Gaussian Processes

Seung-Hoon Na
Chonbuk National University

Gaussian Process Regression

- **Predictions using noisy observations**

$$\begin{pmatrix} \mathbf{y} \\ \mathbf{f}_* \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} \mathbf{K}_y & \mathbf{K}_* \\ \mathbf{K}_*^T & \mathbf{K}_{**} \end{pmatrix}\right)$$

$$\begin{aligned} p(\mathbf{f}_*|\mathbf{X}_*, \mathbf{X}, \mathbf{y}) &= \mathcal{N}(\mathbf{f}_*|\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*) \\ \boldsymbol{\mu}_* &= \mathbf{K}_*^T \mathbf{K}_y^{-1} \mathbf{y} \\ \boldsymbol{\Sigma}_* &= \mathbf{K}_{**} - \mathbf{K}_*^T \mathbf{K}_y^{-1} \mathbf{K}_* \end{aligned}$$

  - The case of a single test input:

$$p(f_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \mathcal{N}(f_*|\mathbf{k}_*^T \mathbf{K}_y^{-1} \mathbf{y}, k_{**} - \mathbf{k}_*^T \mathbf{K}_y^{-1} \mathbf{k}_*)$$

Gaussian Process Regression

- **Computational and numerical issues**
  - It is unwise to directly invert $\mathbf{K}_y$
  - Instead, we use a Cholesky decomposition $\mathbf{K}_y = \mathbf{L}\mathbf{L}^T$

Algorithm 15.1: GP regression

1 $\mathbf{L} = \text{cholesky}(\mathbf{K} + \sigma_y^2 \mathbf{I})$;
2 $\boldsymbol{\alpha} = \mathbf{L}^T \setminus (\mathbf{L} \setminus \mathbf{y})$;
3 $\mathbb{E}[f_*] = \mathbf{k}_*^T \boldsymbol{\alpha}$;
4 $\mathbf{v} = \mathbf{L} \setminus \mathbf{k}_*$;
5 $\text{var}[f_*] = \kappa(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{v}^T \mathbf{v}$;
6 $\log p(\mathbf{y}|\mathbf{X}) = -\frac{1}{2}\mathbf{y}^T \boldsymbol{\alpha} - \sum_i \log L_{ii} - \frac{N}{2}\log(2\pi)$

Marginal probability

Gaussian Process Regression

- $f(x_*) = k_{**} - \boldsymbol{k_*^T K_y^{-1} k_*}$
- $cholesky(\boldsymbol{K_y}) = \boldsymbol{L L^T}$
- $\boldsymbol{k_*^T K_y^{-1} k_* = (L^{-1} k_*)^T (L^{-1} k_*)}$

- $\boldsymbol{v = L^{-1} k_* = L \setminus k_*}$

Cholesky Decomposition

- The Cholesky decomposition(CD)
  - The CD of a symmetric, positive definite matrix A decomposes A into a product of a lower triangular matrix L and its transpose

$$LL^\top = A$$

- Solving linear system using CD:
  - To solve $A\mathbf{x} = \mathbf{b}$
  - We have two steps: $L\mathbf{y} = \mathbf{b}$ $L^\top \mathbf{x} = \mathbf{y}$

$$\mathbf{x} = L^\top \backslash (L \backslash \mathbf{b})$$

- Computing the determinant of a matrix

$$|A| = \prod_{i=1}^{n} L_{ii}^2, \quad \text{or} \quad \log|A| = 2\sum_{i=1}^{n} \log L_{ii}$$

Gaussian Process Classification

- The main difficulty is that the Gaussian prior is not conjugate to the bernoulli/ multinoulli likelihood ➔ several approximations are available
  - **Gaussian approximation**
  - Expectation propagation (Kuss and Rasmussen 2005; Nickisch and Rasmussen 2008)
  - Variational inference (Girolami and Rogers 2006; Opper and Archambeau 2009)
  - MCMC (Neal 1997; Christensen et al. 2006)

Gaussian Process Classification
binary classification

$$p(y_i|\mathbf{x}_i) = \sigma(y_i f(\mathbf{x}_i))$$

- Logistic regression: $\sigma(z) = \mathrm{sigm}(z)$
- Probit regression: $\sigma(z) = \Phi(z)$
- f: GP regression $f \sim \mathrm{GP}(0, \kappa)$

