## Statistical methods in Data Science and AI

Marina Axelson-Fisk 30 september, 2019

## Inference in graphical models

Given a graphical model, we want to answer questions of interest.

- Marginal inference: what is the marginal probability of a given variable $Y$ in our graph, summing out the rest?

$$
P(Y=y)=\sum_{x_{1}} \sum_{x_{2}} \ldots \sum_{x_{n}} P\left(Y=y, X_{1}=x_{1}, X_{2}=x_{2}, \ldots, X_{n}=x_{n}\right)
$$

- Maximum a posteriori (MAP) inference: what is the most likely assignment to the variables in the graph (possibly conditioned on data)?


$$
\max _{x_{1}, \ldots, x_{n}} P\left(Y=y, X_{1}=x_{1}, \ldots, X_{n}=x_{n}\right)
$$

## Inference algorithms in graphical models

## Exact inference

- Variable elimination
- Message passing/belief propagation
- Junction trees

Approximative inference

- Stochastic simulation
- Markov chain Monte Carlo (MCMC)
- Variational algorithms



## Example: variable elimination in a chain graph

Random variables: $A, B, C, D, E$

each taking $n$ possible values $\{1,2, \ldots, n\}$, then the marginal probability of $E$

$$
P(E=e)=\sum_{a=1}^{n} \sum_{b=1}^{n} \sum_{b=1}^{n} \sum_{b=1}^{n} P(A=a, B=b, C=c, D=d, E=e)=\sum_{a, b, c, d} P(a, b, c, d)
$$

we can utilize the chain structure to reduce the number of operations by variable elimination.

## Example: variable elimination in a chain graph

Exploit the structure "inside-out" or from "leaf-to-top"

$$
\begin{gathered}
P(E=e)=\sum_{a} \sum_{b} \sum_{c} \sum_{d} P(A=a, B=b, C=c, D=d, E=e) \\
=\sum_{a} \sum_{b} \sum_{c} \sum_{d} P(a) P(b \mid a) P(c \mid b) P(d \mid c) P(e \mid d) \\
=\sum_{b} \sum_{c} \sum_{d} P(c \mid b) P(d \mid c) P(e \mid d) \sum_{a} P(b \mid a) P(a) \\
=\sum_{b} \sum_{c} \sum_{d} P(c \mid b) P(d \mid c) P(e \mid d) P(b)=\cdots= \\
=\sum_{d} P(e \mid d) P(d)=P(e)
\end{gathered}
$$

## Message passing/belief propagation

- Variable elimination can be seen as "passing a message" (information), or "propagating a belief" from one node to the next.
- This is the basic framework for computing various entities in Hidden Markov Models (HMMs) and Linear dynamical systems (LDS)
- In continuous distributions $p(x \mid \theta)$ message passing corresponds to passing on parameter values $\theta$ between neighboring nodes.


## Inference on trees

- Each node sends out the product of the messages received from the parents to the children
- l.e. message passing is an abstract notion of conditional (in)dependence

undirected tree




## Inference on trees

- The junction tree algorithm (or clique tree algorithm) is a generalization of message passing to arbitrary graphs



## Module 4.1: Kernel methods

## Kernel methods: motivation

- Given a training set $\mathcal{D}=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{N}$
- $y_{i}$ response
- $\mathbf{x}_{i}$ feature vector
- there are numerous tools for detecting linear relations
- Ridge regression
- Support vector machines (SVMs)
- Principal component analysis (PCA)

- But what if the relationship is nonlinear?


## Motivating example: binary classification

- Training set: $\mathcal{D}=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{N}, y_{i} \in\{-1,+1\}$
- Objective: learn a function $f(\mathbf{x})$ such that $y_{i}=\operatorname{sign}\left(f\left(\mathbf{x}_{i}\right)\right)$



## Linear classification

When classes are linearly separable, the boundary is a hyperplane.

$$
\mathbf{w}^{\mathrm{T}} \mathbf{x}+b=0
$$

If $\mathbf{w}^{\mathrm{T}} \mathbf{x}+b \begin{cases}>0 & y=\text { red } \\ <0 & y=\text { green }\end{cases}$

> But which line should we choose?


