := I + W^{1/2}K_{XX}W^{1/2}$  (cf. Rasmussen & Williams).

$$\log p(\mathbf{f}_X | \mathbf{y}) = \sum_{i=1}^n \log \sigma(y_i f_{x_i}) - \frac{1}{2} (\mathbf{f}_X - \mathbf{m}_X)^\top K_{XX}^{-1} (\mathbf{f}_X - \mathbf{m}_X) + \text{const.}$$

$$\nabla \log p(\mathbf{f}_X | \mathbf{y}) = \underbrace{\sum_{i=1}^n \nabla \log \sigma(y_i f_{x_i})}_{=: r} - K_{XX}^{-1} (\mathbf{f}_X - \mathbf{m}_X) \quad \text{with} \quad \frac{\partial \log \sigma(y_i f_{x_i})}{\partial f_{x_j}} = \delta_{ij} \left( \frac{y_i + 1}{2} - \sigma(f_{x_i}) \right)$$

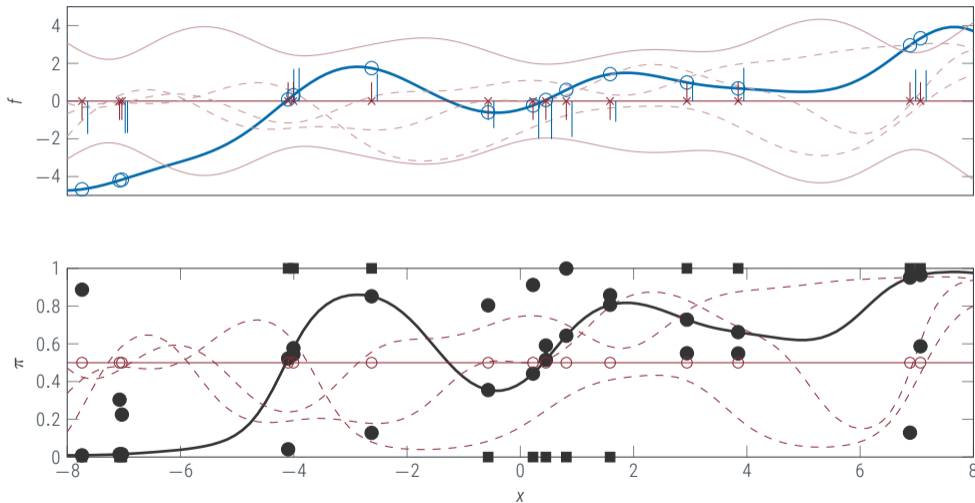
```

1 procedure GP-LOGISTIC-PREDICT( $\hat{\mathbf{f}}, W, R, r, k, \mathbf{x}$ )           //  $\hat{\mathbf{f}}, W, R = \text{Cholesky}(B), r$  handed over from training
2   for  $i = 1, \dots, \text{LENGTH}(\mathbf{x})$  do
3      $\bar{f}_i \leftarrow k_{x_i, X} r$                                // mean prediction (note at minimum,  $0 = \nabla p(\mathbf{f}_X | \mathbf{y}) = r - K_{XX}^{-1} (\mathbf{f}_X - \mathbf{m}_X)$ ).
4      $s \leftarrow R^{-1} (W^{1/2} k_{XX_i})$                    // pre-computation allows this step in  $\mathcal{O}(n^2)$ 
5      $v \leftarrow k_{x_i, x_i} - s^\top s$                        //  $v = \text{cov}(f_x)$ 
6      $\bar{\pi}_i \leftarrow \int \sigma(f_i) \mathcal{N}(f_i, \bar{f}_i, v) df_i$  // predictive probability for class 1 is  $p(y | \bar{\mathbf{f}}) = \int p(y_x | f_x) p(f_x | \bar{\mathbf{f}}) df_x$ 
7   end for
8   return  $\bar{\pi}_X$ 
9 end procedure
    
```

// entire loop is  $\mathcal{O}(n^2 m)$  for  $m$  test cases.