Gaussian Process Classification

- Define the log of the unnormalized posterior

$$\ell(\mathbf{f}) = \log p(\mathbf{y}|\mathbf{f}) + \log p(\mathbf{f}|\mathbf{X})$$

$$= \log p(\mathbf{y}|\mathbf{f}) - \frac{1}{2}\mathbf{f}^T\mathbf{K}^{-1}\mathbf{f} - \frac{1}{2}\log|\mathbf{K}| - \frac{N}{2}\log 2\pi$$

$$J(f) \triangleq -\ell(f)$$

$$\begin{aligned} \mathbf{g} &= -\nabla \log p(\mathbf{y}|\mathbf{f}) + \mathbf{K}^{-1}\mathbf{f} \\ \mathbf{H} &= -\nabla\nabla \log p(\mathbf{y}|\mathbf{f}) + \mathbf{K}^{-1} = \mathbf{W} + \mathbf{K}^{-1} \end{aligned}$$

$$\mathbf{W} \triangleq -\nabla\nabla \log p(\mathbf{y}|\mathbf{f})$$

Gaussian Process Classification

- Formula for $\nabla \log p(\mathbf{y}|\mathbf{f})$ $\nabla\nabla \log p(\mathbf{y}|\mathbf{f})$

$\log p(y_i\|f_i)$	$\frac{\partial}{\partial f_i} \log p(y_i\|f_i)$	$\frac{\partial^2}{\partial f_i^2} \log p(y_i\|f_i)$
$\log \mathrm{sigm}(y_i f_i)$	$t_i - \pi_i$	$-\pi_i(1-\pi_i)$
$\log \Phi(y_i f_i)$	$\frac{y_i \phi(f_i)}{\Phi(y_i f_i)}$	$-\frac{\phi_i^2}{\Phi(y_i f_i)^2} - \frac{y_i f_i \phi(f_i)}{\Phi(y_i f_i)}$

Gaussian Process Classification

- Use IRLS to find the MAP estimate

$$\begin{aligned}\mathbf{f}^{new} &= \mathbf{f} - \mathbf{H}^{-1}\mathbf{g} \\ &= \mathbf{f} + (\mathbf{K}^{-1} + \mathbf{W})^{-1}(\nabla \log p(\mathbf{y}|\mathbf{f}) - \mathbf{K}^{-1}\mathbf{f}) \\ &= (\mathbf{K}^{-1} + \mathbf{W})^{-1}(\mathbf{W}\mathbf{f} + \nabla \log p(\mathbf{y}|\mathbf{f}))\end{aligned}$$

- At convergence, the Gaussian approximation of the posterior:

$$p(\mathbf{f}|\mathbf{X}, \mathbf{y}) \approx \mathcal{N}(\hat{\mathbf{f}}, (\mathbf{K}^{-1} + \mathbf{W})^{-1})$$

Gaussian Process Classification

- **Computing the posterior predictive**
- The predictive mean:

$$\mathbb{E}\left[f_*|\mathbf{x}_*,\mathbf{X},\mathbf{y}\right] = \int \mathbb{E}\left[f_*|\mathbf{f},\mathbf{x}_*,\mathbf{X},\mathbf{y}\right] p(\mathbf{f}|\mathbf{X},\mathbf{y}) d\mathbf{f}$$

$$= \int \mathbf{k}_*^T \mathbf{K}^{-1} \mathbf{f}\; p(\mathbf{f}|\mathbf{X},\mathbf{y}) d\mathbf{f}$$

$$= \mathbf{k}_*^T \mathbf{K}^{-1} \mathbb{E}\left[\mathbf{f}|\mathbf{X},\mathbf{y}\right] \approx \mathbf{k}_*^T \mathbf{K}^{-1} \hat{\mathbf{f}}$$