## Linear classification

- How large margins do we have between the classes?
- How do we maximize that margin?



## How do we maximize the margin?

- Let $\mathbf{x}_{n} \in \mathcal{D}$ be the point closest to the hyperplane $f_{\mathbf{w}}(\mathbf{x})=\mathbf{w}^{T} \mathbf{x}+b=0$
- Normalize w:

$$
\left|\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}\right|=1 \quad \Rightarrow \text { The canonical hyperplane }
$$

- The distance between $x_{n}$ and the plane

$$
\text { distance }=\frac{1}{\|\mathbf{w}\|}\left|\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}-\mathbf{w}^{\mathrm{T}} \mathbf{x}\right|=\frac{1}{\|\mathbf{w}\|}
$$

- Maximize this distance.



## Support vector machines (SVMs)

For two linearly separable classes with class labels $y_{i} \in\{-1,+1\}$

- construct two supporting hyperplanes, one for each class

$$
\begin{aligned}
& \mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}+b \leq-1 \text { for } y_{i}=-1 \\
& \mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}+b \geq+1 \text { for } y_{i}=+1
\end{aligned}
$$

The corresponding supporting hyperplanes are thus

$$
\begin{aligned}
& H_{-1}=\left\{\mathbf{x}_{i}: \mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}+b=-1, y_{i}=-1\right\} \\
& H_{+1}=\left\{\mathbf{x}_{i}: \mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}+b=+1, y_{i}=+1\right\}
\end{aligned}
$$



## Support vector machines (SVMs)

In two dimensions the separating hyperplanes are lines on the form

$$
f_{\mathbf{w}}\left(\mathbf{x}_{i}\right)=\mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}+b=w_{1} x_{1}+w_{2} x_{2}+b=0
$$

Maximizing the margin = minimizing $\|\mathrm{w}\|$ under the constraints

$$
\begin{aligned}
& \mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}+b \leq-1 \text { for } y_{i}=-1 \\
& \mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}+b \geq+1 \text { for } y_{i}=+1
\end{aligned}
$$

or

$$
\mathbf{w}^{*}=\underset{\mathbf{w}, b}{\arg \min }\left\{\frac{1}{2}\|\mathbf{w}\|^{2}: y_{i}\left(\boldsymbol{w}^{\mathrm{T}} \mathbf{x}_{i}+b\right) \geq 1\right\}
$$



## using Lagrange multipliers.

## Lagrange multipliers

- We want to optimize a function $f(x)$ subject to a constraint $g(x)=0$.
- We form the Lagrangian function

$$
L(x, \lambda)=f(x)-\lambda g(x)
$$

where $\lambda$ is a Lagrange multiplier.

- We optimize by computing

$$
\frac{d L}{d x}=0 \text { and } \frac{d L}{d \lambda}=0
$$

and solve the corresponding equation system.


## Support vector machines (SVMs)

Lagrangian function

$$
L(\mathbf{w}, b, \lambda)=\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{i=1}^{N} \lambda_{i}\left(y_{i}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}-b\right)\right)
$$

- $\boldsymbol{\lambda}=\left(\lambda_{1}, \ldots, \lambda_{N}\right)$ : Lagrange multipliers
- $N$ : number of constraints

$$
\begin{array}{r}
\frac{d L}{d w_{i}}=0, i=1, \ldots, N \text { and } \frac{d L}{d b}=0 \\
\Rightarrow \mathbf{w}=\sum_{i=1}^{N} \lambda_{i} y_{i} \mathbf{x}_{i} \text { and } \sum_{i=1}^{N} \lambda_{i} y_{i}=0
\end{array}
$$

$$
\mathbf{w}^{*}=\underset{\mathbf{w}, b}{\arg \min }\left\{\frac{1}{2}\|\mathbf{w}\|^{2}: y_{i}\left(\boldsymbol{w}^{\mathrm{T}} \mathbf{x}_{i}+b\right) \geq 1\right\}
$$



