Gaussian Process Implicit Surfaces

Oliver Williams[1]

Microsoft Research and Trinity Hall
Cambridge, UK

Gaussian Processes in Practice
Bletchley Park, 13 June 2006

[1] Joint work with Andrew Fitzgibbon

Talk outline

Implicit surface

Scalar function $f(x)$ defins a surface wherever it passes through a given value (e.g., 0)

$$\mathcal{S}_0 \triangleq \{x \in \mathbb{R}^d | f(x) = 0\}.$$

Implicit surface

Scalar function $f(x)$ defins a surface wherever it passes through a given value (e.g., 0)

$$\mathcal{S}_0 \triangleq \{x \in \mathbb{R}^d | f(x) = 0\}.$$

Example: Function $f(x)$ for $x \in \mathbb{R}^2$ defines a closed curve

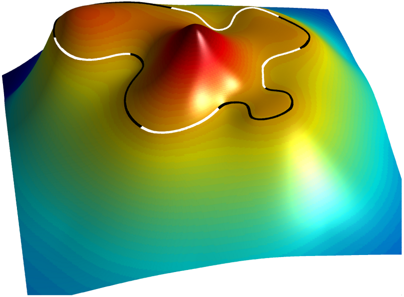

Fitting to data points

Our setting (Turk and O'Brien 1999):

- Given a set of **constraint** points in 2D or 3D $\{x_i\}$, fit $f(x)$
- Have constraints at $f(x_i) = 0$ on the curve and at ± 1 off it e.g.,
  - Simple interior/exterior case
  - Control normals to curve

Fitting to data points

Our setting (Turk and O'Brien 1999):

- Given a set of **constraint** points in 2D or 3D $\{x_i\}$, fit $f(x)$
- Have constraints at $f(x_i) = 0$ on the curve and at ± 1 off it e.g.,
  - Simple interior/exterior case
  - Control normals to curve

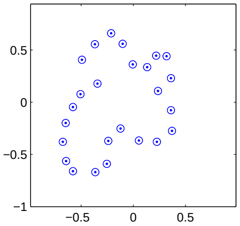

Fitting to data points

Our setting (Turk and O'Brien 1999):

- Given a set of **constraint** points in 2D or 3D $\{x_i\}$, fit $f(x)$
- Have constraints at $f(x_i) = 0$ on the curve and at ± 1 off it e.g.,
  - Simple interior/exterior case
  - Control normals to curve

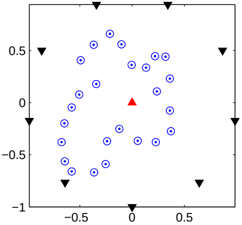

Fitting to data points

Our setting (Turk and O'Brien 1999):

- Given a set of **constraint** points in 2D or 3D $\{x_i\}$, fit $f(x)$
- Have constraints at $f(x_i) = 0$ on the curve and at ± 1 off it e.g.,
  - Simple interior/exterior case
  - Control normals to curve

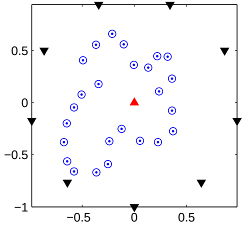

Alternative method: **parametric surface**: $x(t), y(t), [z(t)]$

- What t to assign to data points?
- How to handle different topologies?
- Can represent non-closed curves/surfaces

Topology change

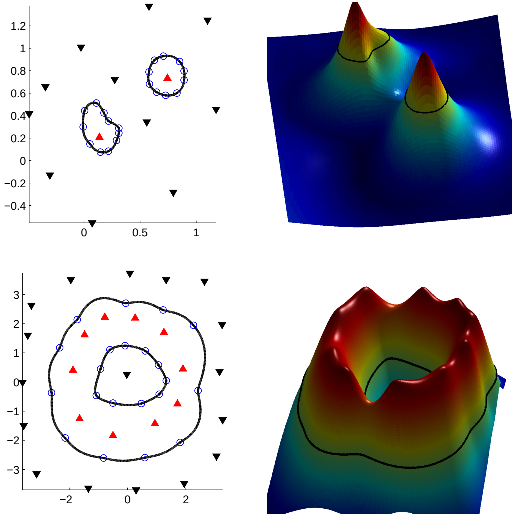

Regularization

Find function passing through constraint points which minimizes **thin-plate spline energy**

$$E(f) = \int_\Omega \left(\nabla^\mathsf{T} \nabla f(x) \right)^2 \, dx$$

Regularization

Find function passing through constraint points which minimizes **thin-plate spline energy**

$$E(f) = \int_\Omega \left(\nabla^\mathsf{T} \nabla f(x)\right)^2 \, dx$$

Fit $f(x)$ with Gaussian process

Regularization

Find function passing through constraint points which minimizes **thin-plate spline energy**

$$E(f) = \int_{\Omega} \left(\nabla^{\mathsf{T}} \nabla f(x) \right)^2 \, dx$$

Fit $f(x)$ with Gaussian process
Use covariance function equivalent to thin-plate spline