Gaussian Process Classification

- The predictive variance:
  - Use the law of total variance

$$\text{var}\left[f_*\right] = \mathbb{E}\left[\text{var}\left[f_*|\mathbf{f}\right]\right] + \text{var}\left[\mathbb{E}\left[f_*|\mathbf{f}\right]\right]$$

$$\mathbb{E}\left[\text{var}\left[f_*|\mathbf{f}\right]\right] = \mathbb{E}\left[k_{**} - \mathbf{k}_*^T\mathbf{K}^{-1}\mathbf{k}_*\right]$$
$$= k_{**} - \mathbf{k}_*^T\mathbf{K}^{-1}\mathbf{k}_*$$

$$\text{var}\left[\mathbb{E}\left[f_*|\mathbf{f}\right]\right] = \text{var}\left[\mathbf{k}_*\mathbf{K}^{-1}\mathbf{f}\right]$$
$$= \mathbf{k}_*^T\mathbf{K}^{-1}\text{cov}\left[\mathbf{f}\right]\mathbf{K}^{-1}\mathbf{k}_*$$

https://www.macroeconomics.tu-berlin.de/fileadmin/fg124/financial_crises/exercise/Variances.pdf

Gaussian Process Classification

- The predictive variance:

$$\mathrm{var}\left[f_*\right] = k_{**} - \mathbf{k}_*^T(\mathbf{K}^{-1} - \mathbf{K}^{-1}\mathrm{cov}\left[\mathbf{f}\right]\mathbf{K}^{-1})\mathbf{k}_*$$

$$\mathrm{cov}\left[f\right] \approx (\mathbf{K}^{-1} + \mathbf{W})^{-1}$$

$$\mathrm{var}\left[f_*\right]$$

$$\approx k_{**} - \mathbf{k}_*^T\mathbf{K}^{-1}\mathbf{k}_* + \mathbf{k}_*^T\mathbf{K}^{-1}(\mathbf{K}^{-1} + \mathbf{W})^{-1}\mathbf{K}^{-1}\mathbf{k}_*$$

$$= k_{**} - \mathbf{k}_*^T(\mathbf{K} + \mathbf{W}^{-1})^{-1}\mathbf{k}_*$$

Matrix inversion lemma

Matrix inversion lemma

- *Consider a general partitioned matrix* $\mathbf{M} = \begin{pmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{G} & \mathbf{H} \end{pmatrix}$

where we assume $\boldsymbol{E}$ *and* $\boldsymbol{H}$ *are invertiable*

$$
\begin{aligned}
(\mathbf{E} - \mathbf{F}\mathbf{H}^{-1}\mathbf{G})^{-1} &= \mathbf{E}^{-1} + \mathbf{E}^{-1}\mathbf{F}(\mathbf{H} - \mathbf{G}\mathbf{E}^{-1}\mathbf{F})^{-1}\mathbf{G}\mathbf{E}^{-1} \\
(\mathbf{E} - \mathbf{F}\mathbf{H}^{-1}\mathbf{G})^{-1}\mathbf{F}\mathbf{H}^{-1} &= \mathbf{E}^{-1}\mathbf{F}(\mathbf{H} - \mathbf{G}\mathbf{E}^{-1}\mathbf{F})^{-1} \\
|\mathbf{E} - \mathbf{F}\mathbf{H}^{-1}\mathbf{G}| &= |\mathbf{H} - \mathbf{G}\mathbf{E}^{-1}\mathbf{F}||\mathbf{H}^{-1}||\mathbf{E}|
\end{aligned}
$$

Gaussian Process Classification

$$p(f_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \mathcal{N}(\mathbb{E}[f_*], \mathrm{var}[f_*])$$

- Convert to a predictive distribution for binary responses

$$\pi_* = p(y_* = 1|\mathbf{x}_*, \mathbf{X}, \mathbf{y})$$
$$\approx \int \sigma(f_*) p(f_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) df_*$$

- This can be approximated using
  - Monte Carlo approximation
  - Probit approximation
  - …

Gaussian Process Classification

- Marginal likelihood
  - Used to optimize the kernel parameters
  - Applying the Laplace approximation, we have:

$$\log p(\mathbf{y}|\mathbf{X}) \approx \ell(\hat{\mathbf{f}}) - \frac{1}{2}\log|\mathbf{H}| + \text{const}$$

$$\log p(\mathbf{y}|\mathbf{X}) \approx \log p(\mathbf{y}|\hat{\mathbf{f}}) - \frac{1}{2}\hat{\mathbf{f}}^T\mathbf{K}^{-1}\hat{\mathbf{f}} - \frac{1}{2}\log|\mathbf{K}| - \frac{1}{2}\log|\mathbf{K}^{-1} + \mathbf{W}|$$