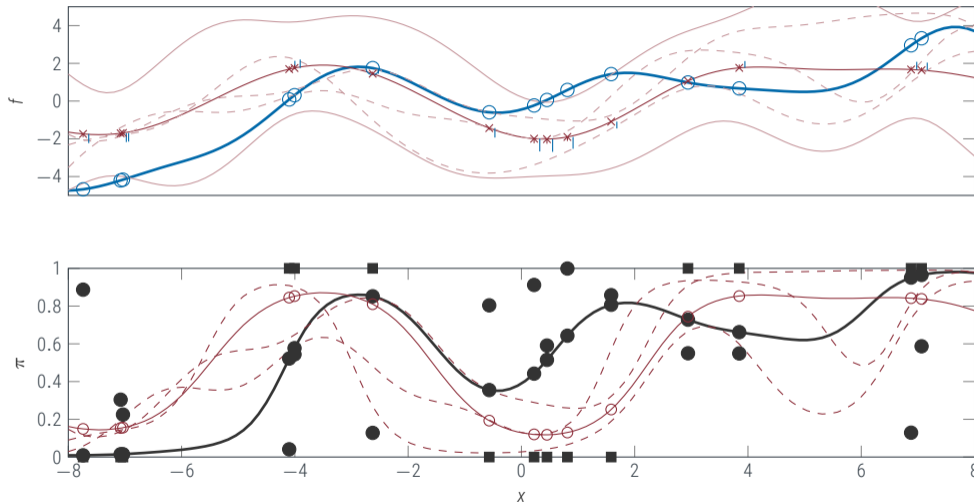
# Pictorial View

Gaussian process logistic regression




# Pictorial View

Gaussian process logistic regression



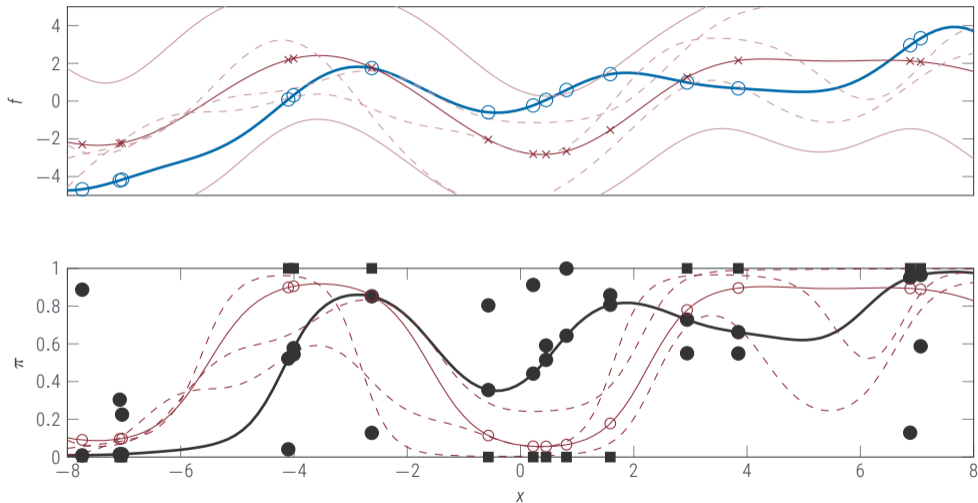
# Pictorial View

Gaussian process logistic regression



# Pictorial View

## Gaussian process logistic regression



## Gaussian Process Classification – (Probabilistic) Logistic Regression:

- ▶ Supervised classification phrased in a **discriminative** model with probabilistic interpretation
- ▶ model binary outputs as a **transformation** of a **latent function** with a Gaussian process prior
- ▶ due to **non-Gaussian likelihood**, the posterior is non-Gaussian; exact inference **intractable**
- ▶ **Laplace approximation**: Find MAP estimator, second order expansion for Gaussian approximation
- ▶ tune code for numerical stability, efficient computations
- ▶ Laplace approximation provides Gaussian posterior on training points, hence evidence, predictions