## Kernelized SVMs

## Dual representation

Instead we maximize

$$
\mathbf{w}=\sum_{i=1}^{N} \lambda_{i} y_{i} \mathbf{x}_{i} \text { and } \sum_{i=1}^{N} \lambda_{i} y_{i}=0
$$

$$
\tilde{L}(\boldsymbol{\lambda})=\sum_{i=1}^{N} \lambda_{i}-\frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \lambda_{i} \lambda_{j} y_{i} y_{j} k\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)
$$

under the constraints

$$
\lambda_{i} \geq 0 \text { and } \sum_{i=1}^{N} \lambda_{i} y_{j}=0
$$

where $k\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)=\mathbf{x}_{i}^{\mathrm{T}} \mathbf{x}_{j}$ is a kernel function.


## CHALMERS

## Kernelized SVMs

To classify a new data point $\mathrm{x}^{*}$ we observe the sign of

$$
f\left(\mathbf{x}^{*}\right)=\sum_{i=1}^{N} \lambda_{i} y_{i} k\left(\mathbf{x}^{*}, \mathbf{x}_{i}\right)+b
$$



## Soft margin SVMs

When the training set cannot be perfectly separated we introduce slack variables

$$
y_{i}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}+b\right) \geq 1-\xi_{i}
$$

such that $\xi_{i} \leq 1$ for points on the correct side.
We want to minimize the misclassification rate, i.e. minimize

$$
\begin{aligned}
& \quad \sum_{i=1}^{N} \mathbb{I}\left\{\xi_{i}-1\right\} \text { where } \mathbb{I}(x)= \begin{cases}1 & \text { if } x \leq 0 \\
0 & \text { if } x>0\end{cases} \\
& \text { NP-complete! }
\end{aligned}
$$



## Soft margin SVMs

Use the upper bound

$$
\sum_{i=1}^{N} \mathbb{I}\left\{\xi_{i}-1\right\} \leq \sum_{i=1}^{N} \xi_{i}
$$

$C$ is a trade-off between misclassification and complexity


## Kernelized soft margin SVMs

Lagrangian function

$$
L(\mathbf{w}, b, \boldsymbol{\lambda})=\frac{1}{2}\|\mathbf{w}\|^{2}+C \sum_{i=1}^{N} \xi_{i}-\sum_{i=1}^{N} \lambda_{i}\left(y_{i}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{i}-1+\xi_{i}\right)\right)-\sum_{i=1}^{N} \mu_{i} \xi_{i}
$$

with Lagrange multipliers
$\boldsymbol{\lambda}=\left(\lambda_{1}, \ldots, \lambda_{N}\right), \boldsymbol{\mu}=\left(\mu_{1}, \ldots, \mu_{n}\right)$


## Kernelized soft margin SVMs

The dual form is the same as before

$$
\tilde{L}(\boldsymbol{\lambda})=\sum_{i=1}^{N} \lambda_{i}-\frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \lambda_{i} \lambda_{j} y_{i} y_{j} k\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)
$$

but the maximization constraints become limited by $C$
$0 \leq \lambda_{i} \leq C$ and $\sum_{i=1}^{N} \lambda_{i} y_{i}=0$


## Nonlinear classification



- There is no linear classifier that can separate red from green.


## Kernel methods: motivation

- Solution:
- map the data into a (possibly high-dimensional) vector space where the relation becomes linear
- apply the linear algorithm in this space



## Nonlinear classification

- There is no linear classifier that can separate red from green.
- However, the following function can separate the regions perfectly

- By mapping $x$ to feature space $\phi(x)=\left(x^{2}, 1\right) \in \mathbb{R}^{2}$ the nonlinear problem has become linear.


