

Foundations of Nonparametric Bayesian Methods

Part II

Peter Orbanz

Overview: Today

1. Bayesian models
2. Construction of stochastic processes
3. Extension of conditional probabilities

Conditioning

Direct approach

Conditional probability of $X(\omega) \in A$ given that $X(\omega) \in B$:

$$\mu(A|B) := \frac{\mu(A \cap B)}{\mu(B)}$$

→ no use if $\mu(B) = 0$ (think of Bayesian model on \mathbb{R}^d)

Abstract Conditional Probabilities

Measure-theoretic definition of conditionals is beyond scope of this talk.

In The Following

We ignore technical details and write $\mu(X|Y)$ or $\mu(A|Y)$ for the conditional probability of X (RV) or A (event) given Y .

Conditional Densities

If X, Y have joint density, $\mu(X|Y)$ has conditional density $p(x|y)$.

Bayesian Models

Parametric Family

Let $X : (\Lambda, \mathcal{A}) \rightarrow (\Omega_X, \mathcal{B}_X)$ and $\Theta : (\Lambda, \mathcal{A}) \rightarrow (\Omega_\theta, \mathcal{B}_\theta)$ be two random variables, and $\mu_X = X(\mathbb{P})$. Then the conditional distribution $\mu_X(X|\Theta)$ is called a *parametric family* of distributions (parameterized by $\theta \in \Omega_\theta$).

Bayesian model

If X observed and Θ unobserved, we call:

- ▶ $\mu_\Theta := \Theta(\mathbb{P})$ the *prior measure*
- ▶ $\mu_\Theta(\Theta|X)$ the *posterior measure*
- ▶ The overall model is called a Bayesian model.

Note: Not defined by a Bayes equation!

Bayes' Theorem

Problem:

Given the prior and the data, how can we determine the posterior? (Without exhaustive knowledge of \mathbb{P} , \mathcal{A} etc)

Bayes Theorem

If the sampling model $\mu_X(X|\Theta)$ has density $p_{X|\theta}$, then:

$$\frac{d\mu_{\Theta|X}}{d\mu_{\Theta}}(\theta|x) = \frac{p_{X|\theta}}{\int p_{X|\theta} d\mu_{\Theta}(\theta)}$$

for all x with $\int p_{X|\theta} d\mu_{\Theta}(\theta) \notin \{0, \infty\}$.

Models With No Bayes Equation

For some models (e.g. DP) posterior \ll prior *not* satisfied \rightarrow

Bayesian model, but no Bayes equation.

Bayesian Nonparametrics

Nonparametric Bayesian model

A Bayesian model with:

1. $\dim(\Omega_\theta) = \dim(\Omega_x) = +\infty$.
2. Model can be evaluated on partial observations.

Partial observation

Random quantity with d dimensions, only $m < d$ are observed.

Example: GP regression

GP draw is function f , but only finite number of values of f known.

Stochastic Process Models

Intuition

Stochastic process = ∞ -dim probability distribution

Typical GP definition

“A Gaussian process is a probability distribution on an infinite collection of random variables X_t such that the marginal distribution for each finite subset (t_1, \dots, t_n) of indices is Gaussian.”

→ Existence? Uniqueness?

Stochastic Process Construction (1)

Stochastic process measure μ^E : Distribution of RV

$$X^E : (\Lambda, \mathcal{A}) \rightarrow (\Omega^E, \mathcal{B}^E)$$

- ▶ E : infinite index set (indexes entries of random vector)
- ▶ Ω_0 : “one-dimensional” sample space
- ▶ $\Omega^E := \prod_{i \in E} \Omega_0$
- ▶ Interpretation: μ^E -draws = mappings $x : E \rightarrow \Omega_0$

Projector

P_{JI} := projection mapping $\Omega^J \rightarrow \Omega^I$ (for $I \subset J \subset E$)

Marginals

Marginal of μ^J on $\Omega^I \subset \Omega^J$:

$$\underbrace{(P_{JI}\mu^J)}_{\text{on } \Omega^I}(A) := \underbrace{\mu^J(P_{JI}^{-1}A)}_{\text{on } \Omega^J}$$

marginals = projections of measures

Stochastic Process Construction (2)

Def: Projective family

Family $\{\mu^I | I \subset E \text{ finite}\}$ such that for all finite I, J with $I \subset J$:

$$P_{II} \mu^J = \mu^I$$

Note: If μ^E given, the finite-dim marginals $\mu^I := P_{EI} \mu^E$ are a projective family.

Kolmogorov's Extension Theorem

If a family $\{\mu^I | I \subset E \text{ finite}\}$ of finite-dimensional measures is projective, there exists a unique measure μ^E on Ω^E with μ^I as its marginals.

Jargon: μ^E is called the *projective limit* of the μ^I .

