Chapter 2 Probability distributions

Introduction to useful probability distributions < discrete / continuous

Discuss key statistical concepts such as Bayesian inference

Density estimation (find $p(x)$)

- iid assumption (independent and identically observations $x_1 \dots x_N$)
- unsupervised learning
- parametric vs non parametric

Conjugate distributions (prior and posterior)

2.1 Discrete random variables

concept, parameters

{ Bernoulli, binomial, beta } distribution

Bernoulli distribution

R.v $x \in \{0, 1\}$, parameter μ denote the prob. of $x = 1$

$$p(x = 1 \mid \mu) = \mu \qquad 0 \leq \mu \leq 1$$

prob.

parameter

$$\text{Bern}(x \mid \mu) = \mu^x (1 - \mu)^{1-x} \qquad (x \in \{0, 1\})$$

Remark

- $p(x=0 \mid \mu) + p(x=1 \mid \mu) = (1-\mu) + \mu = 1$
- $\mathbb{E}[x] = 0 \cdot p(x=0 \mid \mu) + 1 \cdot p(x=1 \mid \mu) = \mu$
- $\mathrm{Var}[x] = \mathbb{E}[x^2] - \mathbb{E}[x]^2 = \mu - \mu^2 = \mu(1-\mu)$

D $\qquad$ $\mathrm{Bern}(x \mid \mu)$

Assume $\{x_1, \ldots x_N\}$ is drawn independently from some Bernoulli

Find a parameter μ $(= p(x=1))$ in frequentist setting

See likelihood function (of μ) ($D = \{x_1, \ldots x_N\}$, $x_i \in \{0, 1\}$)

$$p(D|\mu) = \prod_{n=1}^{N} p(x_n|\mu) = \prod_{n=1}^{N} \mu^{x_n} (1-\mu)^{(1-x_n)} \qquad (2.5)$$

Estimate (find) a parameter μ by maximizing (2.5)

$$\ln p(D|\mu) = \sum_{n=1}^{N} \{ x_n \ln \mu + (1 - x_n) \ln (1-\mu) \}$$

$\ln p(D|\mu)$ only depends on x_n (Σx_n). Find a value for μ of

$$\frac{d}{d\mu} \ln p(D|\mu) = 0$$

$\Rightarrow \quad \mu_{ML} = \frac{1}{N} \sum_{n=1}^{N} x_n = \frac{m}{N}$ sample mean

where m = ※ of $x=1$

(likelihood 를 최댓값으로 만드는 μ)

The sample mean is an example of sufficient statistic.

E.g. Flip a coin 3 times, observe 

$\Rightarrow \quad \mu_{ML} = 1$

Binomial distribution 이항분포

Let $x = 0$ or 1 and N be ※ trials.

R.v $m \in \{0, 1, \dots, N\}$ to be ※ of $x = 1$.

From (2.5)

binomial distribution $\propto \mu^{m}(1-\mu)^{N-m}$

m

m = ※ of $x = 1$, μ : the probability of $x = 1$

To normalize prob. dist. calculate all of possible ※ of obtaining m, $x=1$. Denote by

$$\mathrm{Bin}(m \mid N, \mu) = \binom{N}{m} \mu^m (1-\mu)^{N-m}$$

where $\binom{N}{m} := \dfrac{N!}{(N-m)!\, m!} = {}_N C_m$ $\quad m = 0, 1, 2, \ldots N$

We have

$$\mathbb{E}[m] = \sum_{m=0}^{N} m\, \mathrm{Bin}(m \mid N, \mu)$$

(평균 $x=1$ 의 횟수) $$= \sum_{m=0}^{N} m \binom{N}{m} \mu^m (1-\mu)^{N-m} = N\mu$$

$$\mathrm{Var}[m] = \sum_{m=0}^{N} (m - \mathbb{E}[m])^2 \,\mathrm{Bin}(m \mid N, \mu) = N\mu(1-\mu)$$

2.1.1 The beta distribution

Conjugate distribution

$p(\theta|x)$ and $p(\theta)$ have the same distributions.

$p(\theta)$ is called a conjugate prior for $p(x|\theta)$

likelihood

x가 어떤 분포를 따를경우 가정

Goal parameter θ

⇓

posterior of θ ∝ likelihood function of θ × prior of θ

Recall the $p(\mathcal{D}|\mu)$ of Bernoulli distribution

likelihood

$$p(\mathcal{D}|\mu) = \prod_{n=1}^{N} \mu^{x_n} (1-\mu)^{(1-x_n)} \qquad x_n \in \{0, 1\}$$

$\mu^x (1-\mu)^{1-x}$ ($p(\mu)$ 에 기울 해야되기 때문에 μ 의 관점에서 보면)

To see Bayesian approach, we need to introduce $p(\mu)$.

beta distribution

normalization constant

$$\text{Beta}(\mu|a,b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \mu^{a-1} (1-\mu)^{b-1} \qquad 0 \le \mu \le 1 \qquad (2.13)$$

where $\Gamma(x) := \int_0^\infty u^{x-1} e^{-u} du$ is the gamma function (1.414)

$$\int_0^1 \text{Beta}(\mu|a,b)\, d\mu = 1$$

Remark

- $\int_0^1 \text{Beta}(\mu \mid a, b)\, d\mu = 1$ ✓

- $\mathbb{E}[\mu] = \frac{a}{a+b}$, $\text{Var}[\mu] = \frac{ab}{(a+b)^2(a+b+1)}$ Excercise.

- a, b are called hyperparameters

The posterior dist. of μ has the form as

$$p(\mu \mid m, l, a, b) \propto \mu^{m+a-1}(1-\mu)^{l+b-1}$$

where $l = N - m$.

m : ※ of $x_n = 1$

$\Rightarrow$ posterior

$$p(\mu \mid m, l, a, b) = \frac{\Gamma(m+a+l+b)}{\Gamma(m+a)\Gamma(l+b)} \mu^{m+a-1}(1-\mu)^{l+b-1} \quad (2.18)$$

(# of $x=1$, ※ $x=0$)

a, b : parameters of prior) → posterior

m, l : result of observation

In view of (2.18) and def of beta dist,

a, b can be interpreted as effective ※ observations of $x=1$ and $x=0$.

Sequential approach (Bayesian view point)

x=1 1개

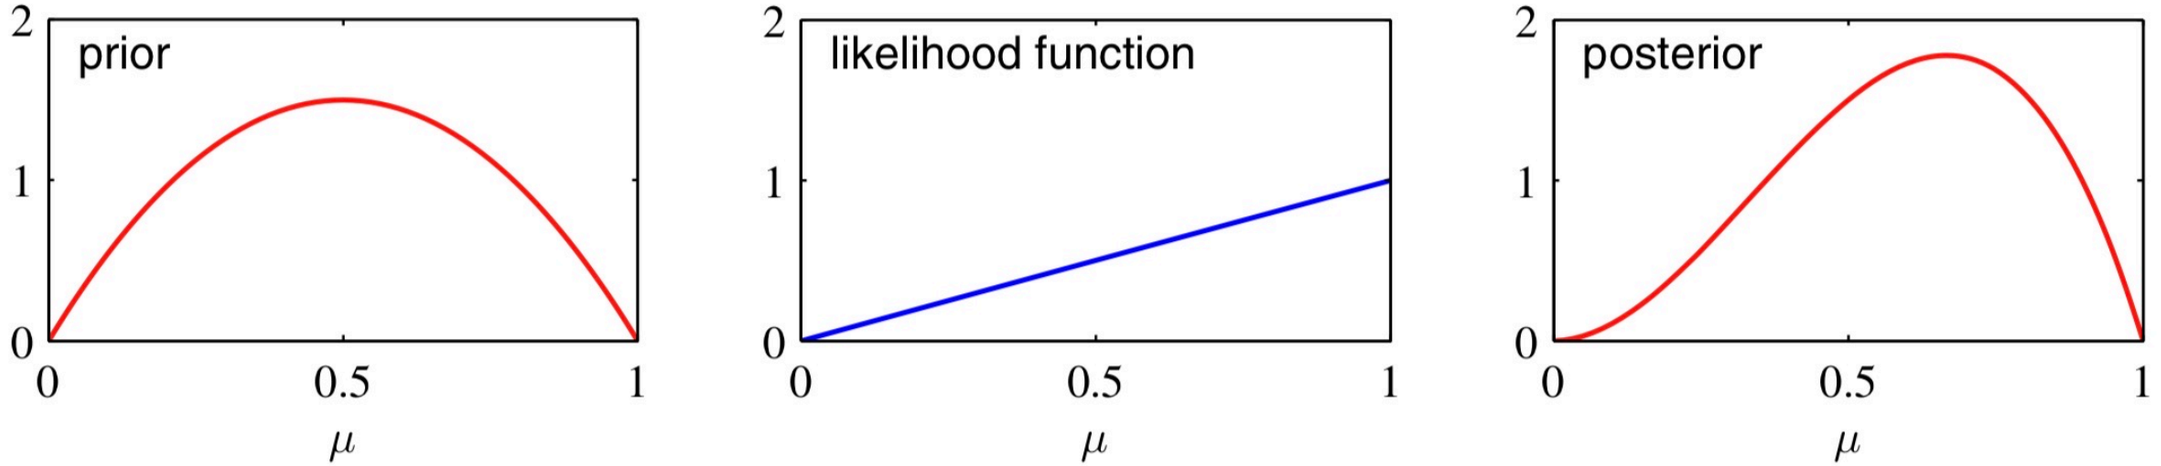

Figure 2.3 Illustration of one step of sequential Bayesian inference. The prior is given by a beta distribution with parameters $a = 2$, $b = 2$, and the likelihood function, given by (2.9) with $N = m = 1$, corresponds to a single observation of $x = 1$, so that the posterior is given by a beta distribution with parameters $a = 3$, $b = 2$.

Bern

$x_1 \dots x_N$

observation

μ

Let us predict the outcome of the next trial

$$p(x=1 \mid D) = \int_0^1 p(x=1, \mu \mid D)\, d\mu = \int_0^1 \frac{p(x=1, \mu, D)}{p(D)}\, d\mu$$

$$= \int_0^1 \frac{p(x=1, \mu, D)}{p(\mu, D)} \cdot \frac{p(\mu, D)}{p(D)}\, d\mu$$

$$= p(x=1 \mid \mu, D) \cdot p(\mu \mid D)$$

conditionally independent

$$= \int_0^1 p(x=1 \mid \mu) \cdot p(\mu \mid D)\, d\mu$$

$$= \int_0^1 \mu\, p(\mu \mid D)\, d\mu = E_\mu[\mu \mid D]$$

$$= \frac{m+a}{m+a+l+b}$$

posterior

$m+l=N$

$A \perp\!\!\!\perp B \mid C$

$p(A, B \mid C) = p(A \mid C) \cdot p(B \mid C)$

$\Leftrightarrow p(A \mid B, C) = p(A \mid C)$

$$p(x=1 \mid D) = \frac{m+a}{m+a+l+b} = \frac{m+a}{N+a+b} \qquad (l = N - m)$$

$\frac{m}{N}$

a, b are hyperparameters of prior.

m, l are from the result of experiment (m = # of $x=1$)

Remark

- $m, l \to \infty$ (huge observations), then ML $\approx$ Bayesian
- $a, b \to \infty$, then variance $\to 0$ (both prior and posterior)

2.2 Multinomial variables

Extend Bernoulli, binomial, beta distributions

Discrete variables that can take on one of K possible case

1-of-K scheme (one hot encoding)

x	orange	apple	grape	
orange	1	0	0	$(1,0,0)^T$
apple	0	1	0	$(0,1,0)^T$
apple	0	1	0	$(0,1,0)^T$
grape	0	0	1	$(0,0,1)^T$
orange	1	0	0	$(1,0,0)^T$
	μ_1	μ_2	μ_3	

K-dim. vector $\mathbf{x} = (x_1, x_2, \ldots x_K)^T$ as $(0, 0, \ldots 1, 0, 0 \ldots 0)^T$

$$\sum_{k=1}^{K} x_k = 1$$

$x \in \{0, 1\}$

Let μ_k be $p(x_k = 1)$. Bernoulli

The distribution of $\mathbf{x}$ (Categorical distribution)

$$p(\mathbf{x} \mid \boldsymbol{\mu}) = \prod_{k=1}^{K} \mu_k^{x_k}$$ (one observation)

where $\boldsymbol{\mu} := (\mu_1, \ldots \mu_K)^T$ with $\mu_k \geq 0$, $\sum_{k=1}^{K} \mu_k = 1$

Remark

- $\mathbf{x}$ can take K possible cases.
- $\sum_{\mathbf{x}} p(\mathbf{x} \mid \boldsymbol{\mu}) = \sum_{k=1}^{K} \mu_k = 1$
- $\mathbb{E}[\mathbf{x} \mid \boldsymbol{\mu}] = \sum_{\mathbf{x}} p(\mathbf{x} \mid \boldsymbol{\mu})\, \mathbf{x} = (\mu_1, \ldots \mu_K)^T = \boldsymbol{\mu}$

Consider D of N independent observations $x_1, \dots x_N$. sampled from categorical

Likelihood

$$p(D|\mu) = \prod_{n=1}^{N} \prod_{k=1}^{K} \mu_k^{x_{nk}} = \prod_{k=1}^{K} \mu_k^{(\sum_n x_{nk})} = \prod_{k=1}^{K} \mu_k^{m_k}$$

where $m_k = \sum_n x_{nk}$ (# of observations $x_k = 1$)

data matrix

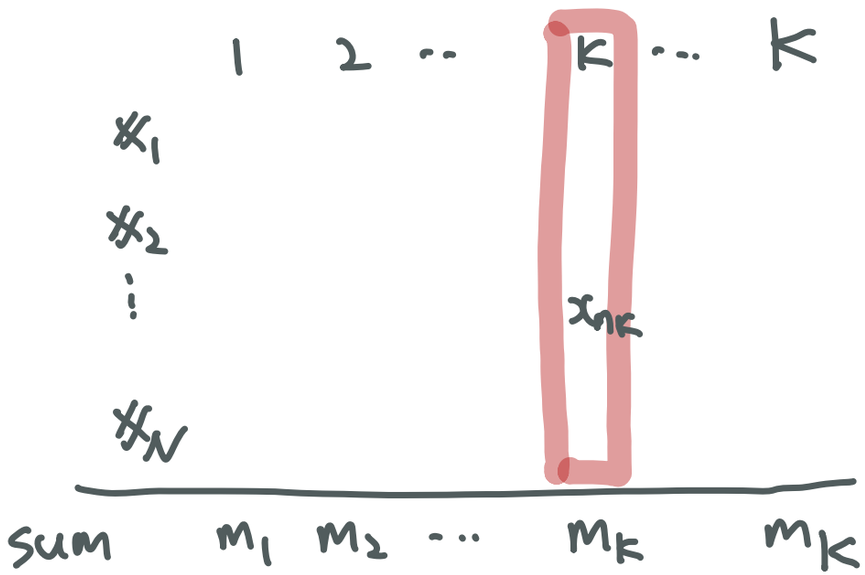

Maximize log likelihood $\ln p(\mathcal{D}|\boldsymbol{\mu})$ for $\boldsymbol{\mu}$ with constraint $\sum \mu_k = 1$. Using Lagrange multiplier and maximizing

$$\sum_{k=1}^{K} m_k \ln \mu_k + \lambda \left(\sum_{k=1}^{K} \mu_k - 1 \right) \qquad (2.31)$$

constraint

Set the derivative of (2.31) wrt μ_k to zero

$$\Rightarrow \quad \mu_k = \frac{m_k}{N}$$

$$\mu_k^{ML} = \frac{m_k}{N}$$

$$\boldsymbol{\mu}_{ML} = (\mu_1^{ML}, \ldots \mu_K^{ML})^T$$

Multinomial distribution $\quad \mathbf{m} := (m_1, \dots m_K)^T$

Joint distribution of the quantities $m_1, \dots m_K$ conditioned on μ and on N of N total observations.

