



10-708 Probabilistic Graphical Models

Machine Learning Department
School of Computer Science
Carnegie Mellon University



Bayesian Nonparametrics: Gaussian Process

Matt Gormley
Lecture 24
Apr. 30, 2021

Reminders

- **Cloud Credits (AWS or GCP)**
 - first request deadline: Thu at 11:59pm
- **Quiz 3**
 - Mon, May 3 during lecture slot
 - Topics: Lectures 16 - 23

QUIZ 2 LOGISTICS

Quiz 2

- **Time / Location**
 - **Time:** In-Class Quiz
Wed, Apr. 14 during lecture time
 - **Location:** The same Zoom meeting as lecture/recitation.
Please arrive online early.
 - Please watch Piazza carefully for announcements.
- **Logistics**
 - Covered material: Lecture 9 – Lecture 15
(and unavoidably some material from Lectures 1 – 8)
 - Format of questions:
 - Multiple choice
 - True / False (with justification)
 - Derivations
 - Short answers
 - Interpreting figures
 - Implementing algorithms on paper
 - Drawing
 - No electronic devices
 - You are allowed to **bring** one 8½ x 11 sheet of notes (front and back)

Quiz 2

- **Advice (for before the exam)**
 - Try out the Gradescope quiz-style interface in the “Fake Quiz” now available
- **Advice (for during the exam)**
 - Solve the easy problems first (e.g. multiple choice before derivations)
 - if a problem seems extremely complicated you’re likely missing something
 - Don’t leave any answer blank!
 - If you make an assumption, write it down
 - If you look at a question and don’t know the answer:
 - we probably haven’t told you the answer
 - but we’ve told you enough to work it out
 - imagine arguing for some answer and see if you like it

Topics for Quiz 1

- Graphical Model Representation
 - Directed GMs vs. Undirected GMs vs. Factor Graphs
 - Bayesian Networks vs. Markov Random Fields vs. Conditional Random Fields
- Graphical Model Learning
 - Fully observed Bayesian Network learning
 - Fully observed MRF learning
 - Fully observed CRF learning
 - Parameterization of a GM
 - Neural potential functions
- Exact Inference
 - Three inference problems:
 - (1) marginals
 - (2) partition function
 - (3) most probably assignment
 - Variable Elimination
 - Belief Propagation (sum-product and max-product)

Topics for Quiz 2

- Learning for Structure Prediction
 - Structured Perceptron
 - Structured SVM
 - Neural network potentials
- (Approximate) MAP Inference
 - MAP Inference via MILP
 - MAP Inference via LP relaxation
- Approximate Inference by Sampling
 - Monte Carlo Methods
 - Gibbs Sampling
 - Metropolis-Hastings
 - Markov Chains and MCMC
- Parameter Estimation
 - Bayesian inference
 - Topic Modeling

Topics for Quiz 3

- Approximate Inference by Optimization
 - Variational Inference
 - Mean Field Variational Inference
 - Coordinate Ascent V.I. (CAVI)
 - Variational EM
 - Variational Bayes
- Deep Generative Models
 - Variational Autoencoders
 - Sigmoid Belief Networks, Restricted Boltzmann Machines, Deep Belief Nets, Deep Boltzmann Machines
- Bayesian Nonparametrics
 - Dirichlet Process
 - DP Mixture Model
 - Indian Buffet Process


GAUSSIAN PROCESS

Motivation: Gaussian Process

Density Estimation

- Given data, estimate a probability density function that best explains it
- A nonparametric prior can be placed over an infinite set of distributions

Prior:




Red: mean density. Blue: median density. Grey: 5-95 quantile.
Others: draws.

Motivation: Gaussian Process

Density Estimation

- Given data, estimate a probability density function that best explains it
- A nonparametric prior can be placed over an infinite set of distributions

Posterior:



Red: mean density. Blue: median density. Grey: 5-95 quantile.
Black: data. Others: draws.