Gaussian distribution

- Follow derivation in (MacKay 2003) in 1D

Gaussian distribution

- Follow derivation in (MacKay 2003) in 1D
- consider the energy as a probability and define D as the linear differential operator

$$E(f) = -\log P(f) + \text{const} = \int_\Omega \left(D^2 f(x)\right)^2 \, dx.$$

Gaussian distribution

- Follow derivation in (MacKay 2003) in 1D
- consider the energy as a probability and define D as the linear differential operator

$$E(f) = -\log P(f) + \text{const} = \int_\Omega \left(D^2 f(x)\right)^2 \, dx.$$

- Use $f(\Omega)$ as vector of function values for all points in Ω:

$$-\log P(f(\Omega)) = f(\Omega)^\mathsf{T} [D^2]^\mathsf{T} D^2 f(\Omega),$$

Gaussian distribution

- Follow derivation in (MacKay 2003) in 1D
- consider the energy as a probability and define D as the linear differential operator

$$E(f) = -\log P(f) + \text{const} = \int_\Omega \left(D^2 f(x)\right)^2 \, dx.$$

- Use $f(\Omega)$ as vector of function values for all points in Ω:

$$-\log P(f(\Omega)) = f(\Omega)^\mathsf{T} [D^2]^\mathsf{T} D^2 f(\Omega),$$

- This is a Gaussian disribution with mean zero and covariance:

$$C = \left([D^2]^\mathsf{T} D^2\right)^{-1} = \left(D^4\right)^{-1}.$$

Covariance function

- Entries of C indexed by $u, v \in \Omega$

$$\int_\Omega D^4(u,w)c(w,v)\,dw = \delta(u-v) \quad \Rightarrow \quad \frac{\partial^4}{\partial r^4}c(r) = \delta(r)$$

where we impose stationarity with $r = u - v$.

Covariance function

- Entries of C indexed by $u, v \in \Omega$

$$\int_\Omega D^4(u,w)c(w,v)\, dw = \delta(u-v) \quad \Rightarrow \quad \frac{\partial^4}{\partial r^4}c(r) = \delta(r)$$

  where we impose stationarity with $r = u - v$.

- Interpret as spectral density and solve

$$\mathcal{F}\big[c(r)\big](\omega) = \omega^{-4}$$

$$\Rightarrow \quad c(r) = \tfrac{1}{6}|r|^3 + a_3 r^3 + a_2 r^2 + a_1 r + a_0.$$

Covariance function

$$c(r) = \frac{1}{6}|r|^3 + a_3 r^3 + a_2 r^2 + a_1 r + a_0$$

Find constants using constraints on $c(\cdot)$

- Symmetry: $a_3 = a_1 = 0$

Covariance function

$$c(r) = \frac{1}{6}|r|^3 + a_3 r^3 + a_2 r^2 + a_1 r + a_0$$

Find constants using constraints on $c(\cdot)$

- Symmetry: $a_3 = a_1 = 0$
- Postive definiteness: simulate by making $c(r) \to 0$ at $\partial\Omega$

$$c(r) = \frac{1}{12}\left(2|r|^3 - 3Rr^2 + R^3\right).$$

where R is the largest magnitude of r within Ω.

1D regression demonstration

GP predicts function values for set of points $\mathcal{U} \subseteq \Omega$

$$P(f(\mathcal{U})|\mathcal{X}) = \text{Normal}\left(f(\mathcal{U}) \mid \mu,\ Q\right)$$

where

$$\mu = C_{ux}^{\mathsf{T}}(C_{xx} + \sigma^2 I)^{-1} t \quad \text{and} \quad Q = C_{uu} - C_{ux}^{\mathsf{T}}(C_{xx} + \sigma^2 I)^{-1} C_{ux}.$$

The matrices are formed by evaluating $c(\cdot,\cdot)$ between sets of points: i.e., $C_{xx} = [c(x_i, x_j)]$, $C_{ux} = [c(u_i, x_j)]$, and $C_{uu} = [c(u_i, u_j)]$.