$$\text{Mult}(m_1, m_2, \dots m_K \mid \mu, N) = \binom{N}{m_1\, m_2 \dots m_K} \prod_{k=1}^{K} \mu_k^{m_k} \qquad 0 \le \mu_k \le 1$$

$$\text{where } \binom{N}{m_1\, m_2 \dots m_K} := \frac{N!}{m_1!\, m_2! \dots m_K!} \qquad \sum_{k=1}^{K} m_k = N$$

K 개 categorical 변수를 갖는 N 개 자료 에서 각 K-class 가 m_K 씩 가질 확률 (μ 가 주어졌을 때)

2.2.1 Dirichlet distribution (multi-dim version of beta)

Consider the prior distributions for the parameters $\{\mu_k\}$ of multinomial distribution. (or categorical distribution)

Recall $\text{Mult}(m_1, \ldots m_K \mid \mu, N) \propto \prod_{k=1}^{K} \mu_k^{m_k}$

$$\Rightarrow \quad p(\mu \mid \alpha) \propto \prod_{k=1}^{K} \mu_k^{\alpha_k - 1}$$

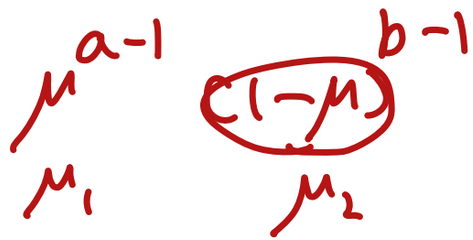

where $0 \le \mu_k \le 1$, $\sum_k \mu_k = 1$. Here $\alpha := (\alpha_1, \ldots \alpha_K)^T$ is the parameter

$\{\mu_k\}$ is confined to a simplex of dim $K-1$

Dirichlet distribution

$$\mathrm{Dir}(\boldsymbol{\mu}\mid\boldsymbol{\alpha}) := \frac{\Gamma(\alpha_0)}{\Gamma(\alpha_1)\cdots\Gamma(\alpha_K)}\prod_{k=1}^{K}\mu_k^{\alpha_k-1} \qquad 0\le\mu_k\le 1$$

where $\Gamma(x)$ is the Gamma function, $\alpha_0 := \sum \alpha_k$

So the posterior distribution for the parameters $\{\mu_k\}$

$$\underbrace{p(\boldsymbol{\mu}\mid\mathbf{D},\boldsymbol{\alpha})}_{\text{posterior}} \propto \underbrace{p(\mathbf{D}\mid\boldsymbol{\mu})}_{\text{likelihood}}\overbrace{p(\boldsymbol{\mu}\mid\boldsymbol{\alpha})}^{\text{prior}} \propto \prod_{k=1}^{K}\mu_k^{\alpha_k+m_k-1}$$

Posterior dist. again follow Dirichlet dist.

$$\Rightarrow \quad p(\mu \mid D, \alpha) = \mathrm{Dir}(\mu \mid \alpha + m) = \frac{\Gamma(\alpha_K + m_K)}{\Gamma(\alpha_1 + m_1) \cdots \Gamma(\alpha_K + m_K)} \prod_{k=1}^{K} \mu_k^{\alpha_k + m_k - 1}$$

prior: α — D: m — prior: α_k — observation: m_k

where $m := (m_1, \ldots m_K)^T$.

As binomial dist. with beta prior, we can interprete α_k of Dirichlet prior as an effective # of $x_k = 1$.

Fre

Bayesian

$t \sim \mathcal{N}(y(x, w), (\beta^{-1}))$

2.3 The Gaussian distribution (a.k.a. normal dist.)

Single real value $x \in \mathbb{R}$

$$\mathcal{N}(x \mid \mu, \sigma^2) := \frac{1}{(2\pi\sigma^2)^{1/2}} \exp\left\{-\frac{1}{2\sigma^2}(x-\mu)^2\right\}$$

where μ : mean, σ^2 : variance

D-dim vector $\mathbf{x} \in \mathbb{R}^D$

$$\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{D/2}} \frac{1}{|\boldsymbol{\Sigma}|^{1/2}} \exp\left\{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma} (\mathbf{x}-\boldsymbol{\mu})\right\}$$

$1 \times D$ $D \times D$ $D \times 1$

where $\boldsymbol{\mu}$: D-dim mean vector, $\boldsymbol{\Sigma}$: $D \times D$ covariance matrix

$|\boldsymbol{\Sigma}|$: determinant of $\boldsymbol{\Sigma}$

Remark: Why Gaussian is important ?

- Fits many natural phenomena
- Maximum entropy in continuous r.v
- Central limit theorem

Fix N. iid random samples of vector $x_1, x_2 \dots x_N$ are drawn from a population with μ, Σ.

$$\Rightarrow \quad \text{R.V.} \quad \bar{x}_N := \frac{1}{N}\sum_{n=1}^{N} x_n \approx \mathcal{N}\left(\mu, \frac{1}{N}\Sigma\right)$$

sample mean

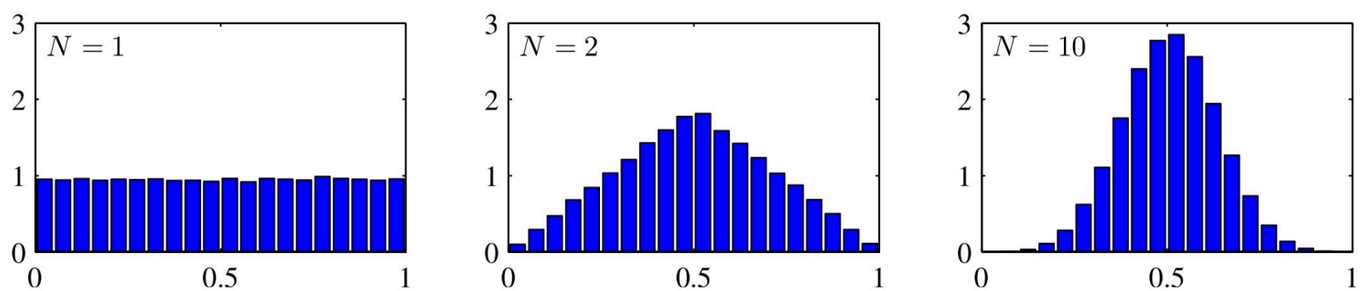

Figure 2.6 Histogram plots of the mean of N uniformly distributed numbers for various values of N. We observe that as N increases, the distribution tends towards a Gaussian.

Let us see geometrical form of Gaussian.

$$\Delta^2 := (x - \mu)^T \Sigma^{-1} (x - \mu) \qquad \mathcal{N}(x|\mu, \Sigma)$$

$1 \times D \quad D \times D \quad D \times 1$

Δ is called Mahalanobis distance from μ to x

WLOG. assume Σ is symmetric (and real).

Consider the eigenvector equation of Σ

$$\Sigma u_i = \lambda_i u_i \qquad i = 1, \ldots, D$$

Δ = constant 인 x의 집합

Eigenvalues $\lambda_1 \ldots \lambda_D$ are real and its eigenvectors can be chosen to form an orthonormal set, so that

$$\mathbf{u}_i^T \mathbf{u}_j = I_{ij}$$

where I_{ij} is the i,j element of the identity matrix.

By eigenvector decomposition, Σ can be expressed as

$$\Sigma = \sum_{i=1}^{D} \lambda_i \mathbf{u}_i \mathbf{u}_i^T$$

value; $D \times 1$ $1 \times D$; $D \times D$; vector

and

$$\Sigma^{-1} = \sum_{i=1}^{D} \frac{1}{\lambda_i} \mathbf{u}_i \mathbf{u}_i^T \quad \text{(inverse)}$$

$$\Rightarrow \quad \Delta^2 = (\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x}-\boldsymbol{\mu}) = (\mathbf{x}-\boldsymbol{\mu})^T \left(\sum_{i=1}^{D} \frac{1}{\lambda_i} \mathbf{u}_i \mathbf{u}_i^T \right) (\mathbf{x}-\boldsymbol{\mu})$$

$1 \times D \quad D \times D \quad D \times 1$

Mahalanobis distance

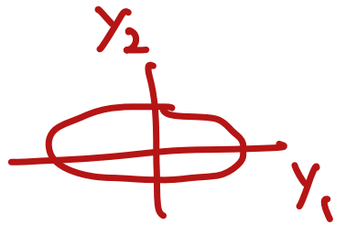

$$= \sum_{i=1}^{D} \frac{y_i^2}{\lambda_i} = C \text{ 인 } \mathbf{x} \text{의 집합}$$

where $y_i := \mathbf{u}_i^T (\mathbf{x}-\boldsymbol{\mu})$ (inner product of $\mathbf{u}_i$, $(\mathbf{x}-\boldsymbol{\mu})$)

$1 \times D \quad D \times 1$

$\mathbf{x}$

y_i : new coordinate system defined by $\mathbf{u}_i$, shifted and rotated

Let $\mathbf{y} := (y_1 \dots y_D)^T$. Then $\mathbf{y} = \mathbf{U}(\mathbf{x}-\boldsymbol{\mu})$ where

$$\mathbf{U} = \begin{pmatrix} \mathbf{u}_1^T \\ \vdots \\ \mathbf{u}_D^T \end{pmatrix}$$ whose rows are $\mathbf{u}_i^T$

(orthogonal matrix)

Figure 2.7 The red curve shows the elliptical surface of constant probability density for a Gaussian in a two-dimensional space $\mathbf{x} = (x_1, x_2)$ on which the density is $\exp(-1/2)$ of its value at $\mathbf{x} = \boldsymbol{\mu}$. The major axes of the ellipse are defined by the eigenvectors $\mathbf{u}_i$ of the covariance matrix, with corresponding eigenvalues λ_i.

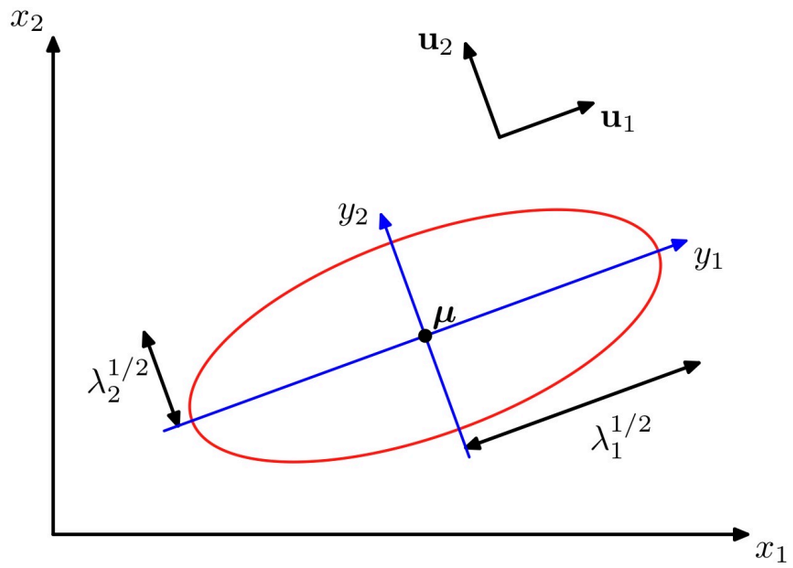

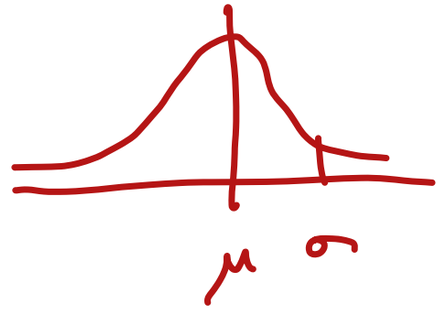

Remark

- If $\pi_i > 0 \quad \forall i = 1 \dots D$, contour surface of Δ is ellipsoid.
- center μ, axes oriented along u_i and scaling factors are given by $\lambda_i^{1/2}$

WLOG, assume all eigenvalues of Σ are strictly positive

Otherwise the distribution cannot be normalized (see ch 12)

i.e. Σ is assume to be positive definite.

$$y := U(x - \mu 1)$$

Now consider the Gaussian dist. in y_i coordinate system.

Jacobian matrix J with

(1.27) $x = g(y)$

$$P_y(y) = P_x(g(y))\,|g'(y)|$$

$$J_{ij} = \frac{\partial x_i}{\partial y_j} = U_{ji}$$

$$U^T y + \mu 1 = x$$

where U_{ji} are element of U^T

So $|J|^2 = |U^T|^2 = |U^T||U| = |U^T U| = |I| = 1$

and hence $|J| = \pm 1$. Also $|\Sigma|$ can be written as

$$|\Sigma| = \prod_{j=1}^{D} \lambda_j$$

$y = U(x - \mu 1)$

Thus in the y_j coordinate system,

$$p(y) = p(x)\,|\det(J)|^{\pm 1} = \prod_{j=1}^{D} \frac{1}{(2\pi\lambda_j)^{1/2}} \exp\left\{-\frac{y_j^2}{2\lambda_j}\right\}$$

$$= \prod_{j=1}^{D} p_{y_j}(y_j)$$

$p(y)$ is the product of D independent univariate Gaussian.

In (1.49), (1.51), we found univariate Gaussian dist has

$$\mathbb{E}[x] = \mu, \quad \text{var}[x] = \sigma^2$$

Now we will intepret parameters μ (D-dim), Σ (D×D).

$$\mathbb{E}[\mathbf{x}] = \frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} \int \exp\left\{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x}-\boldsymbol{\mu})\right\} \mathbf{x}\, d\mathbf{x}$$

$$= \frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} \int \exp\left\{-\frac{1}{2}\mathbf{z}^T \Sigma^{-1} \mathbf{z}\right\} (\mathbf{z}+\boldsymbol{\mu})\, d\mathbf{z}$$

where we have changed variables using $\mathbf{z} = \mathbf{x} - \boldsymbol{\mu}$.