- Computing the derivatives $\dfrac{\partial \log p(\mathbf{y}|\mathbf{X},\boldsymbol{\theta})}{\partial \theta_j}$
  - Now, since $\hat{\boldsymbol{f}}$, $\boldsymbol{W}$, as well as **K**, depend on $\boldsymbol{\theta}$
    - More complex than in the regression case

Gaussian Process Classification

- **Laplace approximation** to the marginal likelihood.

$$
\begin{aligned}
\hat{p}(\boldsymbol{\theta}|\mathcal{D}) &\approx \frac{1}{Z} e^{-E(\boldsymbol{\theta}^*)} \exp\left[-\frac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\theta}^*)^T \mathbf{H} (\boldsymbol{\theta}-\boldsymbol{\theta}^*)\right] \\
&= \mathcal{N}(\boldsymbol{\theta}|\boldsymbol{\theta}^*, \mathbf{H}^{-1}) \\
Z = p(\mathcal{D}) &\approx \int \hat{p}(\boldsymbol{\theta}|\mathcal{D}) d\boldsymbol{\theta} = e^{-E(\boldsymbol{\theta}^*)} (2\pi)^{D/2} |\mathbf{H}|^{-\frac{1}{2}}
\end{aligned}
$$

Gaussian Process Classification

- **Numerically stable computation**
  - To avoid inverting **K** or **W**, introduce using B:

$$\mathbf{B} = \mathbf{I}_N + \mathbf{W}^{\frac{1}{2}}\mathbf{K}\mathbf{W}^{\frac{1}{2}}$$

    - **B**: has eigenvalues bounded below by 1 and can be safely inverted
  - Applying the matrix inversion lemma, we have:

$$(\mathbf{K}^{-1} + \mathbf{W})^{-1} = \mathbf{K} - \mathbf{K}\mathbf{W}^{\frac{1}{2}}\mathbf{B}^{-1}\mathbf{W}^{\frac{1}{2}}\mathbf{K}$$

  - The IRLS update, now:

$$\begin{aligned} \mathbf{f}^{new} &= (\mathbf{K}^{-1} + \mathbf{W})^{-1} \underbrace{(\mathbf{W}\mathbf{f} + \nabla \log p(\mathbf{y}|\mathbf{f}))}_{\mathbf{b}} \\ &= \mathbf{K}(\mathbf{I} - \mathbf{W}^{\frac{1}{2}}\mathbf{B}^{-1}\mathbf{W}^{\frac{1}{2}}\mathbf{K})\mathbf{b} \\ &= \mathbf{K}\underbrace{(\mathbf{b} - \mathbf{W}^{\frac{1}{2}}\mathbf{L}^T \setminus (\mathbf{L} \setminus (\mathbf{W}^{\frac{1}{2}}\mathbf{K}\mathbf{b})))}_{\mathbf{a}} \end{aligned}$$

$\mathbf{B} = \mathbf{L}\mathbf{L}^T$ →

Gaussian Process Classification

- **Numerically stable computation**
  - At convergence, we have:

$$\mathbf{a} = \mathbf{K}^{-1}\hat{\mathbf{f}}$$

  - The log-marginal likelihood is:

$$\log p(\mathbf{y}|\mathbf{X}) = \log p(\mathbf{y}|\hat{\mathbf{f}}) - \frac{1}{2}\mathbf{a}^T\hat{\mathbf{f}} - \sum_i \log L_{ii}$$

  where we exploited:

$$|\mathbf{B}| = |\mathbf{K}||\mathbf{K}^{-1} + \mathbf{W}| = |\mathbf{I}_N + \mathbf{W}^{\frac{1}{2}}\mathbf{K}\mathbf{W}^{\frac{1}{2}}|$$