## Nonlinear classification

- There is no linear classifier that can separate red from green.
- However, a conic section separates them perfectly


$$
\left.\begin{array}{rl}
f(\mathbf{x}) & =a x_{1}^{2}+b x_{1} x_{2}+c x_{2}^{2}+d x_{1}+e x_{2}+g \\
& =\langle\underbrace{(a, b, c, d, e, g}_{\mathbf{w}^{\mathrm{T}}},
\end{array}\right) \underbrace{\left.x_{1}^{2}, x_{1} x_{2}, x_{2}^{2}, x_{1}, x_{2}, 1\right)}_{\phi(\mathbf{x})}\rangle, \quad \phi \in \mathbb{R}^{6})
$$

## Kernel methods: motivation

- Problem:
- computationally difficult to represent data in high dimensions




## Kernel methods: motivation

- Alternative:
- compute similarity measure between vectors in feature space
- apply algorithms based on similarity measures



## Kernel definition

- For input vectors $\mathrm{x}, \mathrm{z} \in \mathcal{X}$ and a mapping $\phi: \mathcal{X} \rightarrow \mathbb{R}^{N}$

$$
k(\mathbf{x}, \mathbf{z})=\phi(\mathbf{x})^{\mathrm{T}} \phi(\mathbf{z})
$$

is a kernel function.

- The kernel trick: we don't need the coordinates of the data in feature space. Just the inner product between vectors.


## Example



Kernel: $k(\mathbf{x}, \mathbf{z})=\mathbf{x}^{\mathrm{T}} \mathbf{z}=\left(x_{1} z_{1}+x_{2} z_{2}\right)^{2}=x_{1}^{2} z_{1}^{2}+2 x_{1} z_{1} x_{2} z_{2}+x_{2}^{2} z_{2}^{2}=$

$$
=\left(x_{1}^{2}, \sqrt{2} x_{1} x_{2}, x_{2}^{2}\right)\left(z_{1}^{2}, \sqrt{2} z_{1} z_{2}, z_{2}^{2}\right)^{\mathrm{T}}=\phi(\mathbf{x})^{\mathrm{T}} \phi(\mathbf{z})
$$

where $\phi(\mathbf{x})=\left(x_{1}^{2}, \sqrt{2} x_{1} x_{2}, x_{2}^{2}\right)^{\mathrm{T}}$ is the nonlinear feature mapping.

## Kernel functions

## Another view:

- a kernel $k(\mathbf{x}, \mathrm{z})$ is a measure of similarity between vectors $\mathrm{x}, \mathrm{z} \in \mathcal{X}$ where $\mathcal{X}$ is some abstract space.
- or simply a distance measure between points in feature space



## Kernel functions

However, a feature space is not unique to a given kernel:
$k(\mathbf{x}, \mathbf{z})=\langle\mathbf{x}, \mathbf{z}\rangle^{2}$ is also the kernel that computes the inner product of the map

$$
\psi\left(x_{1}, x_{2}\right)=\left(x_{1}^{2}, x_{2}^{2}, x_{1} x_{2}, x_{2}, x_{1}\right) \in \mathbb{R}^{4}
$$

Moreover, every prospective kernel does not necessarily correspond to a dot product in some space.


## Kernel functions

Mercer's condition:
A continuous, symmetric, positive semi-definite kernel function can be written as a dot product of vectors in a higher dimension

$$
k(\mathbf{x}, \mathbf{z})=\phi(\mathbf{x})^{\mathrm{T}} \phi(\mathbf{z})
$$

Positive semi-definite: a symmetric matrix with positive eigenvalues


## Let's summarize

- We have a data set $\mathcal{D}=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{N}$ where $\mathbf{x} \in \mathbb{R}^{D}$ and $y_{i} \in\{-1,+1\}$
- We want a nonlinear projection $\phi(x)$ onto a higher dimension

$$
\phi: \mathbb{R}^{D} \rightarrow \mathbb{R}^{D+d}, \quad d>0
$$

where classes have a better chance of being linearly separable.

Cover's theorem (informally):
"A nonlinear projection in a high-dimensional space is more likely to be linearly separable than in a low-dimensional space"


## Let's summarize

- The separating hyperplane in $\mathbb{R}^{D+d}$ will be defined by

$$
\sum_{j=1}^{D+d} w_{j} \phi_{j}(\mathbf{x})+b=\mathbf{w}^{\mathrm{T}} \phi(\mathbf{x})+b=0
$$

- The optimal hyperplane is given by

$$
\mathbf{w}=\sum_{i=1}^{N} \lambda_{i} y_{i} \phi\left(\mathbf{x}_{i}\right) \Leftrightarrow \sum_{i=1}^{N} \lambda_{i} y_{i} k\left(\mathbf{x}_{i}, \mathbf{x}\right)=0
$$

where $k\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)=\phi\left(\mathbf{x}_{i}\right)^{\mathrm{T}} \phi\left(\mathbf{x}_{j}\right)$ is a kernel function and $\lambda_{i}$ Lagrange multipliers.