Note that the exponent is even. So the term z in the factor $(z + \mu_1)$ will vanish.

$$\Rightarrow \quad \mathbb{E}[x] = \mu_1.$$

$x \in \mathbb{R}^D$

Now consider the second order moments of multivariate Gaussian. In univariate case, the second order moment is given by $\mathbb{E}[x^2]$. In multivariate Gaussian, there are D^2 second order moments given by $\mathbb{E}[x_i x_j]$.

$$x = \begin{pmatrix} x_1 \\ \vdots \\ x_D \end{pmatrix} \quad \mathbb{E}[x_i x_j]$$

$$\mathbb{E}[\mathbf{x}\mathbf{x}^T] = \frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} \int \exp\left\{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x}-\boldsymbol{\mu})\right\} \mathbf{x}\mathbf{x}^T d\mathbf{x}$$

$D\times 1 \quad 1\times D$

$$= \frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} \int \exp\left\{-\frac{1}{2}\mathbf{z}^T \Sigma^{-1} \mathbf{z}\right\} \underbrace{(\mathbf{z}+\boldsymbol{\mu})(\mathbf{z}+\boldsymbol{\mu})^T}_{(*)} d\mathbf{z}$$

$\mathbf{z}\mathbf{z}^T$

$\mathbf{z} := \mathbf{x} - \boldsymbol{\mu}$.

$$(*) \quad (\mathbf{z}+\boldsymbol{\mu})(\mathbf{z}^T+\boldsymbol{\mu}^T) = \mathbf{z}\mathbf{z}^T + \mathbf{z}\boldsymbol{\mu}^T + \boldsymbol{\mu}\mathbf{z}^T + \boldsymbol{\mu}\boldsymbol{\mu}^T$$

$\mathbf{z}\mathbf{z}^T$: vanish by symmetry; $\boldsymbol{\mu}\boldsymbol{\mu}^T$: constant

Recall $\mathbf{y} := U(\mathbf{x}-\boldsymbol{\mu})$, rows of U are eigenvectors of Σ

(U: transpose of)

$$\Rightarrow \quad \mathbf{z} = U^T\mathbf{y} = (\mathbf{u}_1 \dots \mathbf{u}_D)\begin{pmatrix} y_1 \\ \vdots \\ y_D \end{pmatrix} = \sum_{j=1}^{D} y_j \mathbf{u}_j$$

$$\frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} \int \exp\left\{-\frac{1}{2} z^T \Sigma^{-1} z\right\} z z^T \, dz$$

$$= \frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} \sum_{i=1}^{D} \sum_{j=1}^{D} u_i u_j^T \int \exp\left\{-\sum_{k=1}^{D} \frac{y_k^2}{2\lambda_k}\right\} y_i y_j \, dy$$

$$= \sum_{i=1}^{D} u_i u_i^T \lambda_i = \Sigma$$

we have used $|\Sigma| = \prod_{i=1}^{D} \lambda_i$ and $\frac{1}{(2\pi)^{1/2}} \frac{1}{\lambda^{1/2}} \exp\left\{-\frac{y^2}{2\lambda}\right\} \sim N(0, \lambda)$

e.g. $D = 2$

$$\sum_i \sum_j \iint \exp\left(-\frac{y_1^2}{2\lambda_1}\right) \exp\left(-\frac{y_2^2}{2\lambda_2}\right) y_i y_j \, dy_1 dy_2$$

will vanish when $i \neq j$

Thus we obtain $\mathbb{E}[\mathbf{x}\mathbf{x}^T] = \boldsymbol{\mu}\boldsymbol{\mu}^T + \boldsymbol{\Sigma}$ ($D \times D$ matrix)

and covariance of $\mathbf{x}$ can be obtained by

$$\begin{aligned}\mathrm{cov}[\mathbf{x}] &= \mathbb{E}[(\mathbf{x} - \mathbb{E}[\mathbf{x}])(\mathbf{x} - \mathbb{E}[\mathbf{x}])^T] \\ &= \boldsymbol{\Sigma}\end{aligned}$$

Remark

D DxD sym. real

$\mathcal{N}(x \mid \mu, \Sigma)$

- # of parameters : $\frac{D(D+3)}{2}$ quadratic

- $\Sigma = \text{diag}(\sigma_i^2)$ or $\Sigma = \sigma^2 I$

  2D D+1

- Unimodal (single maximum)

deep dive into Gaussian

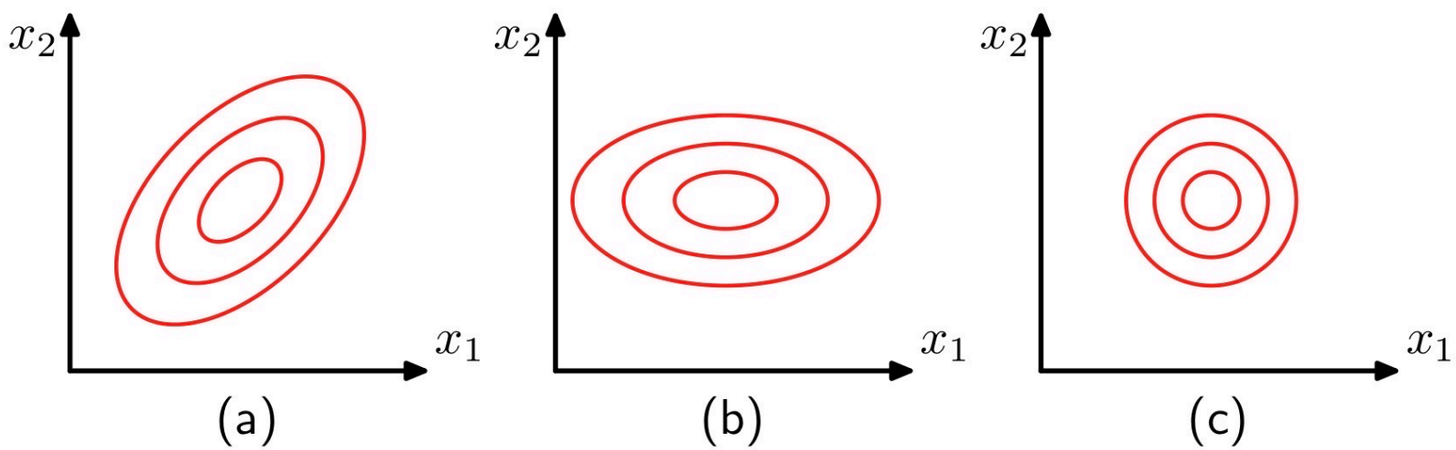

Figure 2.8 Contours of constant probability density for a Gaussian distribution in two dimensions in which the covariance matrix is (a) of general form, (b) diagonal, in which the elliptical contours are aligned with the coordinate axes, and (c) proportional to the identity matrix, in which the contours are concentric circles.

2.3.1 Conditional Gaussian distributions

$\mathbf{x}$: D-dimensional vector with $\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$ which is partitioned into $\mathbf{x}_a$, $\mathbf{x}_b$ with $\mathbf{x}_a \in \mathbb{R}^M$, $\mathbf{x}_b \in \mathbb{R}^{D-M}$

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_a \\ \mathbf{x}_b \end{pmatrix}$$

corresponding partitions of mean vector, covariance matrix

$$\boldsymbol{\mu} = \begin{pmatrix} \boldsymbol{\mu}_a \\ \boldsymbol{\mu}_b \end{pmatrix} \qquad \boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{aa} & \boldsymbol{\Sigma}_{ab} \\ \boldsymbol{\Sigma}_{ba} & \boldsymbol{\Sigma}_{bb} \end{pmatrix}$$

$\boldsymbol{\Sigma}_{aa}$: $M \times M$, $\boldsymbol{\Sigma}_{ab}$: $M \times (D-M)$, $\boldsymbol{\Sigma}_{ba}$: $(D-M) \times M$, $\boldsymbol{\Sigma}_{bb}$: $(D-M) \times (D-M)$

Note that $\boldsymbol{\Sigma}^T = \boldsymbol{\Sigma}$ implies $\boldsymbol{\Sigma}_{aa}$, $\boldsymbol{\Sigma}_{bb}$ are symmetric, $\boldsymbol{\Sigma}_{ba} = \boldsymbol{\Sigma}_{ab}^T$

Let $\Lambda := \Sigma^{-1}$. Inverse of covariance matrix, precision matrix

$$\Lambda = \begin{pmatrix} \Lambda_{aa} & \Lambda_{ab} \\ \Lambda_{ba} & \Lambda_{bb} \end{pmatrix}$$

$M \times M$ (top-left block), $D-M \times D-M$ (bottom-right block)

Note that $\Lambda_{aa}, \Lambda_{bb}$ are symmetric, $\Lambda_{ba} = \Lambda_{ab}^{T}$.

$\Lambda_{aa} \neq \Sigma_{aa}^{-1}$

Find conditional distribution $p(x_a | x_b)$. Fix x_b.

Consider the quadratic form in the exponent.

$$-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu}) \qquad 1\times D \quad D\times D \quad D\times 1 \qquad (\mathbf{x}_b \text{ fixed})$$

$\mathbf{x}_a, \mathbf{x}_b$

$$= -\frac{1}{2}(\mathbf{x}_a-\boldsymbol{\mu}_a)^T\boldsymbol{\Lambda}_{aa}(\mathbf{x}_a-\boldsymbol{\mu}_a) - \frac{1}{2}(\mathbf{x}_a-\boldsymbol{\mu}_a)^T\boldsymbol{\Lambda}_{ab}(\mathbf{x}_b-\boldsymbol{\mu}_b) \qquad (2.70)$$

$$-\frac{1}{2}(\mathbf{x}_b-\boldsymbol{\mu}_b)^T\boldsymbol{\Lambda}_{ba}(\mathbf{x}_a-\boldsymbol{\mu}_a) - \frac{1}{2}(\mathbf{x}_b-\boldsymbol{\mu}_b)^T\boldsymbol{\Lambda}_{bb}(\mathbf{x}_b-\boldsymbol{\mu}_b)$$

First, $p(\mathbf{x}_a|\mathbf{x}_b)$ will be M-dim Gaussian, because density function is a quadratic form of exponent.

Now we are going to find its mean vector and covariance by "completeting the square".

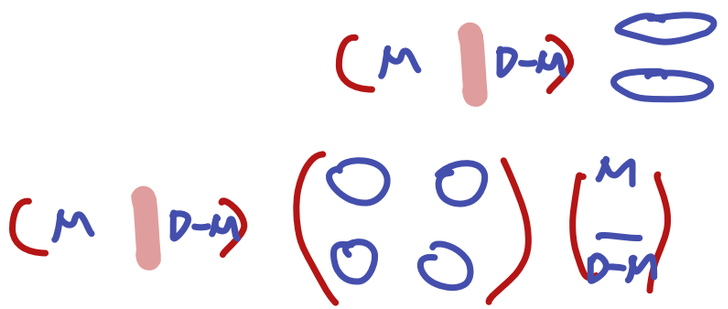

E.g. in case $p(x) \propto \exp(ax^2 + bx + c)$ → x ≒ gaussian

① ⇒

$$p(x) \propto \exp\left\{ a\left(x^2 + \frac{b}{a}x + \frac{b^2}{4a^2}\right) - \frac{b^2}{4a} + c \right\}$$

$$\propto \exp\left\{ a\left(x + \frac{b}{2a}\right)^2 \right\}$$

$$\Rightarrow \quad p(x) = \mathcal{N}(x \mid \mu, \sigma^2) \qquad \mu = -\frac{b}{2a}, \quad \sigma^2 = -\frac{1}{2a}$$

② Since $\mathcal{N}(x \mid \mu, \sigma^2) \propto \exp\left\{ -\frac{1}{2\sigma^2}(x^2 - 2\mu x + \mu^2) \right\}$

$$\propto \exp\left\{ -\frac{1}{2\sigma^2}x^2 + \frac{\mu}{\sigma^2}x \right\}$$

$$\Rightarrow \quad a = -\frac{1}{2\sigma^2}, \quad b = \frac{\mu}{\sigma^2} \qquad \left(\mu = -\frac{b}{2a}, \quad \sigma^2 = -\frac{1}{2a} \right)$$

Likewise, the exponent in D-dim Gaussian can be written

$$-\frac{1}{2}(x-\mu_1)^T \Sigma^{-1}(x-\mu_1) = -\frac{1}{2}x^T\Sigma^{-1}x + x^T\Sigma^{-1}\mu_1 + \text{constant}$$

In view of (2.70), (second order in x_a)

$$-\frac{1}{2}x_a^T \Lambda_{aa} x_a$$

So we obtain covariance of $p(x_a|x_b)$ is given by

$$\Sigma_{a|b} := \Lambda_{aa}^{-1} \quad (\neq \Sigma_{aa})$$

$\mu_{a|b}$?

Now consider all of the terms in (2.70) that are linear in x_a

$$x_a^T \{ \Lambda_{aa} \mu_a - \Lambda_{ab} (x_b - \mu_b) \}$$

where we have used $\Lambda_{ba}^T = \Lambda_{ab}$.

Since $$\Sigma_{a|b}^{-1} \mu_{a|b} = \Lambda_{aa} \mu_a - \Lambda_{ab} (x_b - \mu_b),$$

$$\mu_{a|b} = \Sigma_{a|b} \{ \Lambda_{aa} \mu_a - \Lambda_{ab} (x_b - \mu_b) \}$$

$$= \mu_a - \Lambda_{aa}^{-1} \Lambda_{ab} (x_b - \mu_b) .$$

$$\therefore \quad \mu_{a|b} = \mu_a - \Lambda_{aa}^{-1}\Lambda_{ab}(x_b - \mu_b), \qquad \Sigma_{a|b} = \Lambda_{aa}^{-1}$$

Let us find Λ_{aa} and Λ_{ab}.

Recall $\Lambda = \Sigma^{-1}$,

$$\begin{pmatrix} \Sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ba} & \Sigma_{bb} \end{pmatrix}^{-1} = \begin{pmatrix} \Lambda_{aa} & \Lambda_{ab} \\ \Lambda_{ba} & \Lambda_{bb} \end{pmatrix}$$

Use the following identification for the inverse of a partitioned matrix

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}^{-1} = \begin{pmatrix} M & -MBD^{-1} \\ -D^{-1}CM & D^{-1} + CMB^{-1} \end{pmatrix} \qquad (2.76)$$

where $M := (A - BD^{-1}C)^{-1}$

So we have

$$\Lambda_{aa} = (\Sigma_{aa} - \Sigma_{ab}\Sigma_{bb}^{-1}\Sigma_{ba})^{-1}$$

$$\Lambda_{ba} = -(\Sigma_{aa} - \Sigma_{ab}\Sigma_{bb}^{-1}\Sigma_{ba})^{-1}\Sigma_{ab}\Sigma_{bb}^{-1}$$

and hence

$$\begin{cases} \mu_{a|b} = \mu_a + \Sigma_{ab}\Sigma_{bb}^{-1}(x_b - \mu_b) \\ \Sigma_{a|b} = \Sigma_{aa} - \Sigma_{ab}\Sigma_{bb}^{-1}\Sigma_{ba} \end{cases}$$

x_b fix

Remark

- $\mu_{a|b}$ is a linear function of x_b
- $\Sigma_{a|b}$ is independent of x_b

2.3.2 Marginal Gaussian distributions

Consider the following marginal distribution $\mathcal{N}(x \mid \mu, \Sigma)$

$$p(x_a) = \int p(x_a, x_b)\, dx_b \qquad x = \begin{pmatrix} x_a \\ x_b \end{pmatrix} \begin{matrix} M \\ D-M \end{matrix}$$

(x_b를 적분해서 남은 x_a가 어떤 분포인지?)

Strategy: focus on quadratic form of exponent and identify the mean vector and covariance matrix of $p(x_a)$

Recall (2.70)

$$-\frac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)$$

$$= -\frac{1}{2}(x_a-\mu_a)^T \Lambda_{aa}(x_a-\mu_a) - \frac{1}{2}(x_a-\mu_a)^T \Lambda_{ab}(x_b-\mu_b)$$

$$-\frac{1}{2}(x_b-\mu_b)^T \Lambda_{ba}(x_a-\mu_a) - \frac{1}{2}(x_b-\mu_b)^T \Lambda_{bb}(x_b-\mu_b)$$

In order to integrate out x_b, pick out those terms involving x_b

$$-\frac{1}{2} x_b^T \Lambda_{bb} x_b + x_b^T m_1 = -\frac{1}{2}(x_b - \Lambda_{bb}^{-1} m_1)^T \Lambda_{bb} (x_b - \Lambda_{bb}^{-1} m_1)$$

(2.84)

$$+ \underbrace{\frac{1}{2} m_1^T \Lambda_{bb}^{-1} m_1}_{\text{indep. of } x_b} \qquad \text{(square expression)}$$

where $m_1 := \Lambda_{bb} \mu_{1b} - \Lambda_{ba}(x_a - \mu_{1a})$

For $p(x_a) = \int p(x_a, x_b)\, dx_b$,

(2.86) $$\int \exp\left\{ -\frac{1}{2}(x_b - \Lambda_{bb}^{-1} m_1)^T \Lambda_{bb} (x_b - \Lambda_{bb}^{-1} m_1) \right\} dx_b$$

which is an inverse of the normalization coefficient.

As seen before, this coefficient is independent of mean.