Gaussian Process Classification

- **Numerically stable computation**
  - Compute the predictive distr. $\mathbb{E}[f_*] = \mathbf{k}_*^T \mathbf{K}^{-1} \hat{\mathbf{f}}$
  - Here, at the mode, $\nabla \ell = 0$, so $\hat{\mathbf{f}} = \mathbf{K}(\nabla \log p(\mathbf{y}|\hat{\mathbf{f}}))$
  - Thus, the predictive mean:

$$\mathbb{E}[f_*] = \mathbf{k}_*^T \nabla \log p(\mathbf{y}|\hat{\mathbf{f}})$$

  - Also, we use:

$$(\mathbf{K} + \mathbf{W}^{-1})^{-1} = \mathbf{W}^{\frac{1}{2}} \mathbf{W}^{-\frac{1}{2}} (\mathbf{K} + \mathbf{W}^{-1})^{-1} \mathbf{W}^{-\frac{1}{2}} \mathbf{W}^{\frac{1}{2}}$$
$$= \mathbf{W}^{\frac{1}{2}} \mathbf{B}^{-1} \mathbf{W}^{\frac{1}{2}}$$

  - Thus, the predictive variance:

$$\text{var}[f_*] = k_{**} - \mathbf{k}_*^T \mathbf{W}^{\frac{1}{2}} (\mathbf{L}\mathbf{L}^T)^{-1} \mathbf{W}^{\frac{1}{2}} \mathbf{k}_* = k_{**} - \mathbf{v}^T \mathbf{v}$$

$$\mathbf{v} = \mathbf{L} \setminus (\mathbf{W}^{\frac{1}{2}} \mathbf{k}_*)$$

Gaussian Process Classification

Algorithm 15.2: GP binary classification using Gaussian approximation

// First compute MAP estimate using IRLS;
$\mathbf{f} = \mathbf{0}$;
repeat
  $\mathbf{W} = -\nabla\nabla \log p(\mathbf{y}|\mathbf{f})$;
  $\mathbf{B} = \mathbf{I}_N + \mathbf{W}^{\frac{1}{2}}\mathbf{K}\mathbf{W}^{\frac{1}{2}}$;
  $\mathbf{L} = \text{cholesky}(\mathbf{B})$;
  $\mathbf{b} = \mathbf{W}\mathbf{f} + \nabla \log p(\mathbf{y}|\mathbf{f})$;
  $\mathbf{a} = \mathbf{b} - \mathbf{W}^{\frac{1}{2}}\mathbf{L}^T \setminus (\mathbf{L} \setminus (\mathbf{W}^{\frac{1}{2}}\mathbf{K}\mathbf{b}))$;
  $\mathbf{f} = \mathbf{K}\mathbf{a}$;
until *converged*;
$\log p(\mathbf{y}|\mathbf{X}) = \log p(\mathbf{y}|\mathbf{f}) - \frac{1}{2}\mathbf{a}^T\mathbf{f} - \sum_i \log L_{ii}$;
// Now perform prediction ;
$\mathbb{E}[f_*] = \mathbf{k}_*^T \nabla \log p(\mathbf{y}|\mathbf{f})$;
$\mathbf{v} = \mathbf{L} \setminus (\mathbf{W}^{\frac{1}{2}}\mathbf{k}_*)$;
$\text{var}[f_*] = k_{**} - \mathbf{v}^T\mathbf{v}$;
$p(y_* = 1) = \int \text{sigm}(z)\mathcal{N}(z|\mathbb{E}[f_*], \text{var}[f_*])dz$;

Gaussian Process Classification

The posterior predictive probability for the red circle class generated by a GP with an SE kernel. Thick black line is the decision boundary if we threshold at a probability of 0.5.

– Manual parameters, short length scale.