Gaussian Process

Whiteboard:

- Parametric vs. Nonparametric learning
- High level idea of GP regression
- GP Regression
 - Example prior
 - Strawman inference algorithm
 - Example posterior
- GP Classification
 - approximate inference
 - Example posterior

Background: Multivariate Gaussians

Whiteboard:

- Marginal of multivariate Gaussian
- Conditional of multivariate Gaussian

Gaussian Process Regression

Whiteboard:

- Function-space view
 - definition of Gaussian Process
 - mean function
 - covariance function
- Example kernels
- Weight-space view
 - linear regression (linear model + Gaussian noise)
 - ridge regression (adding a Gaussian prior)
 - Bayesian linear regression
 - Bayesian kernel regression (aka. GP Regression)


Gaussian Process Regression

Whiteboard:

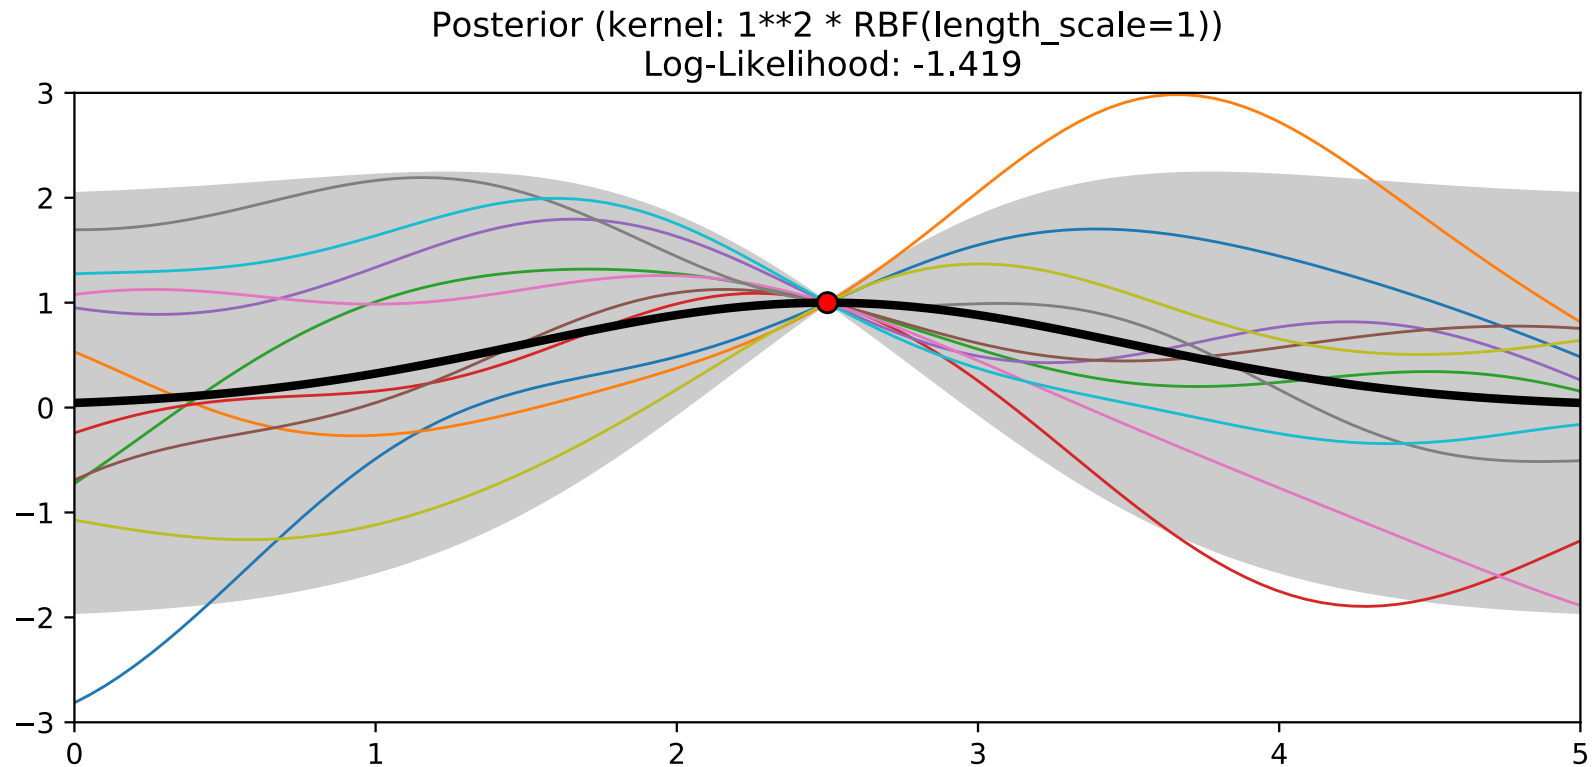
- MBR Decoding
- Computational complexity