Combining the last term $(\frac{1}{2} m^T \Lambda_{bb}^{-1} m)$ in (2.84) with remaining terms in (2.70) depending on x_a, we obtain

$$-\frac{1}{2} m^T \Lambda_{bb}^{-1} m - \frac{1}{2} x_a^T \Lambda_{aa} x_a + x_a^T (\Lambda_{aa} \mu_a + \Lambda_{ab} \mu_b) + \text{constant}$$

$$= -\frac{1}{2} [\Lambda_{bb} \mu_b - \Lambda_{ba} (x_a - \mu_a)]^T \Lambda_{bb}^{-1} [\Lambda_{bb} \mu_b - \Lambda_{ba} (x_a - \mu_a)]$$

$$+ x_a^T (\Lambda_{aa} \mu_a + \Lambda_{ab} \mu_b) + \text{constant}$$

$$= -\frac{1}{2} x_a^T (\Lambda_{aa} - \Lambda_{ab} \Lambda_{bb}^{-1} \Lambda_{ba}) x_a + x_a^T (\Lambda_{aa} - \Lambda_{ab} \Lambda_{bb}^{-1} \Lambda_{ba}) \mu_a$$

$$+ \text{constant}$$

Recall the exponent in D-dim Gaussian can be written

$$-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu) = -\frac{1}{2}x^T\Sigma^{-1}x + x^T\Sigma^{-1}\mu + \text{constant}$$

Denote the covariance of $p(x_a)$ by Σ_a and Σ_a is given by

$$\Sigma_a = \underbrace{(\Lambda_{aa} - \Lambda_{ab}\Lambda_{bb}^{-1}\Lambda_{ba})}_{(\Sigma_a)^{-1}}{}^{-1}$$

Similarly, mean vector is given by

$$\Sigma_a \underbrace{(\Lambda_{aa} - \Lambda_{ab}\Lambda_{bb}^{-1}\Lambda_{ba})}_{\Sigma_a^{-1}} \mu_a = \mu_a$$

To simplify $\Sigma_a = (\Lambda_{aa} - \Lambda_{ab}\Lambda_{bb}^{-1}\Lambda_{ba})^{-1}$,

recall $$\begin{pmatrix} \Lambda_{aa} & \Lambda_{ab} \\ \Lambda_{ba} & \Lambda_{bb} \end{pmatrix}^{-1} = \begin{pmatrix} \Sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ba} & \Sigma_{bb} \end{pmatrix}$$

and use (2.76) expression of the inverse of a partitioned matrix

$$\Sigma_a = (\Lambda_{aa} - \Lambda_{ab}\Lambda_{bb}^{-1}\Lambda_{ba})^{-1} = \Sigma_{aa}$$

Thus we have $\mathbb{E}[\mathbf{x}_a] = \boldsymbol{\mu}_a$, $\mathrm{cov}[\mathbf{x}_a] = \Sigma_{aa}$

where $p(\mathbf{x}_a) = \int p(\mathbf{x}_a, \mathbf{x}_b)\, d\mathbf{x}_b$

$\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Lambda} := \boldsymbol{\Sigma}^{-1}$ D-dim $\mathbf{x}$

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_a \\ \mathbf{x}_b \end{pmatrix} \quad \boldsymbol{\mu} = \begin{pmatrix} \boldsymbol{\mu}_a \\ \boldsymbol{\mu}_b \end{pmatrix} \quad \boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{aa} & \boldsymbol{\Sigma}_{ab} \\ \boldsymbol{\Sigma}_{ba} & \boldsymbol{\Sigma}_{bb} \end{pmatrix} \quad \boldsymbol{\Lambda} = \begin{pmatrix} \boldsymbol{\Lambda}_{aa} & \boldsymbol{\Lambda}_{ab} \\ \boldsymbol{\Lambda}_{ba} & \boldsymbol{\Lambda}_{bb} \end{pmatrix}$$

Conditioned distribution

$$p(\mathbf{x}_a \mid \mathbf{x}_b) = \mathcal{N}(\mathbf{x}_a \mid \boldsymbol{\mu}_{a|b}, \boldsymbol{\Lambda}_{aa}^{-1})$$

$$\boldsymbol{\mu}_{a|b} = \boldsymbol{\mu}_a - \boldsymbol{\Lambda}_{aa}^{-1} \boldsymbol{\Lambda}_{ab} (\mathbf{x}_b - \boldsymbol{\mu}_b)$$

Marginal distribution

$$p(\mathbf{x}_a) = \mathcal{N}(\mathbf{x}_a \mid \boldsymbol{\mu}_a, \boldsymbol{\Sigma}_{aa})$$

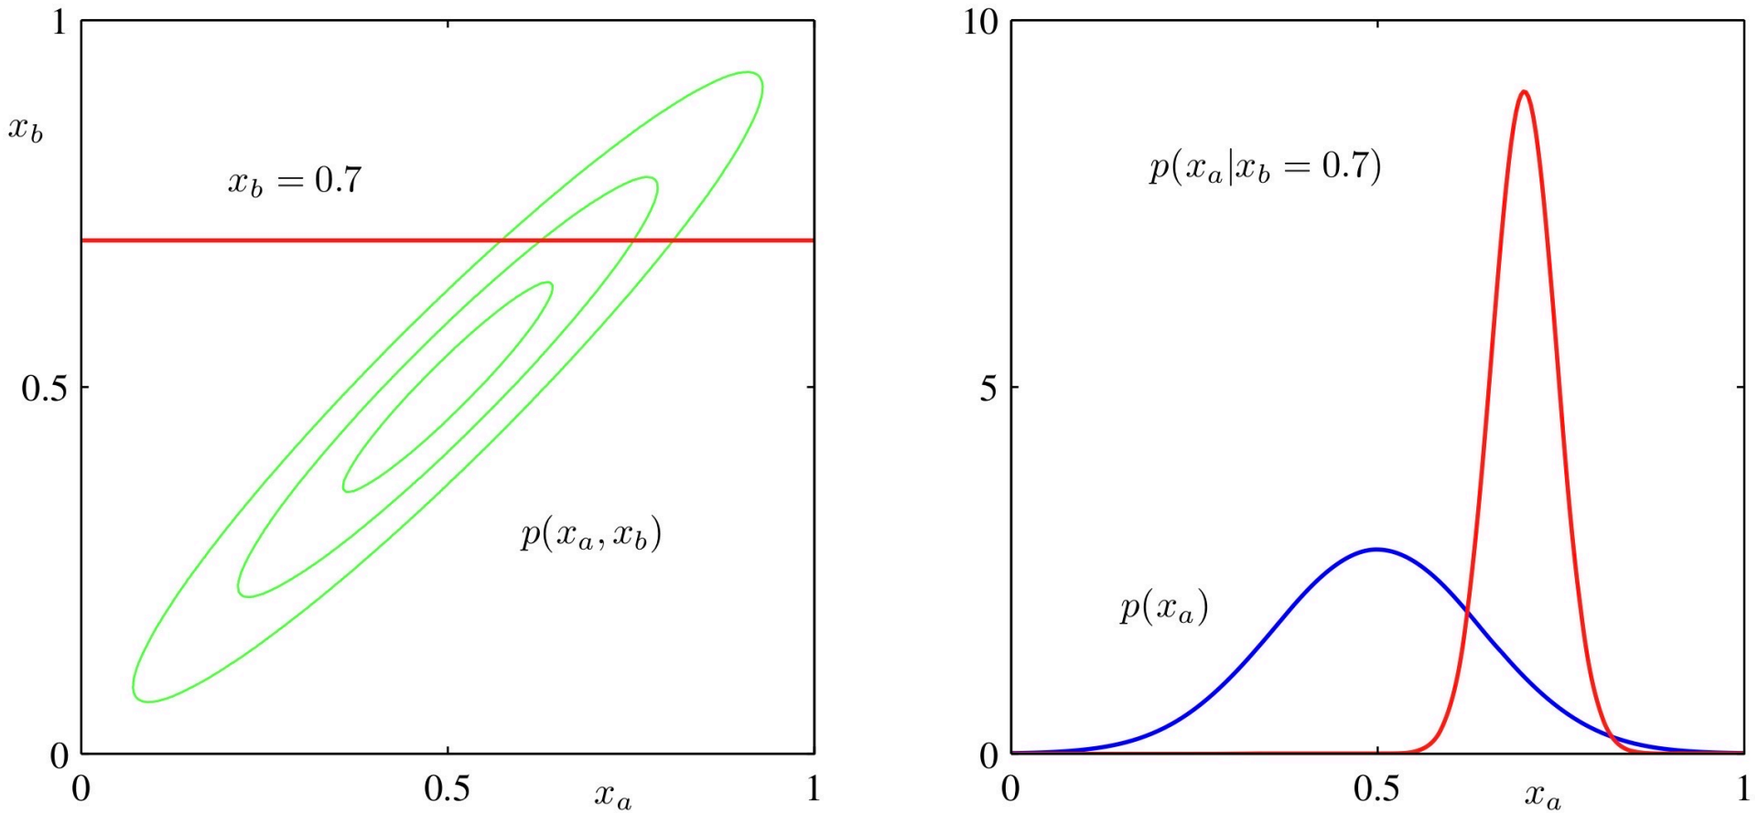

Figure 2.9 The plot on the left shows the contours of a Gaussian distribution $p(x_a, x_b)$ over two variables, and the plot on the right shows the marginal distribution $p(x_a)$ (blue curve) and the conditional distribution $p(x_a|x_b)$ for $x_b = 0.7$ (red curve).

2.3.3 Bayes' theorem for Gaussian variables

Linear Gaussian model example

Gaussian marginal dist. $p(\mathbf{x})$, Gaussian conditional dist $p(\mathbf{y}|\mathbf{x})$

$p(\mathbf{y}|\mathbf{x})$ has a mean as a linear function of $\mathbf{x}$ and

a covariance which is independent of $\mathbf{x}$.

i.e. $$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Lambda}^{-1})$$

$$p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}|\mathbf{A}\mathbf{x}+\mathbf{b}, \mathbf{L}^{-1})$$

$\mathbf{x} \in$ M-dim

$\mathbf{y} \in$ D-dim

where $\boldsymbol{\mu}$, $\mathbf{A}$ and $\mathbf{b}$ are parameters governing the means,

and $\boldsymbol{\Lambda}$ and $\mathbf{L}$ are precision matrices.

We will find $p(y)$ and $p(x|y)$.

marginal conditional

Known $p(x)$, $p(y|x)$

Let $z := \begin{pmatrix} x \\ y \end{pmatrix}$ and us consider the joint prob. dist

$$p(z) = p(x, y) = p(y|x)p(x) \tag{2.102}$$

$$\ln p(z) = \ln p(x) + \ln p(y|x)$$

$$= -\frac{1}{2}(x - \mu_1)^T \Lambda (x - \mu_1)$$

$$-\frac{1}{2}(y - Ax - b)^T L (y - Ax - b) + \text{const}$$

const: indep. of x, y.

This is a quadratic function of the components of x, y in z,

hence $p(z)$ is a Gaussian

Consider the second term in (2.102)

$$-\frac{1}{2}\mathbf{x}^T(\Lambda + A^T L A)\mathbf{x} - \frac{1}{2}\mathbf{y}^T L \mathbf{y} + \frac{1}{2}\mathbf{y}^T L A \mathbf{x} + \frac{1}{2}\mathbf{x}^T A^T L \mathbf{y}$$

$$= -\frac{1}{2}\begin{pmatrix}\mathbf{x}\\ \mathbf{y}\end{pmatrix}^T \begin{pmatrix}\Lambda + A^T L A & -A^T L\\ -LA & L\end{pmatrix}\begin{pmatrix}\mathbf{x}\\ \mathbf{y}\end{pmatrix} = -\frac{1}{2}\mathbf{z}^T R \mathbf{z}$$

$\mathbf{z}$ has precision (inverse of covariance) matrix given by

$$\mathbf{z} = \begin{pmatrix}\mathbf{x}\\ \mathbf{y}\end{pmatrix} \qquad \begin{pmatrix}\Lambda + A^T L A & -A^T L\\ -LA & L\end{pmatrix}$$

$$\Rightarrow \quad \mathrm{cov}[\mathbf{z}] = R^{-1} = \begin{pmatrix}\Lambda^{-1} & \Lambda^{-1} A^T\\ A\Lambda^{-1} & L^{-1} + A\Lambda^{-1} A^T\end{pmatrix} \qquad (2.105)$$

Consider the linear term in (2.102)

$$x^T \Lambda \mu - x^T A^T L b + y^T L b = \begin{pmatrix} x \\ y \end{pmatrix}^T \begin{pmatrix} \Lambda\mu - A^T L b \\ L b \end{pmatrix}$$

$$\Rightarrow \quad \mathbb{E}[z] = R^{-1} \begin{pmatrix} \Lambda\mu - A^T L b \\ L b \end{pmatrix}$$

$$z = \begin{pmatrix} x \\ y \end{pmatrix} \qquad = \begin{pmatrix} \mu \\ A\mu + b \end{pmatrix} \begin{matrix} x \\ y \end{matrix} \qquad (2.108)$$

$Ax+b$

Using section 2.3.2 and $p(y) = \int p(z)\, dx$.

$$\mathbb{E}[y] = A\mu + b$$

$$\text{cov}[y] = L^{-1} + A\Lambda^{-1}A^T$$

Now we can find an expression for $p(\mathbf{x}|\mathbf{y})$.

$$\mathbb{E}[\mathbf{x}|\mathbf{y}] = (\Lambda + \mathbf{A}^T\mathbf{L}\mathbf{A})^{-1}\{\mathbf{A}^T\mathbf{L}(\mathbf{y}-\mathbf{b}) + \Lambda\boldsymbol{\mu}\}$$

$$\mathrm{cov}[\mathbf{x}|\mathbf{y}] = (\Lambda + \mathbf{A}^T\mathbf{L}\mathbf{A})^{-1}$$

Marginal and Conditional Gaussians

Given a marginal Gaussian distribution for $\mathbf{x}$ and a conditional Gaussian distribution for $\mathbf{y}$ given $\mathbf{x}$ in the form

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Lambda}^{-1}) \tag{2.113}$$

$$p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}|\mathbf{A}\mathbf{x} + \mathbf{b}, \mathbf{L}^{-1}) \tag{2.114}$$

the marginal distribution of $\mathbf{y}$ and the conditional distribution of $\mathbf{x}$ given $\mathbf{y}$ are given by

$$p(\mathbf{y}) = \mathcal{N}(\mathbf{y}|\mathbf{A}\boldsymbol{\mu} + \mathbf{b}, \mathbf{L}^{-1} + \mathbf{A}\boldsymbol{\Lambda}^{-1}\mathbf{A}^{\mathrm{T}}) \tag{2.115}$$

$$p(\mathbf{x}|\mathbf{y}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\Sigma}\{\mathbf{A}^{\mathrm{T}}\mathbf{L}(\mathbf{y} - \mathbf{b}) + \boldsymbol{\Lambda}\boldsymbol{\mu}\}, \boldsymbol{\Sigma}) \tag{2.116}$$

where

$$\boldsymbol{\Sigma} = (\boldsymbol{\Lambda} + \mathbf{A}^{\mathrm{T}}\mathbf{L}\mathbf{A})^{-1}. \tag{2.117}$$

2.3.4 Maximum likelihood for the Gaussian

Data set $\mathbf{X} = (\mathbf{x}_1 \dots \mathbf{x}_N)^T$, $\{\mathbf{x}_n\}$ iid samples of D-dimensional Gaussian. The log likelihood function is given by

$\ln p(\mathbf{X} \mid \mu, \Sigma)$ $\quad$ $\mathbf{X}$ $N \times D$ matrix

$$= -\frac{ND}{2}\ln(2\pi) - \frac{N}{2}\ln|\Sigma| - \frac{1}{2}\sum_{n=1}^{N}(\mathbf{x}_n - \mu)^T \Sigma^{-1}(\mathbf{x}_n - \mu)$$

Note that likelihood function depends only on the following two quantities

$$\sum_{n=1}^{N} \mathbf{x}_n \qquad \sum_{n=1}^{N} \mathbf{x}_n \mathbf{x}_n^T$$

These are known as sufficient statistics for Gaussian

$$\nabla_{\mu} \ln p(\mathbf{X} \mid \mu, \Sigma) = \sum_{n=1}^{N} \Sigma^{-1} (\mathbf{x}_n - \mu)$$