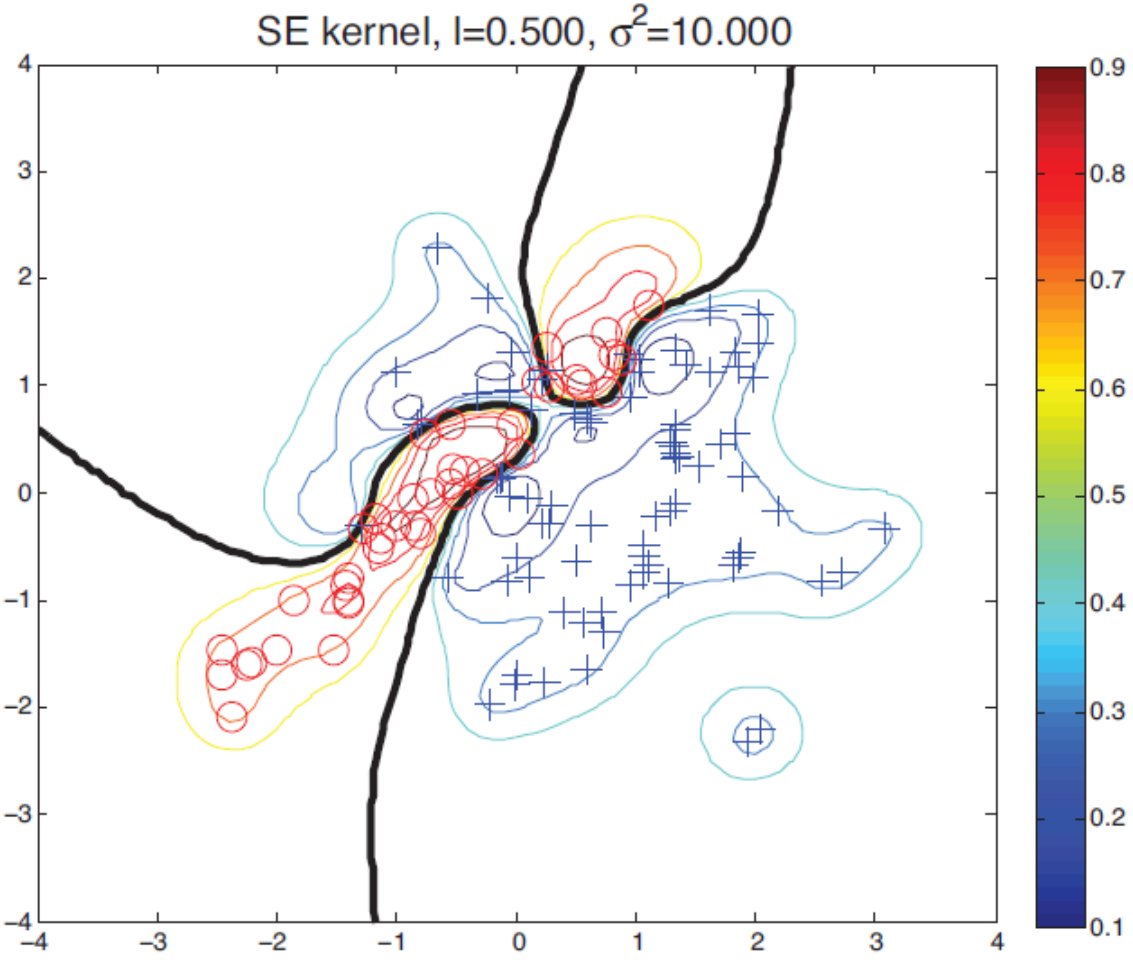

Gaussian Process Classification

- Learned parameters, long length scale

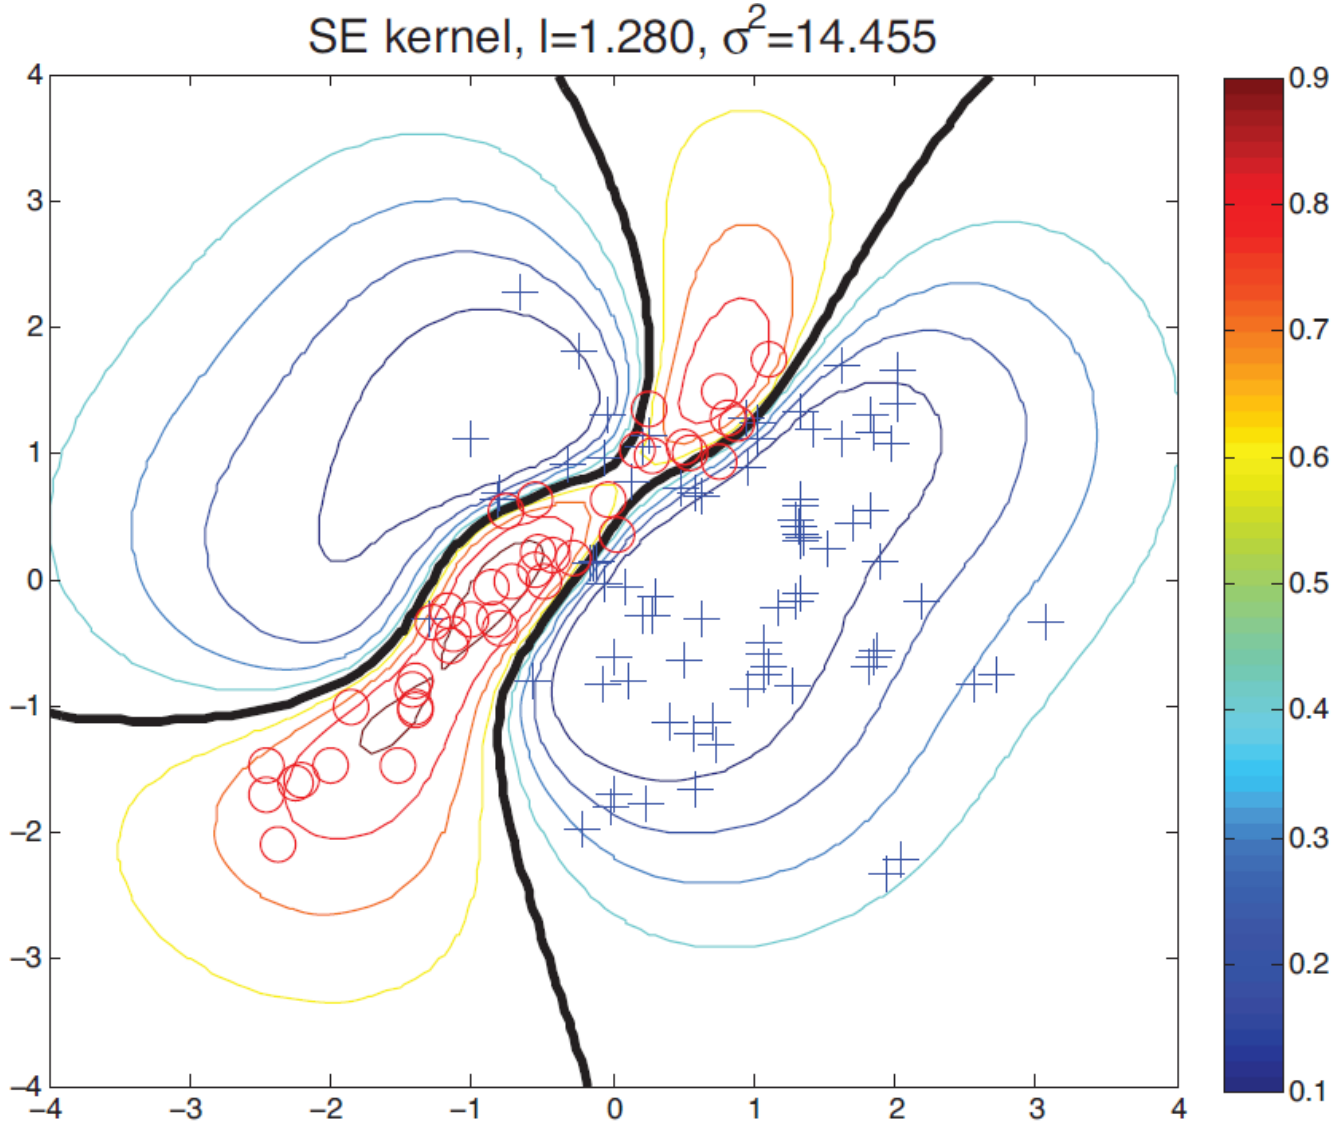

Gaussian Process Classification: Multi-class classification

$$p(y_i|\mathbf{x}_i) = \mathrm{Cat}(y_i|\mathcal{S}(\mathbf{f}_i))$$
$$\mathbf{f}_i = (f_{i1}, \ldots, f_{iC})$$
$$f_{.c} \sim \mathrm{GP}(0, \kappa_c)$$

– Again, we will use a Gaussian approximation to the posterior
  - 1) Use IRLS to compute the mode
  - 2) Apply the Gaussian approximation at the mode