Example: GP construction

Choice of components

- ▶ $\Omega_0 := \mathbb{R}$ and index set $E = \mathbb{R}$
- ▶ $P_{|I|}$: Euclidean projector from $\mathbb{R}^{|J|}$ to $\mathbb{R}^{|I|}$.
- ▶ Marginal family: μ^I are $|I|$ -dimensional Gaussians

Ensure marginals projective

- ▶ Start with mean function $m(\cdot)$ and covariance $k(\cdot, \cdot)$.
- ▶ Note: $E = \mathbb{R}$, finite $I = \{t_1, \dots, t_{|I|}\} \subset \mathbb{R}$
- ▶ $\mu^I =$ Gaussian, mean $(m(t_1), \dots, m(t_{|I|}))$ and $\Sigma_{ij} = k(t_i, t_j)$

Apply Extension Theorem

GP measure μ^E exists and is unique.

Note: μ^E has mean m and covariance function k , but that is *not* an immediate consequence of theorem!

Extensions Theorem: Caveat

Problem

If dimension E is uncountable, the projective limit measure μ^E is basically useless.

Explanation

- ▶ Domain of μ^E : \mathcal{B}^E (generated by product topology)
- ▶ Sets in \mathcal{B}^E : “axes-parallel” in all but countably many dimensions
- ▶ E uncountable $\rightarrow \mathcal{B}^E$ too coarse for meaningful modeling

A Note of Caution:

Problem is often neglected in literature.

Example: Original paper on the DP (Ferguson, 1973).

Uncountable Dimensions

Intuition:

Objects of interest *effectively* have countably many degrees of freedom.

Examples

- ▶ **Continuous functions:** Completely defined by values on dense subset (e.g. \mathbb{Q} in \mathbb{R})
- ▶ **Probability measures:** Completely defined by values on countable system of sets.

Strategies

1. Modify theorem to directly define measure on “interesting” space (eg space of continuous functions).
2. Use Kolmogorov theorem, then restrict μ^E to interesting subspace.

Extension of Conditional Probabilities

Motivation

Bayesian estimation deals with conditional probabilities or parametric families, rather than individual distributions.

Extension Result

Assumptions:

- ▶ E countable
- ▶ Conditionals on subspaces Ω^I satisfy

$$\mu^J(\mathbf{P}_{JI}^{-1} \cdot | \Theta^J) = \mu^I(\cdot | \Theta^I) \quad \text{for } I \subset J$$

Then there is a conditional distribution $\mu^E(X^E | \Theta^E)$ on Ω^E with marginals $\mu^I(\cdot | \Theta^I)$.

Disclaimer

Result statement above neglects some technical details.

Conjugate Models

Definition 1

A likelihood and a family of priors are *conjugate* if all possible posteriors are elements of the prior family. (“Closure under sampling”)

Definition 2

Likelihood and prior family are conjugate if there exists a measurable mapping of the form

$$\text{Prior parameters} \times \text{Data} \rightarrow \text{Posterior parameters}$$

In Exponential Family Models

Mapping T to posterior parameters:

$$(\lambda, \mathbf{y}) \xrightarrow{T} \left(\lambda + n, \mathbf{y} + \sum_{i=1}^n \mathbf{S}(\mathbf{x}_i) \right)$$

Conjugate Projective Limits

Extension Result: In Short

If mappings to posterior parameters satisfy projection relation, they define corresponding mapping for projective limit model.

In Detail

- ▶ $T^l(x^l, y^l)$ mappings to posterior parameters
- ▶ Fix y^E and write $T_y^l = T^l(\cdot, P_{EI}y^E)$

If some mapping T^E satisfies

$$P_{EI}^{-1} \circ T_y^{l,-1} = T_y^{E,-1} \circ P_{EI}^{-1}$$

then T^E defines functional conjugacy for limit model.

For Exponential Family Marginals

If S^E sufficient for extension:

$$(\lambda, y^E) \xrightarrow{T^E} (\lambda + n, y^E + \sum_i S^E(x_i^E))$$

Projective Limits of Bayes Equations

$$\begin{array}{ccc} \mu^E(\Theta^E | X^E, Y^E) & \xleftarrow[\mathcal{T}^E]{x_1^E, \dots, x_n^E} & \mu^E(\Theta^E | Y^E) \\ \text{P}_{EI} \downarrow & \lim \leftarrow & \uparrow \text{P}_{EI} \\ \mu^I(\Theta^I | X^I, Y^I) & \xleftarrow[\mathcal{T}^I]{x_1^I, \dots, x_n^I} & \mu^I(\Theta^I | Y^I) \end{array}$$

Model Constructions

Example Models

Marginals

Bernoulli/beta

Multin./Dirichlet

Gauss/Gauss

Mallows/conj.

Proj Limit

IBP/beta process

CRP/DP

GP/GP

(exists)

Observations

Binary arrays

Discrete dist.

cont. functions

Bijections of \mathbb{N}

Construction Recipe

- ▶ Choose finite-dimensional observation (eg permutations)
- ▶ Choose exponential family model on observations
- ▶ Choose canonical conjugate prior
- ▶ Check: Model and sufficient statistic projective

Warning: More difficult in uncountable dimensions.