D-dim vector

Set this gradient to zero vector, we obtain

$$\mu_{ML} = \frac{1}{N} \sum_{n=1}^{N} \mathbf{x}_n$$

(sample mean)

solution of maximum likelihood estimator

MLE

$$\Sigma_{ML} = \frac{1}{N} \sum_{n=1}^{N} (\mathbf{x}_n - \mu_{ML})(\mathbf{x}_n - \mu_{ML})^T$$

(sample covariance)

Remark

- Σ_{ML} involves μ_{ML}
- μ_{ML} is independent of Σ_{ML}

$$\hat{\sigma}^2 = \frac{1}{N-1}\sum (x - \mu_{..})^2$$

$$E[\hat{\sigma}^2] = \sigma^2$$

Evaluate the expectations of thes solutions under the true distribution. Then we obtain

$$E[\mu_{ML}] = \mu$$ unbiased estimate

$$E[\Sigma_{ML}] = \frac{N-1}{N}\Sigma$$

Σ

2.3.5 Sequential estimation

Sequential estimation for maximum likelihood

This method allows data points to be proceed one at time and then discarded and are important for on-line applications

Consider

$$\mu_{ML} = \frac{1}{N} \sum_{n=1}^{N} x_n$$

which we will denote by $\mu_{ML}^{(N)}$ based on N observations

$$\mu_{ML}^{(N)} = \frac{1}{N}\sum_{n=1}^{N} x_n$$

$$= \frac{1}{N} x_N + \frac{1}{N}\sum_{n=1}^{N-1} x_n$$

$$= \frac{1}{N} x_n + \frac{N-1}{N}\mu_{ML}^{(N-1)}$$

$$= \mu_{ML}^{(N-1)} + \frac{1}{N}\left(x_N - \mu_{ML}^{(N-1)}\right)$$

error signal

x_N, $x_1 \dots x_{N-1}$

$\Downarrow$

MLE

$\mu_{ML}^{(N-1)} = \frac{1}{N-1}\sum_{n=1}^{N-1} x_n$

$\mu_{ML}^{(N-1)}$

$$\theta_{ML}^{(N)} = \theta_{ML}^{(N-1)} - a_N \nabla(\theta^{(N-1)})$$

General formulation of sequential learning (Robbins-Monro)

Two r.v. z, θ governed by a joint distribution $p(z, \theta)$

Define deterministic function $f(\theta)$ by

$$f(\theta) := \mathbb{E}[z|\theta] = \int z\, p(z|\theta)\, dz$$

$\mathbb{E}[t|x]$

conditional expectation

which is a function of θ (called regression function)

Find the root θ^* at which $f(\theta^*) = 0$

Suppose we can observe values of z one at a time

Figure 2.10 A schematic illustration of two correlated random variables z and θ, together with the regression function $f(\theta)$ given by the conditional expectation $\mathbb{E}[z|\theta]$. The Robbins-Monro algorithm provides a general sequential procedure for finding the root $\theta^\star$ of such functions.

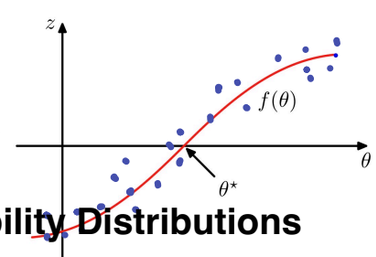

Assume the conditional variance of z is finite

$$E[(z - f)^2 | \theta] < \infty$$

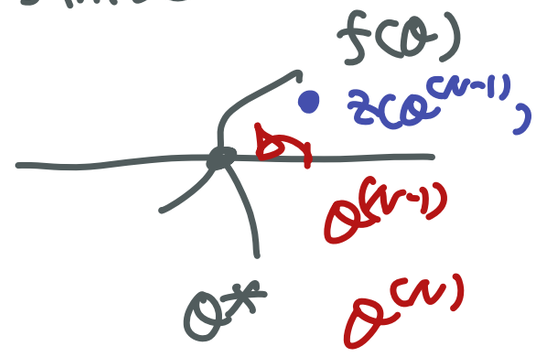

θ^* solution

and wlog $f(\theta) > 0$ for $\theta > \theta^*$ and $f(\theta) < 0$ for $\theta < \theta^*$

A sequence of successive estimates of the root θ^* given by

$$\theta^{(N)} := \theta^{(N-1)} - a_{N-1} z(\theta^{(N-1)}) \qquad (2.129)$$

where $z(\theta^{(N)})$ is an observed value of z when $\theta = \theta^{(N)}$

$\{a_N\}$ represents a seq of positive numbers satisfying

$$\lim_{N\to\infty} a_N = 0$$

$$\sum_{N=1}^{\infty} a_N = \infty \qquad \text{e.g. } \frac{1}{N}$$

$$\sum_{N=1}^{\infty} a_N^2 < \infty$$

of $f(\theta)$

By [Robbins - Monro], (2.129) converges to the root with probability one.

Remark

- Third condition ensures that the accumulated noise has finite variance and hence does not spoil convergence.

General Maximum likelihood problem

$$f(\theta) = \int z\, p(z|\theta)\, dz$$

By definition of θ_{ML}, θ_{ML} satisfies

$$= \mathbb{E}_x\left[\frac{\partial}{\partial\theta} \ln p(x|\theta)\right]$$

$$\frac{\partial}{\partial\theta}\left\{-\frac{1}{N}\sum_{n=1}^{N} \ln p(x_n|\theta)\right\}\Big|_{\theta_{ML}} = 0$$

Taking $N \to \infty$ and exchanging derivative and summation,

$$-\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N}\frac{\partial}{\partial\theta}\ln p(x_n|\theta) = \mathbb{E}_x\left[-\frac{\partial}{\partial\theta}\ln p(x|\theta)\right]$$

$$\mathbb{E}_x[f(x)] \approx \lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N} f(x_n) \quad \text{observations } x_n$$

I.e. find the root of a regression function

Apply Robbins – Monro procedure

$$\theta^{(N)} := \theta^{(N-1)} - a_{N-1} \frac{\partial}{\partial \theta^{(N-1)}} [-\ln p(x_N | \theta^{(N-1)})] \quad (2.135)$$

Specific example: sequential estimation of the mean of Gaussian distribution

In this case $\theta^{(N)}$ is the $\mu_{ML}^{(N)}$ mean of the Gaussian and z is given by

(2.136)

$$z = \frac{\partial}{\partial \mu_{ML}} \ln p(x | \mu_{ML}, \sigma^2) = -\frac{1}{\sigma^2}(x - \mu_{ML})$$

Substituting (2.136) into (2.135) with $a_N = \sigma^2/N-1$

then we obtain (2.126)

2.3.6 Bayesian inference for the Gaussian

MLE method gave point estimates for μ, Σ (section 2.3.4)

Now develop a Bayesian treatment

Single Gaussian random variable x. Suppose σ^2 is known

Aim to inference μ given N observations $\mathbf{x} = \{x_1, \ldots x_N\}$

The likelihood function is given by

$$p(\mathbf{x}|\mu) = \prod_{n=1}^{N} p(x_n|\mu) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp\left\{-\frac{1}{2\sigma^2}\sum_{n=1}^{N}(x_n-\mu)^2\right\}$$

Note that this function is the form of the exponential of a quadratic form of μ.

We will choose a prior $p(\mu)$ given by Gaussian because the product of two exponentials of quadratic function of μ will also be Gaussian

Take prior prob. $p(\mu)$ to be

$$p(\mu) = \mathcal{N}(\mu \mid \mu_0, \sigma_0^2)$$

μ_0, σ_0^2 hyperparameters

Posterior

$$p(\mu \mid x) \propto p(x \mid \mu) \cdot \underbrace{p(\mu)}_{\mathcal{N}(\mu \mid \mu_0, \sigma_0^2)}$$

Exercise 2.38. we obtain

posterior parameters

$$p(\mu|X) = \mathcal{N}(\mu|\mu_N, \sigma_N^2)$$

where $\mu_N = \frac{\sigma^2}{N\sigma_0^2 + \sigma^2}\mu_0 + \frac{N\sigma_0^2}{N\sigma_0^2 + \sigma^2}\mu_{ML}$ // $\frac{1}{N}\sum x_n$

$$\frac{1}{\sigma_N^2} = \frac{1}{\sigma_0^2} + \frac{N}{\sigma^2}$$

Remark

- μ_N is a compromise between μ_0 and μ_{ML}
- Effect of change in value N

- Precision is additive, if $N \to \infty$, $\sigma_N^2 \to 0$
- When N is finite, if $\sigma_0^2 \to \infty$, then the poterior mean reduces to μ_{ML} and variance σ_N^2 becomes $\frac{\sigma^2}{N}$

Sequential inference in Bayesian paradigm

$$p(\mu \mid \mathcal{X}) \propto \left[p(\mu) \prod_{n=1}^{N-1} p(x_n \mid \mu) \right] p(x_N \mid \mu)$$

Posterior

$\propto$ posterior distribution after observing $N-1$ data

$\mathcal{X} = \{ x_1, \ldots x_N \}$

② Now we wish to infer the variance and assume mean is known.

Let the precision $\lambda := 1/\sigma^2$. The likelihood function of λ

$$p(\mathbf{x}|\lambda) = \prod_{n=1}^{N} \mathcal{N}(x_n|\mu, \lambda^{-1}) \propto \lambda^{N/2} \exp\left\{-\frac{\lambda}{2}\sum_{n=1}^{N}(x-\mu)^2\right\}$$

i.e. the form of $\lambda^{N/2} \cdot \exp(-\frac{\lambda}{2})$

The corresponds to gamma distribution

$$\mathrm{Gam}(\lambda|a,b) = \frac{1}{\Gamma(a)} b^a \lambda^{a-1} \exp(-b\lambda), \qquad \lambda > 0$$

Here $\Gamma(a)$ is a gamma function $\Gamma(a) = \int_0^\infty u^{a-1} e^{-u} du$

Remark

- If $a > 0$, gamma distribution has finite integral.
- If $a \geq 1$, the distribution itself is finite.
- $$\mathbb{E}[\lambda] = \frac{a}{b} \qquad \mathrm{var}[\lambda] = \frac{a}{b^2}$$

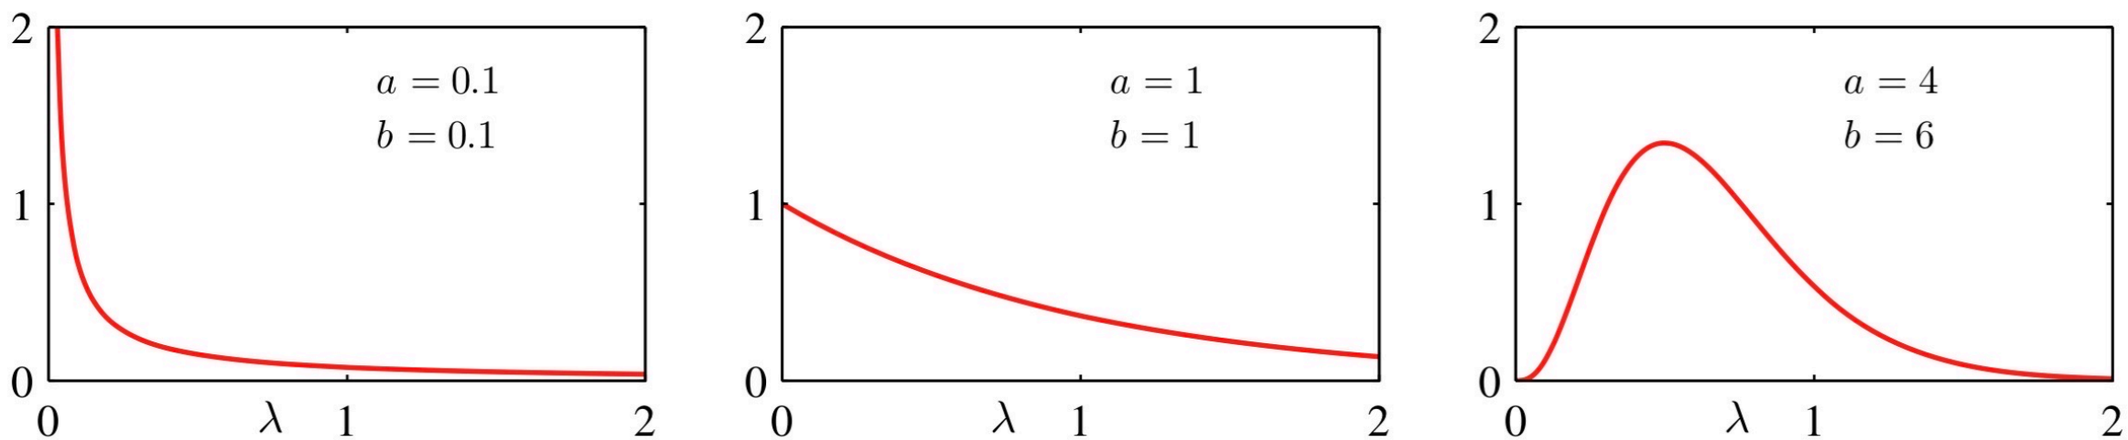

Figure 2.13 Plot of the gamma distribution $\mathrm{Gam}(\lambda|a, b)$ defined by (2.146) for various values of the parameters a and b.

Consider the prior dist. $\mathrm{Gam}(\lambda \mid a_0, b_0)$. ($a_0, b_0$: hyperparameter)

The posterior dist. of λ is as below

$$p(\lambda \mid \mathbf{x}) \propto \underbrace{\lambda^{N/2} \exp\left\{-\frac{\lambda}{2}\sum_{n=1}^{N}(x_n-\mu)^2\right\}}_{\text{likelihood function of } \lambda} \cdot \underbrace{\mathrm{Gam}(\lambda \mid a_0, b_0)}_{\text{prior of } \lambda}$$

$$\propto \lambda^{a_0-1}\lambda^{N/2}\exp\left\{-b_0\lambda - \frac{\lambda}{2}\sum_{n=1}^{N}(x_n-\mu)^2\right\} \qquad (2.149)$$

$$\Rightarrow \quad p(\lambda \mid \mathbf{x}) = \mathrm{Gam}(\lambda \mid a_N, b_N) \quad \text{where}$$

$$a_N = a_0 + \frac{N}{2}$$

$$b_N = b_0 + \frac{1}{2}\sum_{n=1}^{N}(x_n-\mu)^2 = b_0 + \frac{N}{2}\sigma_{ML}^2$$

Remark

- Effect of observing N data points
  - increases the value of a by $\frac{N}{2}$
  - " b by $\frac{N}{2}\sigma_{ML}^2$
- We can interprete the parameter a_0 in terms of $2a_0$ 'effective' prior observations.
- $\mathbb{E}[\lambda|\mathcal{X}] = \frac{a_N}{b_N} = \frac{2a_0 + N}{2b_0 + N\sigma_{ML}^2}$

$\text{var}[\lambda|\mathcal{X}] = \frac{a_N}{b_N^2}$ $\qquad \mathbb{E}[\lambda] \longrightarrow \frac{1}{\sigma_{ML}^2}$ $\qquad \lambda = \frac{1}{\sigma^2}$

Now suppose that the both μ and λ are unknown

Consider the dependence of the likelihood function on μ and λ

$$p(\mathbf{x} \mid \mu, \lambda) = \prod_{n=1}^{N} \left(\frac{\lambda}{2\pi}\right)^{1/2} \exp\left\{-\frac{\lambda}{2}(x_n - \mu)^2\right\}$$

$$\lambda, \mu \qquad \propto \left[\lambda^{1/2} \exp\left(-\frac{\lambda\mu^2}{2}\right)\right]^N \exp\left[\lambda\mu \sum_{n=1}^{N} x_n - \frac{\lambda}{2} \sum_{n=1}^{N} x_n^2\right]$$

Thus the prior distribution should take the form

$$p(\mu, \lambda) \propto \left[\lambda^{1/2} \exp\left(-\frac{\lambda\mu^2}{2}\right)\right]^{\beta} \exp\{c\lambda\mu - d\lambda\}$$

$$= \exp\left\{-\frac{\beta\lambda}{2}(\mu - c/\beta)^2\right\} \lambda^{\beta/2} \exp\left\{-\left(d - \frac{c^2}{2\beta}\right)\lambda\right\} \tag{2.153}$$

μ, λ $\qquad$ λ

where c, d and β are constants. Use $p(\mu, \tau) = p(\mu|\tau)\,p(\tau)$.