GAUSSIAN PROCESS INFERENCE

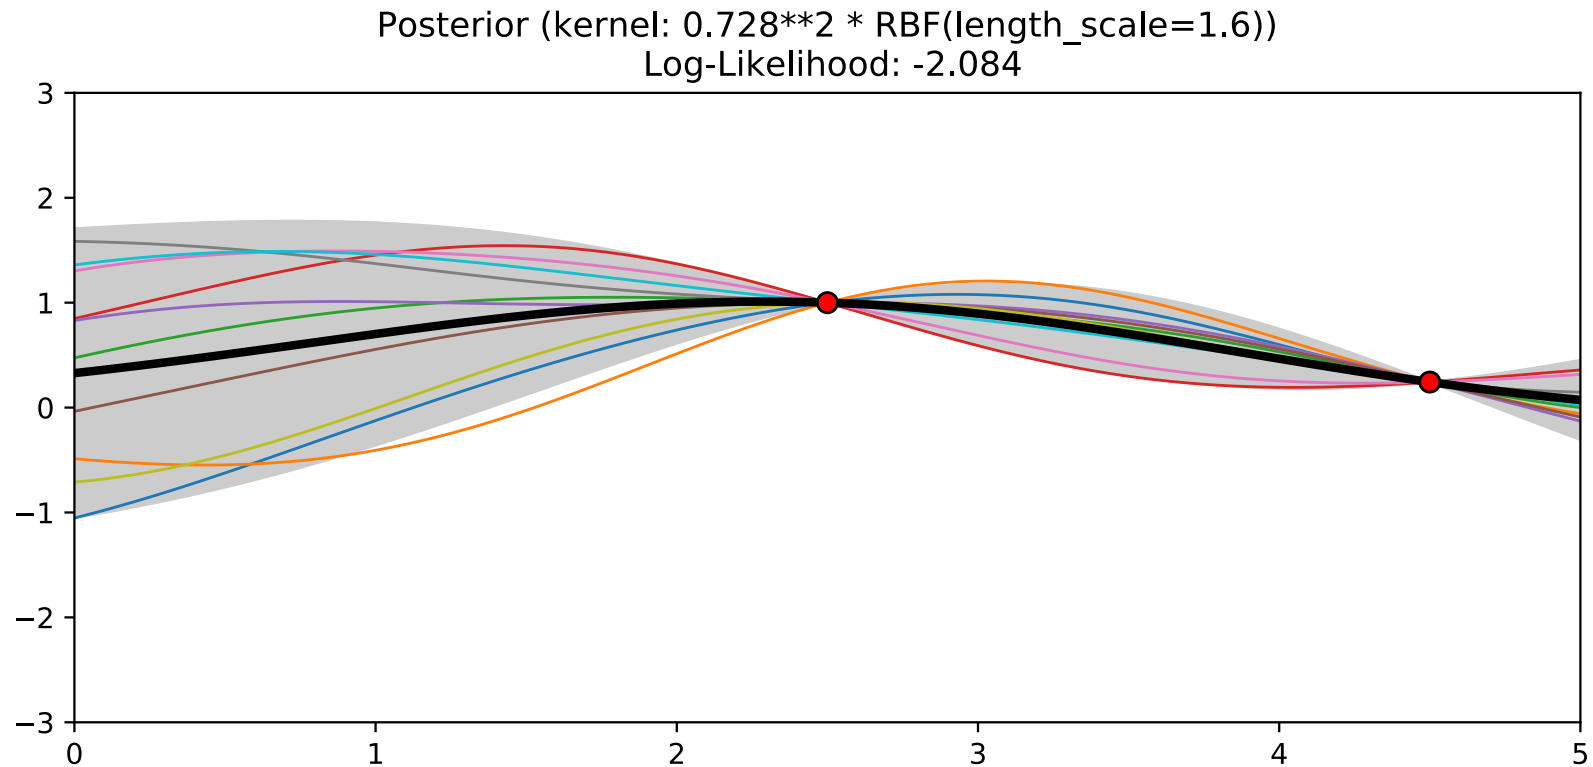
Gaussian Process Example



Gaussian Process Example



Gaussian Process Example



Gaussian Process Example