## How to choose the mapping $\phi(x)$ ?

- Choosing an optimal feature space is non-trivial
- The kernel trick reduces this to choosing the best kernel, and determine the corresponding implicit mapping $\phi(x)$.
- Performance of the algorithm highly depends on the kernel
- The best kernel depends on the specific problem
- Kernels can be applied to
- Numeric vectors
- Strings
- Trees
- Graphs



## How to choose the best kernel

We want the kernel to be

- Valid: an implicit mapping must exist
$=$ a kernel that can be expressed as the dot product of two vectors
= satisfy the Mercer's condition of positive semi-definiteness
- Accurate: embody the "true" similarity between objects
- Appropriate: generalize outside training data
- Efficient: computations should be feasible


## Which kernels meet Mercer's condition?

- Linear kernels

$$
k(\mathbf{x}, \mathbf{z})=\mathbf{x}^{\mathrm{T}} \mathbf{z}
$$

- Polynomial kernels

$$
k(\mathbf{x}, \mathbf{z})=\left(1+\mathbf{x}^{\mathrm{T}} \mathbf{z}\right)^{n}
$$

- Radial basis function (RBF) kernels

$$
k(\mathbf{x}, \mathbf{z})=\exp \left(-\frac{1}{2}\|\mathbf{x}-\mathbf{z}\|^{2}\right)
$$



## Radial basis function (RBF) kernels

- The squared exponential (SE) or Gaussian kernel

$$
k(\mathbf{x}, \mathbf{z})=\exp \left(-\frac{1}{2}(\mathbf{x}-\mathbf{z})^{\mathrm{T}} \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\mathbf{z})\right)
$$

- If the covariance matrix $\Sigma$ is diagonal, we get

$$
k(\mathbf{x}, \mathbf{z})=\exp \left(-\frac{1}{2} \sum_{j} \frac{1}{\sigma^{2}}\left(x_{j}-z_{j}\right)^{2}\right)
$$



- If $\boldsymbol{\Sigma}$ is spherical

$$
k(\mathbf{x}, \mathbf{z})=\exp \left(-\frac{1}{2 \sigma^{2}}\|\mathbf{x}-\mathbf{z}\|^{2}\right)
$$



## Kernels for comparing text documents

- If $x_{i j}=$ the number of times word $j$ occurs in document $i$

$$
k(\mathbf{x}, \mathbf{z})=\frac{\mathbf{x}^{\mathrm{T}} \mathbf{z}}{\|\mathbf{x}\|_{2}\|\mathbf{z}\|_{2}}=\cos (\theta)
$$

called the cosine similarity.


## CHALMERS

## Matern kernel

- Commonly used in Gaussian processes

$$
k(r)=\frac{2^{1-v}}{\Gamma(v)}\left(\frac{\sqrt{2 v} r}{l}\right)^{v} B_{v}\left(\frac{\sqrt{2 v} r}{l}\right) \rightarrow \text { SE kernel as } v \rightarrow \infty
$$

where $r=\|\mathbf{x}-\mathbf{z}\|, v \geq 0, l>0$ and $B_{v}$ a modified Bessel function.