$p(\mu|\tau)$: a Gaussian whose precision is a linear function of τ

$p(\tau)$: a gamma distribution. So we take a prior

$$p(\mu, \tau) = \mathcal{N}(\mu | \mu_0, (\beta\tau)^{-1})\,\mathrm{Gam}(\tau | a, b) \qquad (2.154)$$

where $\mu_0 := c/\beta$, $a := (1+\beta)/2$, $b := d - c^2/2\beta$

(2.154) is called normal gamma or Gaussian gamma.

Note that it is not the simply the product of an independent Gaussian prior and gamma prior.

Multivariate Gaussian $\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Lambda}^{-1})$ for D-dim $\mathbf{x}$

First, when precision matrix $\boldsymbol{\Lambda}$ is known, the conjugate prior distribution is again a Gaussian.

Second, for known mean and unknown precision matrix $\boldsymbol{\Lambda}$, the conjugate prior distribution is the Wishart distribution given by

$$\mathcal{W}(\boldsymbol{\Lambda} \mid \mathbf{W}, \nu) = B|\boldsymbol{\Lambda}|^{(\nu-D-1)/2} \exp\left(-\frac{1}{2}\mathrm{Tr}(\mathbf{W}^{-1}\boldsymbol{\Lambda})\right)$$

trace of matrix

where ν is called the number of degrees of freedom

W is a $D \times D$ scale matrix. The normalization constant B is given by

$$B(W, \nu) = |W|^{-\nu/2} \left(2^{\nu D/2} \pi^{D(D-1)/4} \prod_{i=1}^{D} \Gamma\left(\frac{\nu + 1 - i}{2}\right) \right)^{-1}.$$

If both the mean and precision are unknown, the conjugate prior is given by

$$p(\mu, \Lambda \mid \mu_0, \beta, W, \nu) = \mathcal{N}(\mu \mid \mu_0, (\beta\Lambda)^{-1}) \mathcal{W}(\Lambda \mid W, \nu)$$

which is known as the normal-Wishart or Gaussian-Wishart

2.3.7 Student's t-distribution

Conjugate prior for the precision of a Gaussian is given by a gamma distribution.

Consider univariate Gaussian $\mathcal{N}(x \mid \mu, \tau^{-1})$ with Gamma prior $\mathrm{Gam}(\tau \mid a, b)$. Integrate out the precision

$$p(x \mid \mu, a, b) = \int_0^\infty \mathcal{N}(x \mid \mu, \tau^{-1}) \, \mathrm{Gam}(\tau \mid a, b) \, d\tau$$

$$= \int_0^\infty \frac{b^a e^{(-b\tau)} \tau^{a-1}}{\Gamma(a)} \left(\frac{\tau}{2\pi}\right)^{1/2} \exp\left\{-\frac{\tau}{2}(x-\mu)^2\right\} d\tau$$

$$= \frac{b^a}{\Gamma(a)} \left(\frac{1}{2\pi}\right)^{1/2} \left[b + \frac{(x-\mu)^2}{2}\right]^{-a-1/2} \Gamma(a+1/2)$$

where we have made the change of variable $z = \tau[b + (x-\mu)^2/2]$

Define new parameters $\nu = 2a$ and $\lambda = a/b$.

$$St(x|\mu,\lambda,\nu) = \frac{\Gamma(\nu/2+1/2)}{\Gamma(\nu/2)}\left(\frac{\lambda}{\pi\nu}\right)^{1/2}\left[1+\frac{\lambda(x-\mu)^2}{\nu}\right]^{-\nu/2-1/2}$$

known as Student's t-distribution. λ is called precision and ν is called the degree of freedom.

When $\nu = 1$, t-distribution reduces to the Cauchy dist.

While in the limit $\nu \to \infty$, t-distribution becomes Gaussian $\mathcal{N}(x|\mu,\lambda^{-1})$

Remark

- t-dist. can be interpreted infinite mixture of Gaussian
- Longer tail, robustness property

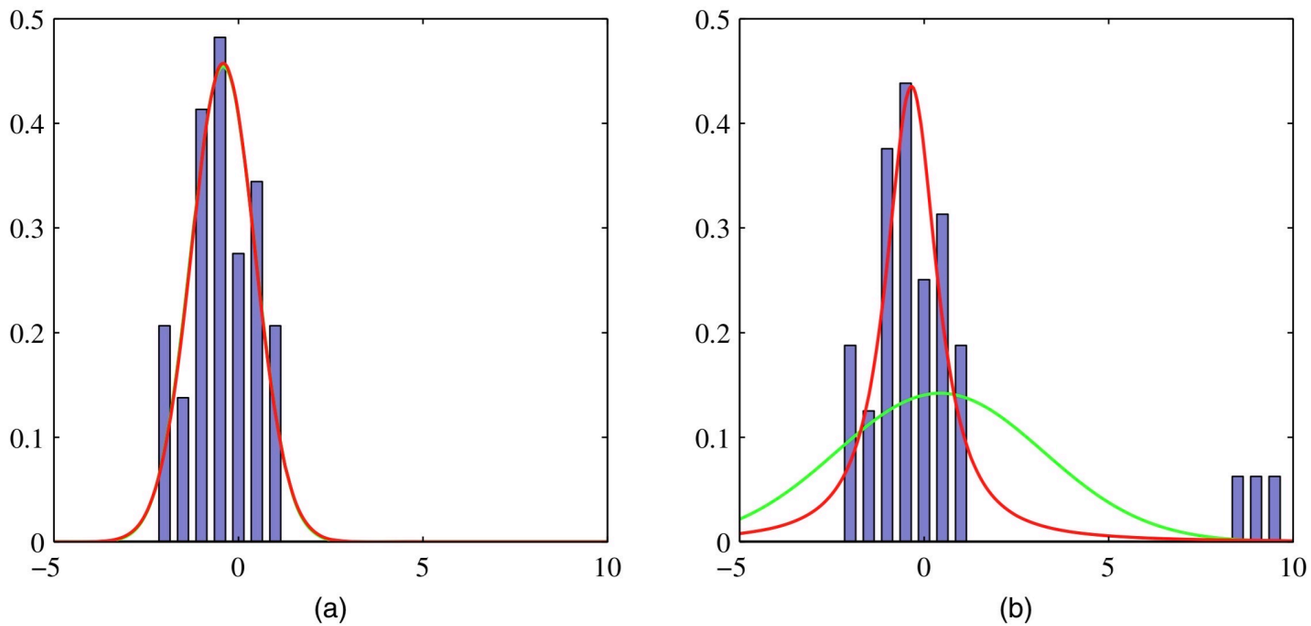

Figure 2.16 Illustration of the robustness of Student's t-distribution compared to a Gaussian. (a) Histogram distribution of 30 data points drawn from a Gaussian distribution, together with the maximum likelihood fit obtained from a t-distribution (red curve) and a Gaussian (green curve, largely hidden by the red curve). Because the t-distribution contains the Gaussian as a special case it gives almost the same solution as the Gaussian. (b) The same data set but with three additional outlying data points showing how the Gaussian (green curve) is strongly distorted by the outliers, whereas the t-distribution (red curve) is relatively unaffected.

Multivariate Student's t-distribution

$$\mathrm{St}(x \mid \mu, \Lambda, \nu) = \int_0^\infty \mathcal{N}(x \mid \mu, (\eta\Lambda)^{-1}) \, \mathrm{Gam}(\eta \mid \nu/2, \nu/2) \, d\eta$$

$$= \frac{\Gamma(D/2 + \nu/2)}{\Gamma(\nu/2)} \frac{|\Lambda|^{1/2}}{(\pi\nu)^{D/2}} \left[1 + \frac{\Delta^2}{\nu} \right]^{-D/2 - \nu/2}$$

where $\Delta := (x - \mu)^T \Lambda (x - \mu)$

Remark

- $\mathbb{E}[x] = \mu$ if $\nu > 1$
- $\mathrm{cov}[x] = \frac{\nu}{(\nu - 2)} \Lambda^{-1}$ if $\nu > 2$
- $\mathrm{mode}[x] = \mu$

2.3.8 Periodic variables

Consider an angular (polar) coordinate $0 \leq \theta < 2\pi$ and the problem of evaluating the mean of observations $D = \{\theta_1, \ldots \theta_N\}$.

① Simple average $(\theta_1 + \cdots \theta_N)/N$ is strongly coordinate dependent.

② Set angular observations as points on unit circle.

Let x_i be two-dim vector with $x_i = (\cos\theta_i, \sin\theta_i)$

embedding

Average the vectors $\{\mathbf{x}_n\}$ instead to give

$$\bar{\mathbf{x}} = \frac{1}{N}\sum_{n=1}^{N}\mathbf{x}_n$$

$$= \bar{r}(\cos\bar{\theta}, \sin\bar{\theta})$$

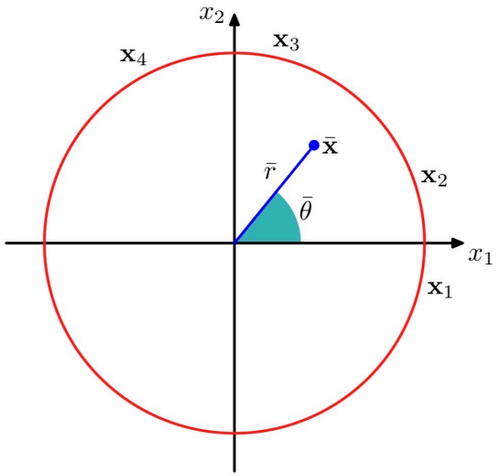

Figure 2.17 Illustration of the representation of values θ_n of a periodic variable as two-dimensional vectors $\mathbf{x}_n$ living on the unit circle. Also shown is the average $\bar{\mathbf{x}}$ of those vectors.

i.e. $\bar{r}\cos\bar{\theta} = \frac{1}{N}\sum_{n=1}^{N}\cos\theta_n$, $\bar{r}\sin\bar{\theta} = \frac{1}{N}\sum_{n=1}^{N}\sin\theta_n$

Thus we can solve for $\bar{\theta}$ to give

$$\bar{\theta} = \tan^{-1}\left\{\frac{\sum_n \sin\theta_n}{\sum_n \cos\theta_n}\right\}$$

Consider $p(\theta)$ that have period 2π and must satisfies

$$p(\theta) \geq 0$$

$$\int_0^{2\pi} p(\theta)\, d\theta = 1$$

$$p(\theta + 2\pi) = p(\theta)$$

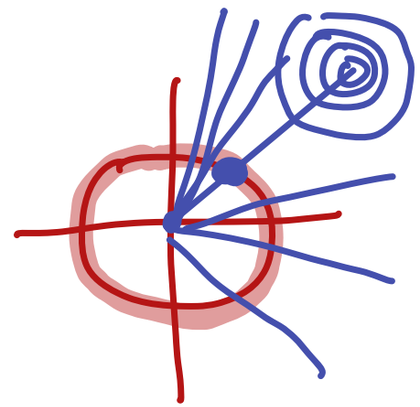

We can easily obtain a Gaussian - like distribution.

Consider a Gaussian over $\mathbf{x} = (x_1, x_2)$ having mean $\boldsymbol{\mu} = (\mu_1, \mu_2)$

and covariance matrix $\Sigma = \sigma^2 I$ so that

$$p(x_1, x_2) = \frac{1}{2\pi\sigma^2} \exp\left\{ -\frac{(x_1 - \mu_1)^2 + (x_2 - \mu_2)^2}{2\sigma^2} \right\} \qquad (2.173)$$

Map $\mathbf{x} = (x_1, x_2)$ and $\boldsymbol{\mu}$ into polar coordinates

$$\begin{cases} x_1 = r\cos\theta, & x_2 = r\sin\theta \\ \mu_1 = r_0\cos\theta_0, & \mu_2 = r_0\sin\theta_0 \end{cases} \qquad \text{fixed } r_0, \theta_0$$

Substitute these transformation into (2.173) with $r = 1$ condition

The exponent in (2.173)

$$-\frac{1}{2\sigma^2}\{(r\cos\theta - r_0\cos\theta_0)^2 + (r\sin\theta - r_0\sin\theta_0)^2\} \qquad (r = 1)$$

$$= -\frac{1}{2\sigma^2}\{1 + r_0^2 - 2r_0\cos\theta\cos\theta_0 - 2r_0\sin\theta\sin\theta_0\}$$

$$= \frac{r_0}{\sigma^2}\cos(\theta - \theta_0) + \text{const}$$

Define $m = r_0/\sigma^2$. Then we obtain the expression for the distribution of $p(\theta)$ along unit circle

$$p(\theta \mid \theta_0, m) = \frac{1}{2\pi} \frac{1}{I_0(m)} \exp\{ m \cos(\theta - \theta_0) \}$$

which called von Mises distribution. Here θ_0 represents the mean and $m = r_0/\sigma^2$ is called concentration parameter.

$I_0(m)$: zeroth - order Bessel function of the first kind

$$I_0(m) := \frac{1}{2\pi} \int_0^{2\pi} \exp\{ m \cos\theta \} \, d\theta$$

Now consider the maximum likelihood for θ_0 and m

Observations $D = \{\theta_1, \ldots \theta_N\}$ is given

$$\ln p(D \mid \theta_0, m) = \prod_{n=1}^{N} p(\theta_n \mid \theta_0, m) \qquad (2.181)$$

$$= -N \ln(2\pi) - N \ln I_0(m) + m \sum_{n=1}^{N} \cos(\theta_n - \theta_0)$$

Set the derivative w.r.t θ_0 equal to zero gives

$$\sum_{n=1}^{N} \sin(\theta_n - \theta_0) = 0$$

Thus, we obtain

$$\theta_0^{ML} = \tan^{-1}\left\{\frac{\sum \sin\theta_n}{\sum \cos\theta_n}\right\}$$

Similarly maximizing (2.181) w.r.t m. Set the derivative of (2.181) w.r.t m, then we have

$$A(m) = \frac{1}{N} \sum_{n=1}^{N} \cos(\theta_n - \theta_0^{ML}) \qquad (2.185)$$

where we used $I_0'(m) = I_1(m)$ and have defined

$$A(m) := \frac{I_1(m)}{I_0(m)}$$

We can rewrite (2.185) in the form

$$A(m_{ML}) = \left(\frac{1}{N} \sum_{n=1}^{N} \cos\theta_n \right) \cos\theta_0^{ML} + \left(\frac{1}{N} \sum_{n=1}^{N} \sin\theta_n \right) \sin\theta_0^{ML}$$

Here A(m) can be inverted unmerically.

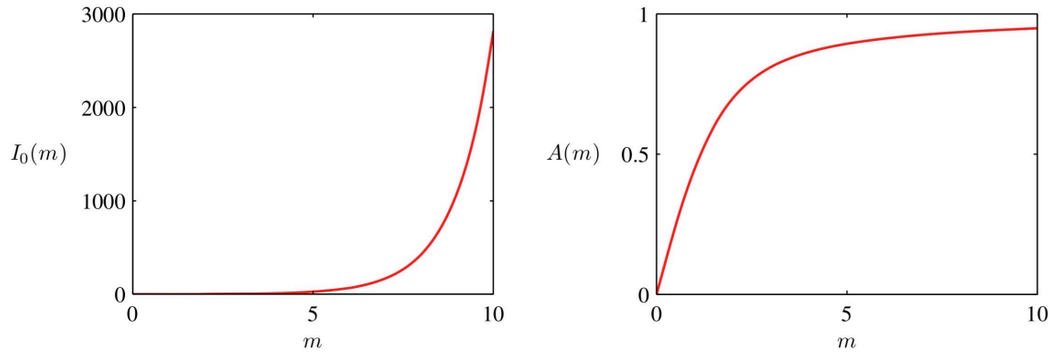

Figure 2.20 Plot of the Bessel function $I_0(m)$ defined by (2.180), together with the function $A(m)$ defined by (2.186).