Gaussian Process Classification: Multi-class classification

– The unnormalized log posterior:

$$\ell(\mathbf{f}) = -\frac{1}{2}\mathbf{f}^T\mathbf{K}^{-1}\mathbf{f} + \mathbf{y}^T\mathbf{f} - \sum_{i=1}^{N}\log\left(\sum_{c=1}^{C}\exp f_{ic}\right) - \frac{1}{2}\log|\mathbf{K}| - \frac{CN}{2}\log 2\pi$$

$$\mathbf{f} = (f_{11},\ldots,f_{N1},f_{12},\ldots,f_{N2},\cdots,f_{1C},\ldots,f_{NC})^T$$

– $\boldsymbol{y}$: a dummy encoding of y_i's with the same layout as $\boldsymbol{f}$

– $\boldsymbol{K}$: a block diagonal matrix containing $\boldsymbol{K}_c$

$$\mathbf{K}_c = \left[\kappa_c(\mathbf{x}_i,\mathbf{x}_j)\right]$$

Gaussian Process Classification: Multi-class classification

$$\nabla \ell = -\mathbf{K}^{-1}\mathbf{f} + \mathbf{y} - \boldsymbol{\pi}$$

$$\nabla\nabla \ell = -\mathbf{K}^{-1} - \mathbf{W}$$

$$\mathbf{W} \triangleq \mathrm{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}\boldsymbol{\Pi}^T$$

– Use IRLS to compute the mode

$$\mathbf{f}^{new} = (\mathbf{K}^{-1} + \mathbf{W})^{-1}(\mathbf{W}\mathbf{f} + \mathbf{y} - \boldsymbol{\pi})$$

Gaussian Process Classification: Multi-class classification

- The posterior predictive:

$$\mathbb{E}\left[f_{*c}\right] = \mathbf{k}_c(\mathbf{x}_*)^T \mathbf{K}_c^{-1} \hat{\mathbf{f}}_c = \mathbf{k}_c(\mathbf{x}_*)^T (\mathbf{y}_c - \hat{\boldsymbol{\pi}}_c)$$

$$\mathbb{E}\left[\mathbf{f}_*\right] = \mathbf{Q}_*{}^T \left(\mathbf{y} - \hat{\boldsymbol{\pi}}\right)$$

$$\mathbf{Q}_* = \begin{pmatrix} \mathbf{k}_1(\mathbf{x}_*) & \dots & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & \dots & \mathbf{k}_C(\mathbf{x}_*) \end{pmatrix}$$

Gaussian Process Classification: Multi-class classification

– The covariance of the latent response

$$\begin{aligned}\operatorname{cov}[\mathbf{f}_*] &= \boldsymbol{\Sigma} + \mathbf{Q}_*^T \mathbf{K}^{-1}(\mathbf{K}^{-1} + \mathbf{W})^{-1}\mathbf{K}^{-1}\mathbf{Q}_* \\ &= \operatorname{diag}(\mathbf{k}(\mathbf{x}_*, \mathbf{x}_*)) - \mathbf{Q}_*^T(\mathbf{K} + \mathbf{W}^{-1})^{-1}\mathbf{Q}_*\end{aligned}$$

$$\Sigma_{cc} = \kappa_c(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_c^T(\mathbf{x}_*)\mathbf{K}_c^{-1}\mathbf{k}_c(\mathbf{x}_*)$$

$$\mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) = [\kappa_c(\mathbf{x}_*, \mathbf{x}_*)]$$

- Computer the posterior predictive for the visible response:

$$p(y|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) \approx \int \operatorname{Cat}(y|\mathcal{S}(\mathbf{f}_*))\mathcal{N}(\mathbf{f}_*|\mathbb{E}[\mathbf{f}_*], \operatorname{cov}[\mathbf{f}_*])d\mathbf{f}_*$$

Gaussian Process Classification: Multi-class classification

- **Computing the marginal likelihood**
  - similar to the binary case

$$\log p(\mathbf{y}|\mathbf{X}) \approx -\frac{1}{2}\hat{\mathbf{f}}^T\mathbf{K}^{-1}\hat{\mathbf{f}} + \mathbf{y}^T\hat{\mathbf{f}} - \sum_{i=1}^{N}\log\left(\sum_{c=1}^{C}\exp\hat{f}_{ic}\right)$$

$$- \frac{1}{2}\log|\mathbf{I}_{CN} + \mathbf{W}^{\frac{1}{2}}\mathbf{K}\mathbf{W}^{\frac{1}{2}}|$$