$$
k(r)=\exp (-r / l) \text { when } v=1 / 2
$$



## String kernels

- Consider two strings x and z of lengths $D_{\mathrm{x}}$ and $D_{z}$, defined on a protein alphabet
- $\mathcal{A}=\{\mathrm{A}, \mathrm{R}, \mathrm{N}, \mathrm{D}, \mathrm{C}, \mathrm{E}, \mathrm{Q}, \mathrm{G}, \mathrm{H}, \mathrm{I}, \mathrm{L}, \mathrm{K}, \mathrm{M}, \mathrm{F}, \mathrm{P}, \mathrm{S}, \mathrm{T}, \mathrm{W}, \mathrm{Y}, \mathrm{V}\}$ where
$\mathbf{x}\left(D_{\mathbf{x}}=110\right)$ :
IPTSALVKETLALLSTHRTLLIANETLRIPVPVHKNHQLCTEEIFQGIGTL ESQTVQGGTVERLFKNLSLIKKYIDGQKKKCGEERRRVNQFLDYLQEFLGV MNTEWI
$\mathbf{z}\left(D_{\mathbf{z}}=153\right):$
PHRRDLCSRSIWLARKIRSDLTALTESYVKHQGLWSELTEAERLQENLQAY RTFHVLLARLLEDQQVHFTPTEGDFHQAIHTLLLQVAAFAYQIEELMILLE YKIPRNEADGMLFEKKLWGLKVLQELSQWTVRSIHDLRFISSHQTGIP

Similarity measure:
number of common substrings

$$
k(\mathbf{x}, \mathbf{z})=\sum_{s \in \mathcal{A}^{*}} w_{s} \phi_{s}(\mathbf{x}) \phi_{s}(\mathbf{z})
$$

where $s$ is a substring, $w_{s} \geq 0$ and $\mathcal{A}^{*}$ the set of all substrings from $\mathcal{A}$.

## String kernels

- If $w_{s}=1$ for a nonempty substring $|s|>0$ :
$\phi(\mathrm{x})=$ number of times each char in $\mathcal{A}$ occurs in x
- bag-of-characters model
- If we require each substring $s$ to be surrounded by white space $\phi(\mathbf{x})=$ number of times each word $s$ occurs in $\mathbf{x}$
- bag-of-words model

Similarity measure:
number of common substrings

$$
k(\mathbf{x}, \mathbf{z})=\sum_{s \in \mathcal{A}^{*}} w_{s} \phi_{s}(\mathbf{x}) \phi_{s}(\mathbf{z})
$$

where $s$ is a substring, $w_{s} \geq 0$ and $\mathcal{A}^{*}$ the set of all substrings from $\mathcal{A}$.

- If we only consider strings with $|s|=k$ we get the $k$-spectrum kernel


## CHALMERS

## An SVM as a neural network



## Gaussian processes

- Linear regression: determine relation $f$ between response $y$ and independent variable $x$

$$
y=f(x)+\epsilon
$$

where $f$ is assumed to be linear: $f(x)=\beta_{0}+\beta_{1} x$


- Bayesian linear regression: determine a posterior distribution over the parameters $\beta_{0}$ and $\beta_{1}$ that gets updated whenever new data is available.
- Gaussian processes: finds a posterior distribution over the possible functions $f(x)$ consistent with the observed data and a suitable prior.


## Gaussian processes

- The current example isn't really linear.
- Quadratic function

$$
y=\beta_{0}+\beta_{1} x+\beta_{2} x^{2}+\epsilon
$$

- Three parameters to estimate: $\beta_{0}, \beta_{1}, \beta_{2}$
- But what if we don't know how many parameters we should use?

- Instead of searching for suitable parameter values for a fixed number of parameters (and a fixed function), we want to search among all functions that fit our data.


## Gaussian processes

- We need to define a prior over the function space.
- Assume we limit our $x$-values: $-5 \leq x \leq 5$.
- In that domain we want to sample functions that are reasonablye smooth.
- We use a covariance matrix to ensure that points close together in input space will produce output values thar are also close together.




## Gaussian processes

- A Gaussian process defines a prior over functions, which, given observed data, can be converted into a posterior.
- Instead of explicitly representing a distribution over a function, choose a finite set of points $\left\{x_{1}, \ldots, x_{n}\right\}$.

- A GP assumes that the function values $f\left(x_{1}\right), \ldots, f\left(x_{n}\right)$ has a jointly Gaussian distribution with some mean $\mu(x)$ and covariance $\Sigma(x)$, given by

$$
\boldsymbol{\Sigma}_{i j}=k\left(x_{i}, x_{j}\right)
$$

where $k$ is a positive definite kernel.