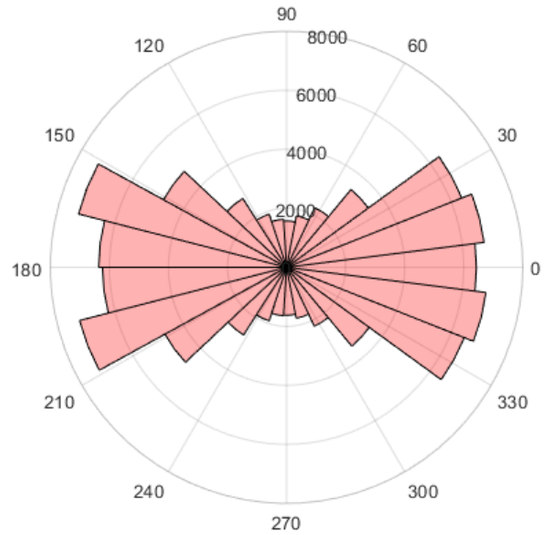

Remark: other techniques to construct periodic variable

- Histogram in polar coordinates
- Mixtures of von Mises distributions

2.3.9 Mixtures of Gaussians

Limitations of a Gaussian (unimodal)

Figure 2.21 Plots of the 'old faithful' data in which the blue curves show contours of constant probability density. On the left is a single Gaussian distribution which has been fitted to the data using maximum likelihood. Note that this distribution fails to capture the two clumps in the data and indeed places much of its probability mass in the central region between the clumps where the data are relatively sparse. On the right the distribution is given by a linear combination of two Gaussians which has been fitted to the data by maximum likelihood using techniques discussed Chapter 9, and which gives a better representation of the data.

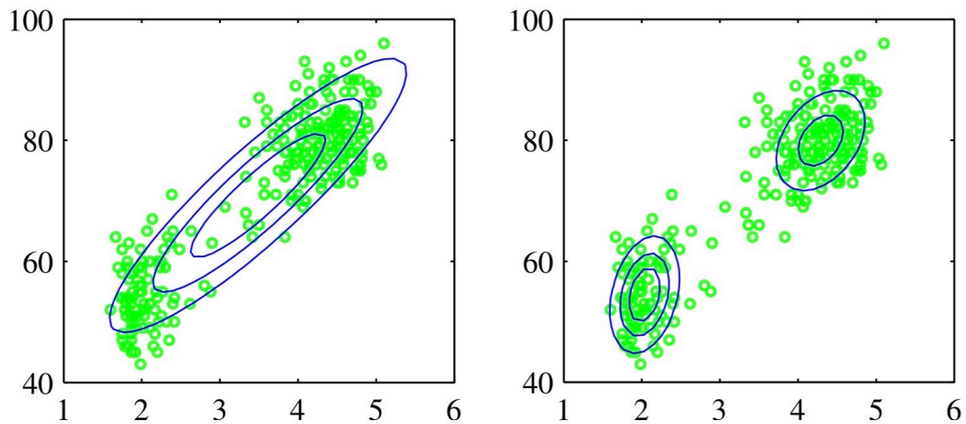

Mixture distribution: linear combinations of basic distributions.

Mixture of Gaussians : superposition of k Gaussians

$$p(x) = \sum_{k=1}^{K} \pi_k \mathcal{N}(x \mid \mu_k, \Sigma_k)$$

Each $\mathcal{N}(x \mid \mu_k, \Sigma_k)$ is called a component of mixture

The parameter π_k are called mixing coefficients and satisfies

$$\sum_{k=1}^{K} \pi_k = 1 \quad , \qquad 0 \leq \pi_k \leq 1$$

Figure 2.22 Example of a Gaussian mixture distribution in one dimension showing three Gaussians (each scaled by a coefficient) in blue and their sum in red.

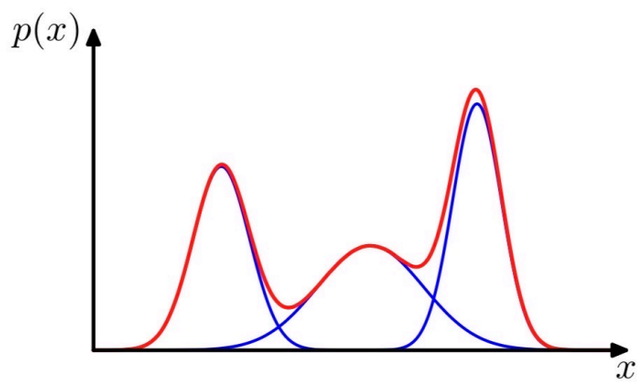

$p(x)$ can be rewrite in the form

$$p(x) = \sum_{k=1}^{K} p(k)\, p(x \mid k)$$

$\pi_k = p(k)$ prior probability of picking the k^{th} component

$\mathcal{N}(x \mid \mu_k, \Sigma_k) = p(x \mid k)$ probability of x conditioned on k.

Consider the posterior $p(k|x)$ a.k.a. responsibilities

$$\gamma_k(x) := p(k|x) = \frac{p(k,x)}{p(x)} = \frac{p(k)\,p(x|k)}{\sum_\ell p(\ell)\,p(x|\ell)}$$

$$= \frac{\pi_k\,\mathcal{N}(x|\mu_k,\Sigma_k)}{\sum_\ell \pi_\ell\,\mathcal{N}(x|\mu_\ell,\Sigma_\ell)}$$

Gaussian mixture is governed by $\pi := \{\pi_1 \dots \pi_K\}$, $\mu := \{\mu_1 \dots \mu_K\}$ and $\Sigma := \{\Sigma_1, \dots \Sigma_K\}$

One way to set these parameters is to use maximum likelihood.

$$\ln p(X \mid \pi, \mu, \Sigma) = \sum_{n=1}^{N} \ln \left\{ \sum_{k=1}^{K} \pi_k \mathcal{N}(x_n \mid \mu_k, \Sigma_k) \right\}$$

Maximum likelihood solution for the parameters no longer has a closed-form analytical solution.

Expectation maximization (chapter 9)

2.4 The exponential family

Broad class of distributions called the exponentials family

Random vector $\mathbf{x}$, parameters $\boldsymbol{\eta}$ (called natural parameters)

$$p(\mathbf{x} \mid \boldsymbol{\eta}) := h(\mathbf{x})\, g(\boldsymbol{\eta})\, \exp\{\boldsymbol{\eta}^T u(\mathbf{x})\} \tag{2.194}$$

Here $u(\mathbf{x})$ is some function of $\mathbf{x}$ and $g(\boldsymbol{\eta})$ can be interpreted as the normalization coefficient. i.e.

$$g(\boldsymbol{\eta}) \int h(\mathbf{x}) \exp\{\boldsymbol{\eta}^T u(\mathbf{x})\}\, d\mathbf{x} = 1 \tag{2.195}$$

Recall Bernoulli distribution

$$p(x|\mu) = \text{Bern}(x|\mu) = \mu^x (1-\mu)^{1-x} \qquad x = 0, 1$$

$$= \exp\{ x \ln \mu + (1-x) \ln(1-\mu) \}$$

$$= (1-\mu) \exp\left\{ \ln\left(\frac{\mu}{1-\mu}\right) x \right\}$$

Comparison with (2.194)

$$\eta = \ln\left(\frac{\mu}{1-\mu}\right)$$

Solve for μ to give

$$\mu = \sigma(\eta) := \frac{1}{1 + \exp(-\eta)}$$ logistic sigmoid

Thus, Bernoulli distribution can be rewrited in the form

$$p(x|\eta) = \sigma(-\eta)\exp(\eta x)$$

we have used $\sigma(-\eta) = 1 - \sigma(\eta)$. Comparison with (2.194)

$$u(x) = x$$

$$h(x) = 1$$

$$g(\eta) = \sigma(-\eta)$$

Next consider the multinomial distribution

$$p(\mathbf{x}|\boldsymbol{\mu}) = \prod_{k=1}^{M} \mu_k^{x_k} = \exp\left\{ \sum_{k=1}^{M} x_k \ln \mu_k \right\}$$

where $\mathbf{x} = (x_1, \ldots x_M)^T$ (one-hot vector)

The standard representation (2.194) so that

$$p(\mathbf{x}|\boldsymbol{\eta}) = \exp(\boldsymbol{\eta}^T \mathbf{x})$$

where $\boldsymbol{\eta} = (\eta_1 \ldots \eta_M)^T$ with $\eta_k = \ln \mu_k$. i.e.

$$u(\mathbf{x}) = \mathbf{x}, \quad h(\mathbf{x}) = 1, \quad g(\boldsymbol{\eta}) = 1$$

Since $\sum_{k=1}^{M} \mu_k = 1$, parameters η_k are not independent.

i.e. $\mu_M = 1 - \sum_{k=1}^{M-1} \mu_k$. left $M-1$ parameters

$$0 \le \mu_k \le 1 \quad , \quad \sum_{k=1}^{M-1} \mu_k \le 1$$

$$-\sum x_k \ln\left(1 - \Sigma \mu_k\right)$$

$$\sum x_k \ln \mu_k$$

So the multinomial distribution becomes $1 - \sum_{k=1}^{M-1} x_k$

$$\exp\left\{ \sum_{k=1}^{M} x_k \ln \mu_k \right\} = \exp\left\{ \sum_{k=1}^{M-1} x_k \ln \mu_k + x_M \underbrace{\ln \mu_M}_{\ln\left(1 - \sum_{k=1}^{M-1} \mu_k\right)} \right\}$$

$$= \exp\left\{ \sum_{k=1}^{M-1} x_k \ln\left(\frac{\mu_k}{1 - \sum_{j=1}^{M-1} \mu_j} \right) + \ln\left(1 - \sum_{k=1}^{M-1} \mu_k \right) \right\}$$

Now identify

$$\ln\left(\frac{\mu_k}{1-\sum_j \mu_j}\right) = \eta_k$$

and

$$\mu_k = \frac{\exp(\eta_j)}{1+\sum_j \exp(\eta_j)}$$

softmax
normalized exponential

In this representation, multinomial distribution

$$p(x|\eta) = \left(1+\sum_{k=1}^{M-1}\exp(\eta_k)\right)^{-1}\exp(\eta^T x)$$

$\eta = (\eta_1 \dots \eta_{M-1}, 0)^T$. ($M$-dimensional)

$$u(x) = x, \quad h(x) = 1, \quad g(\eta) = \left(1 + \sum_{k=1}^{M-1} \exp(\eta_k)\right)^{-1}$$

Finally, consider the univariate Gaussian distribution.

$$p(x|\mu, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{1/2}} \exp\left\{-\frac{1}{2\sigma^2}(x-\mu)^2\right\}$$

$$= \frac{1}{(2\pi\sigma^2)^{1/2}} \exp\left\{-\frac{1}{2\sigma^2}x^2 + \frac{\mu}{\sigma^2}x - \frac{1}{2\sigma^2}\mu^2\right\}$$

$$\eta = \begin{pmatrix} \mu/\sigma^2 \\ -1/2\sigma^2 \end{pmatrix}, \quad u(x) = \begin{pmatrix} x \\ x^2 \end{pmatrix}, \quad h(x) = (2\pi)^{-1/2}, \quad g(\eta) = (-2\eta_2)^{1/2} \exp\left(\frac{\eta_1^2}{4\eta_2}\right)$$

2.4.1 Maximum likelihood and sufficient statistics

Consider the exponential family of distributions over $\mathbf{x}$

$$p(\mathbf{x}|\boldsymbol{\eta}) = h(\mathbf{x})\, g(\boldsymbol{\eta}) \exp\{\boldsymbol{\eta}^T \mathbf{u}(\mathbf{x})\} \qquad (2.194)$$

Taking the gradient of both side of

$$g(\boldsymbol{\eta}) \int h(\mathbf{x}) \exp\{\boldsymbol{\eta}^T \mathbf{u}(\mathbf{x})\}\, d\mathbf{x} = 1$$

w.r.t $\boldsymbol{\eta}$, we have

$$\nabla g(\boldsymbol{\eta}) \int h(\mathbf{x}) \exp\{\boldsymbol{\eta}^T \mathbf{u}(\mathbf{x})\}\, d\mathbf{x}$$

$$+ g(\boldsymbol{\eta}) \int h(\mathbf{x}) \exp\{\boldsymbol{\eta}^T \mathbf{u}(\mathbf{x})\}\, \mathbf{u}(\mathbf{x})\, d\mathbf{x} = 0$$

Using (2.195) then

$$-\frac{1}{g(\eta)}\nabla g(\eta) = g(\eta)\int h(x)\exp\{\eta^T u(x)\}\,u(x)\,dx = \mathbb{E}[u(x)]$$

We therefore obtain the result

$$-\nabla \ln g(\eta) = \mathbb{E}[u(x)]$$

Now consider iid samples denoted by $\mathbf{X} = \{x_1 \ldots x_N\}$

for which likelihood

$$p(\mathbf{X}|\eta) = \left(\prod_{n=1}^{N} h(x_n)\right) g(\eta)^N \exp\left\{\eta^T \sum_{n=1}^{N} u(x_n)\right\}$$

Setting the gradient of $\ln p(\mathbf{X}|\eta)$ w.r.t η to zero.

We get the following condition to be satisfied by η_{ML}

$$-\nabla \ln g(\eta_{ML}) = \frac{1}{N} \sum_{n=1}^{N} u(\mathbf{x}_n)$$

Note that MLE depends on the data only through $\sum_n u(\mathbf{x}_n)$

called sufficient statistic of (2.194)

Do not need to store the entire data.

E.g. Bernoulli $u(x) = x$. sum of $\{x_n\}$

Gaussian $u(x) = (x, x^2)^T$ sum of $\{x_n\}$ and $\{x_n^2\}$.

2.4.2 Conjugate prior

For a given prob. density $p(x|\eta)$, seek a prior $p(\eta)$ that is conjugate to the likelihood function.

(the posterior has the same functional form as the prior)

For exponential family (2.194), $\exists$ conjugate prior of η

$$p(\eta \mid \chi, \nu) = f(\chi, \nu)\, g(\eta)^{\nu} \exp\{\nu \eta^{T} \chi\}$$

where $f(\chi, \nu)$ is a normalization coefficient and $g(\eta)$ is the same function in (2.194)

The posterior

$$p(\eta \mid \mathbf{X}, \chi, \nu) \propto p(\mathbf{X} \mid \eta) \cdot p(\eta \mid \chi, \nu)$$

$$\propto g(\eta)^{\nu+N} \exp\left\{ \eta^T \left(\sum_{n=1}^{N} u(x_n) + \nu\chi \right) \right\}$$

2.4.3 Noninformative priors

Intend to have as litte influence on the posterior as possible.

Let density or likelihood is given by $p(x|\lambda)$.

Consider noninformative prior $p(\lambda)$

First, $p(\lambda) =$ constant

- If the domain of λ is unbounded, prior cannot be normalized. Such prior is called improper.
- Transformation behavior of density under a nonlinear change of variables

Example 1.

Density of x takes the form

$$p(x \mid \mu) = f(x - \mu)$$

μ is known as location parameter. E.g. $\mathcal{N}(x \mid \mu, \sigma^2)$

Translation invariance

If $x \to \hat{x} := x + c$, then

$$\hat{p}(\hat{x} \mid \hat{\mu}) = f(\hat{x} - \hat{\mu})$$

where we have defined $\hat{\mu} := \mu + c$

Thus $p(x|\mu) = p(\hat{x}|\hat{\mu})$ so density is independent of origin.

Prior distribution should satisfy this translation invariance property.

$$\Rightarrow \quad \int_A^B p(\mu)\,d\mu = \int_{A-c}^{B-c} p(\mu)\,d\mu = \int_A^B p(\mu - c)\,d\mu \qquad \forall\, A, B$$

So we have $p(\mu - c) = p(\mu)$

Example of location parameter: mean of a Gaussian

The conjugate prior for μ is again Gaussian $p(\mu \mid \mu_0, \sigma_0^2)$

and we obtain noninformative prior by taking $\sigma_0^2 \to \infty$.

Example 2.

Density of x takes the form

$$p(x|\sigma) = \frac{1}{\sigma} f\left(\frac{x}{\sigma}\right) \qquad \sigma > 0$$

σ is known as scale parameter. E.g. $\mathcal{N}(x|\mu, \sigma^2)$

Scale invariance

If $x \to \hat{x} := cx$, then

$$\hat{p}(\hat{x}|\hat{\sigma}) = \frac{1}{\hat{\sigma}} f\left(\frac{\hat{x}}{\hat{\sigma}}\right)$$

where we have defined $\hat{\sigma} = c\sigma$

So this transformation corresponds to a change of scale.

Prior distribution should satisfy this scale invariance property.

$$\Rightarrow \quad \int_A^B p(\sigma)\,d\sigma = \int_{A/c}^{B/c} p(\sigma)\,d\sigma = \int_A^B p\left(\frac{1}{c}\sigma\right)\frac{1}{c}\,d\sigma \quad \forall A, B$$

So we have $p(\sigma) = p\left(\frac{1}{c}\sigma\right)\frac{1}{c}$ and hence $p(\sigma) \propto \frac{1}{\sigma}$

Note that this is an improper prior because of $0 < \sigma < \infty$

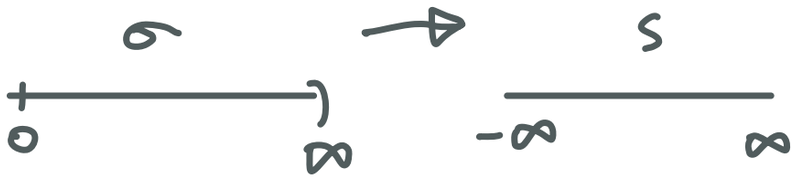

$$p_s(\ln\sigma) = \underbrace{p(\sigma)}_{\propto \frac{1}{\sigma}} \cdot \left|\frac{d\sigma}{ds}\right| = \text{constant}$$

$$s(\sigma) = \ln\sigma \qquad ds = \frac{1}{\sigma}\,d\sigma$$

Convenient to think of prior for scale parameter in terms of the density of the log of the parameter.

Example of scale parameter: standard deviation σ of a Gaussian

$$\mathcal{N}(x \mid \mu, \sigma^2) \propto \sigma^{-1} \exp\{-(\tilde{x}/\sigma)^2\}$$

where $\tilde{x} := x - \mu$

More convenient to work in terms of the precision

$\eta = 1/\sigma^2$ rather than σ itself $\quad d\eta = \frac{1}{\sigma^3} d\sigma$

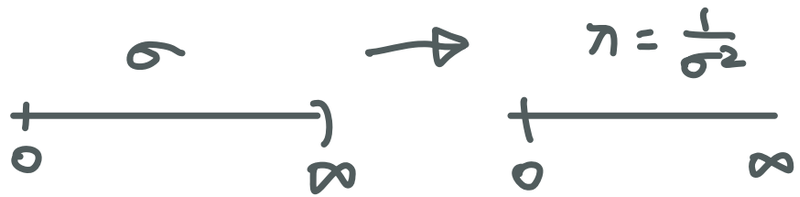

$$p(\eta) = p(\sigma) \cdot \left|\frac{d\sigma}{d\eta}\right| \propto \frac{1}{\eta} \qquad (p(\sigma) \propto \tfrac{1}{\sigma})$$

We have seen the conjugate prior for λ was $\text{Gam}(\lambda | a_0, b_0)$

The non informative is obtained as the special case $a_0 = b_0 = 0$.

2.5 Nonparametric Methods

Approaches to density modeling

Parametric vs Nonparametric (few assumptions)

Histogram method for density estimation

Single continuous random variable x. Partition x into distinct bins of width Δ_i and then count the number n_i of observations of x falling in bin i

We obtain probability values for each bin given by

$$p_i = \frac{n_i}{N \Delta_i}$$

where N is ✱ of total observations. So a model for the density $p(x)$ is piecewise constant over the width Δ_i of each bin, and often the bins are chosen to have the same width $\Delta_i = \Delta$

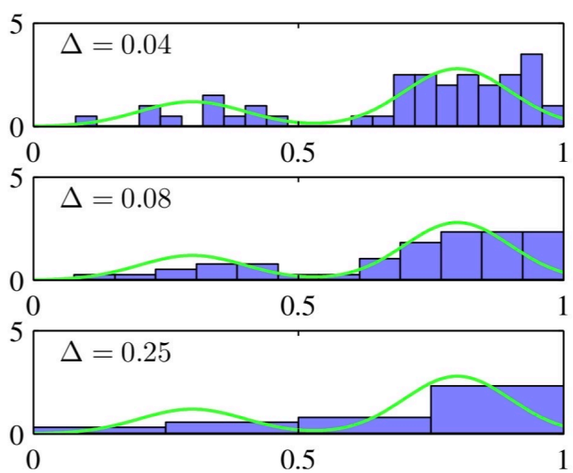

Figure 2.24 An illustration of the histogram approach to density estimation, in which a data set of 50 data points is generated from the distribution shown by the green curve. Histogram density estimates, based on (2.241), with a common bin width Δ are shown for various values of Δ.

Remark

- Effect of a choice of width Δ (smoothing parameter)
- After compting histogram, the data set can be discarded.
- Useful tool for a quick visualization of 1-d or 2-d data
- Limitation of high dimensional data M^D, M bins in D-dim

2.5.1 Kernel density estimator

Estimate unknown probability density $p(x)$ in D-dim space.

Consider some small region R containing x.

Then the probability mass associated with this R

$$P = \int_R p(x)\,dx \qquad \text{(true prob.)}$$

Suppose we observed N data set drawn from $p(x)$.

Since each point has a probability P of falling within R

$$\text{Bin}(K \mid N, P) = \binom{N}{K} P^K (1-P)^{N-K} \qquad K = 0, 1, \ldots N$$

$$\Rightarrow \quad \mathbb{E}[K/N] = p, \qquad \mathrm{var}[K/N] = p(1-p)/N$$

For large N,

$$K \simeq Np$$

If the region R is sufficiently small that $p(x)$ is roughly constant over R, then

$$p \simeq p(x)V$$

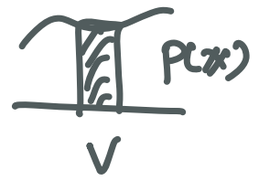

where V is the volume of R.

Combining these expressions we obtain the density estimate

$$p(x) = \frac{K}{NV} \qquad (2.246)$$

R 알아서

Remark: two contradictory assumptions on R and K

We can exploit (2.246) in two different ways

- K nearest neighbour method (fix K)
- kernel density estimator (fix V)

Kernel method in detail

Take the region R to be a small hypercube centered on x

To count the number k of points falling within R, define

$$k(u) := \begin{cases} 1 & |u_i| \le \frac{1}{2} \quad i = 1, \dots D \\ 0 & \text{otherwise} \end{cases}$$

k(u) is an example of kernel function. i.e. the quantity

$k((x - x_n)/h) = 1$ if x_n lies in a cube of side h centered on x otherwise 0.

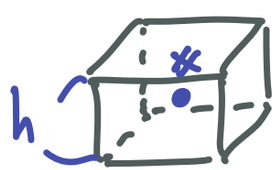

The total number of points lying in the cube

$$K := \sum_{n=1}^{N} k\left(\frac{x - x_n}{h}\right)$$

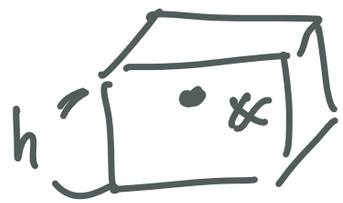

Substituting this expression into (2.246),

$$p(x) = \frac{1}{N} \sum_{n=1}^{N} \frac{1}{h^D} k\left(\frac{x - x_n}{h}\right)$$

where we have used $V = h^D$. (example of kernel density estimator)

We can obtain a smoother density model (Gaussian kernel)

$$p(x) = \frac{1}{N} \sum_{n=1}^{N} \frac{1}{(2\pi h^2)^{D/2}} \exp\left\{ -\frac{\|x - x_n\|^2}{2h} \right\}$$

where h represents the standard deviation of Gaussian component and plays the role of a smoothing parameter.

Generally, we can choose any other kernel function $k(u)$ subject to

$$k(u) \geq 0$$

$$\int k(u)\, du = 1$$

Remark

- No computation involved in the training phase
- Computational cost grows linearly with the data size.

Figure 2.25 Illustration of the kernel density model (2.250) applied to the same data set used to demonstrate the histogram approach in Figure 2.24. We see that h acts as a smoothing parameter and that if it is set too small (top panel), the result is a very noisy density model, whereas if it is set too large (bottom panel), then the bimodal nature of the underlying distribution from which the data is generated (shown by the green curve) is washed out. The best density model is obtained for some intermediate value of h (middle panel).

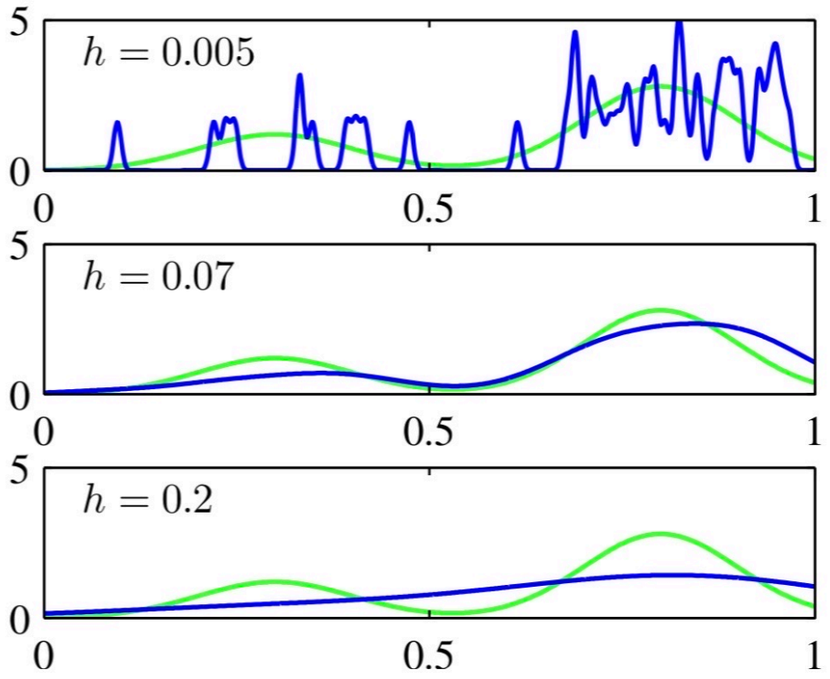

2.5.2 Nearest - neighbour methods

K nearest neighbours

For local density estimation, fix value of K and use the data to find an appropriate value for V.

Consider a sphere centered on x and allow the radius to grow until it contains K data points. i.e. the radius is not determined (fixed)

The value of K governs the degree of smoothing.

Use (2.246) $p(x) = \frac{K}{NV}$ with KNN method for density estimation

KNN method can be extended to classification.

Apply KNN to each class separately and then make use of Baye's theorem.

N : total # of data set, N_k : # of points in C_k

i.e. $\sum_{k=1}^{K} N_k = N$

New point x (fixed). Draw a sphere centered on x containing precisely K points irrespective of their class.

This sphere has the volume V and contains K_k points from the class C_k.

$$p(\mathbf{x} \mid C_k) = \frac{K_k}{N_k V}$$

Similarly, the unconditioned density is given by

$$p(\mathbf{x}) = \frac{K}{N V}$$

and class prior $p(C_k) = N_k / N$

$K = 5$

Combining these equations and using Baye's theorem

$$p(C_k \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid C_k)\, p(C_k)}{p(\mathbf{x})} = \frac{K_k}{K}$$

$p(C_2 \mid \mathbf{x}) = \frac{4}{5}$ $p(C_1 \mid \mathbf{x}) = \frac{1}{5}$

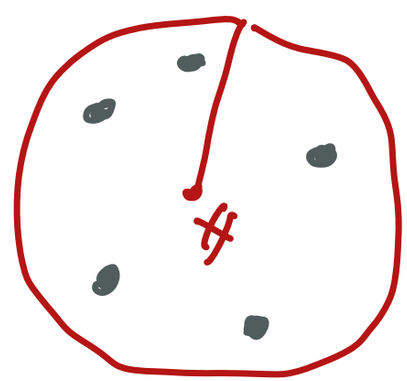

To minimize the probability of misclassification, assign $\mathbf{x}$ to the class having the largest posterior probability $p(C_k|\mathbf{x})$

The particular case of $K=1$ is called nearest-neighbour

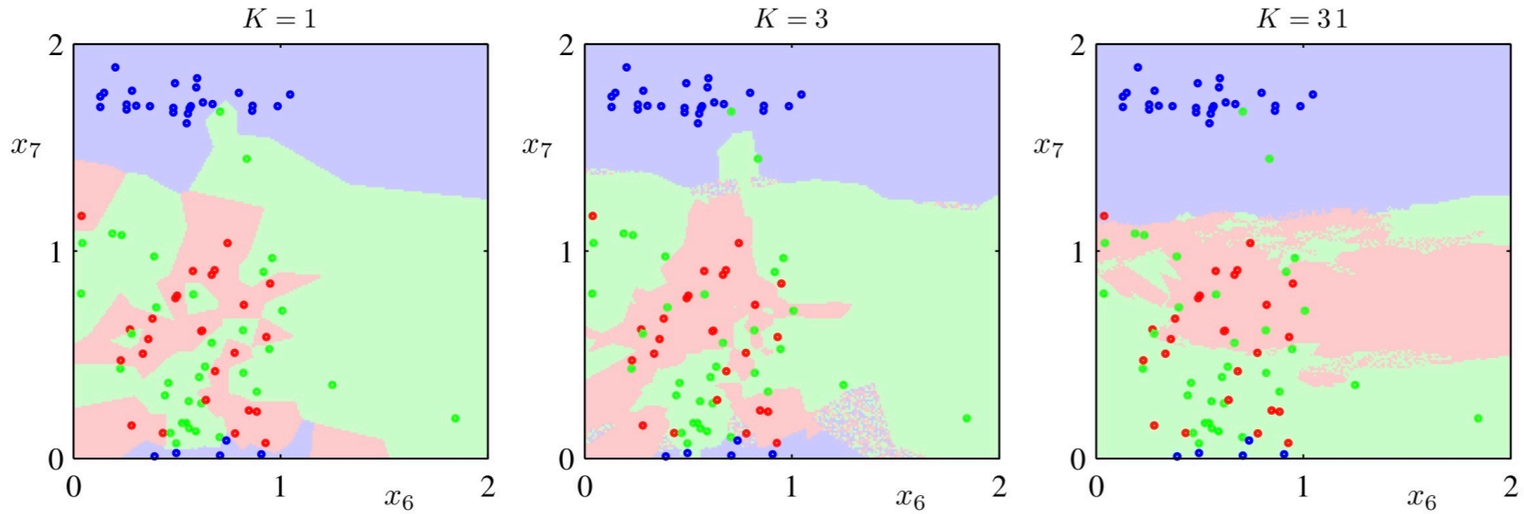

Figure 2.28 Plot of 200 data points from the oil data set showing values of x_6 plotted against x_7, where the red, green, and blue points correspond to the 'laminar', 'annular', and 'homogeneous' classes, respectively. Also shown are the classifications of the input space given by the K-nearest-neighbour algorithm for various values of K.

What happen if $K = N$.

Remark

- K controls the degree of smoothing
- KNN and kernel density methods require the entire data set to stored