## Multivariate normal distribution

- Gaussian processes are based on the multivariate normal (Gaussian) distribution
$\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})=\frac{1}{\sqrt{|2 \pi \boldsymbol{\Sigma}|}} \exp \left[-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right]$
where $\Sigma$ is a covariance matrix.
- The covariance matrix determines the shape of the "bell"



## Multivariate normal distribution

- Viewing from above, if the contours form a perfect circle, the variables are independent, and the covariance is zero

$$
\boldsymbol{\Sigma}=\left[\begin{array}{ll}
\boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\
\boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22}
\end{array}\right]=\left[\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right]
$$

- With covariance $\Sigma_{12} \neq 0$, the contour will have a more oval shape.



## Gaussian processes

- Assume that we want to learn a function $f$ from data $\mathcal{D}=\left\{\left(x_{i}, y_{i}\right): i=1, \ldots, D\right\}$
- Assume we have a distribution $p(f)$ over functions,
- Now: $p(f)$ is a Gaussian process if for a finite subset $\left\{x_{1}, \ldots, x_{n}\right\}$ the marginal distribution over that subset

$$
p\left(f\left(x_{1}\right), \ldots, f\left(x_{n}\right)\right)
$$

has a Gaussian distribution.

- Now, if the prior is Gaussian, so is the posterior

$$
p(f \mid \mathcal{D})=\frac{p(\mathcal{D} \mid f) p(f)}{p(\mathcal{D})}
$$

## Gaussian processes

- The Gaussian process is parameterized by a mean vector and a covariance matrix

$$
p\binom{f\left(x_{1}\right)}{f\left(x_{2}\right)} \sim N(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma})
$$

where

$$
\boldsymbol{\mu}(\mathbf{x})=\left[\begin{array}{l}
\mu\left(x_{1}\right) \\
\mu\left(x_{2}\right)
\end{array}\right], \quad \Sigma=\left[\begin{array}{ll}
k\left(x_{1}, x_{1}\right) & k\left(x_{1}, x_{2}\right) \\
k\left(x_{2}, x_{1}\right) & k\left(x_{2}, x_{2}\right)
\end{array}\right]
$$

and $k\left(x_{1}, x_{2}\right)$ is a kernel function.

## Gaussian processes

- So: we have the following data

$$
\mathcal{D}=\left\{\left(x_{i}, y_{i}\right): i=1, \ldots, D\right\}
$$

and for some new point $x^{*}$ we want to predict $y^{*}$.

- To do this we want to find a function $f$ such that

$$
y_{i}=f\left(x_{i}\right)
$$

- Instead we assume that $\left\{f\left(x_{1}\right), \ldots, f\left(x_{D}\right)\right\}$ follow a joint normal distribution, and use it to compute the posterior distribution

$$
p\left(f\left(x^{*}\right) \mid x^{*}, \mathcal{D}, f\right)=N\left(\mu^{*}, \Sigma^{*}\right)
$$

## Gaussian processes

- Intuitively, be begin by sampling from our prior distribution


## Gaussian processes

- Intuitively, be begin by sampling from our prior distribution
- We use our training data to represent the outputs of the unknown function.


## Gaussian processes

- Intuitively, be begin by sampling from our prior distribution
- We use our training data to represent the outputs of the unknown function.
- And update the posterior.



## Gaussian processes

- Where does the kernel $k$ come from?
- The covariance matrix characterizes the similarities between nearby points.
- The same range of kernels available as for SVMs.
- The squared exponential (SE)
- The radial basis function (RBF)
- The Matern


## Gaussian processes



Gaussian Process


Neural Net


Random Forest


## Advantages of GPs over SVMs

- GPs handle uncertainty in unknown function $f$ by averaging, not minimizing
- GPs can learn kernel parameters from data, no matter how flexible we want to make the kernel
- GPs can learn regularization parameters without crossvalidation.
- Can incorporate interpretable noise models and priors over functions, and can sample from prior to get intuitions about the model assumptions.
- We can combine automatic feature selection with learning
 using Automatic Relevance Determination (ARD)