1D regression demonstration

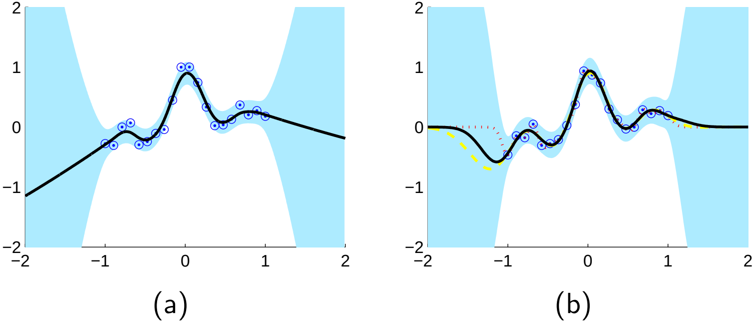

Figure: **Thin plate vs. squared exponential covariance.** *Mean (solid line) and 3 s.d. error bars (filled region) for GP regression (a) Thin-plate covariance; (b) Squared exponential covariance function* $c(u_i, u_j) = e^{-\alpha \|u_i - u_j\|^2}$ *with* $\alpha = 2$, 10 *and* 100*; error bars correspond to* $\alpha = 10$.

Covariance in 2D

- In 2D the Green's equation is

$$\left(\nabla^{\mathsf{T}}\nabla\right)^2 c(r) = \delta(r)$$

where now $c(u, v) = c(r)$ with $r \triangleq \|u - v\|$.

Covariance in 2D

- In 2D the Green's equation is

$$\left(\nabla^{\mathsf{T}}\nabla\right)^2 c(r) = \delta(r)$$

where now $c(u, v) = c(r)$ with $r \triangleq \|u - v\|$.

- Solution (with similar constraints at the boundary of Ω)

$$c(r) = 2r^2 \log|r| - \left(1 + 2\log(R)\right) r^2 + R^2$$

Demonstration

- Set constraint points
-

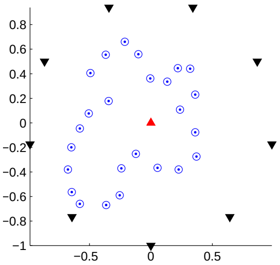

$\mathcal{X}$: $\{\odot = 0,\ \blacktriangle = +1, \blacktriangledown = -1\}$

Demonstration

- Set constraint points
- Fit GP to points

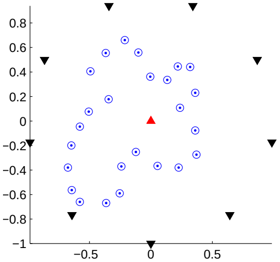

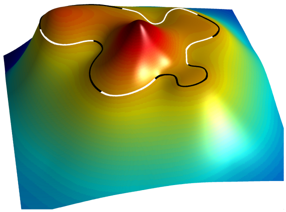

$\mathcal{X}$: $\{\odot = 0, \blacktriangle = +1, \blacktriangledown = -1\}$

Mean function

Probabilistic interpretation

Gaussian process makes probabilistic prediction:

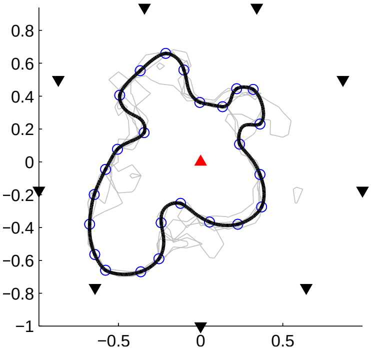

Samples from posterior

Probabilistic interpretation

Gaussian process makes probabilistic prediction:

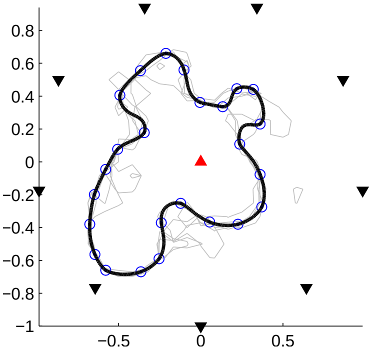

Samples from posterior

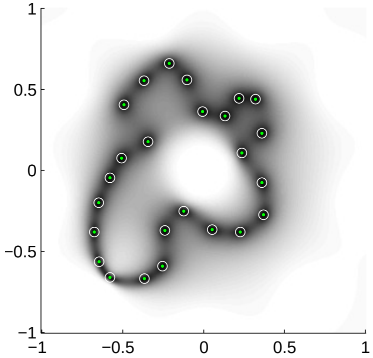

$P(x \in \mathcal{S}_0) \equiv P(f(x) = 0).$

With different topology

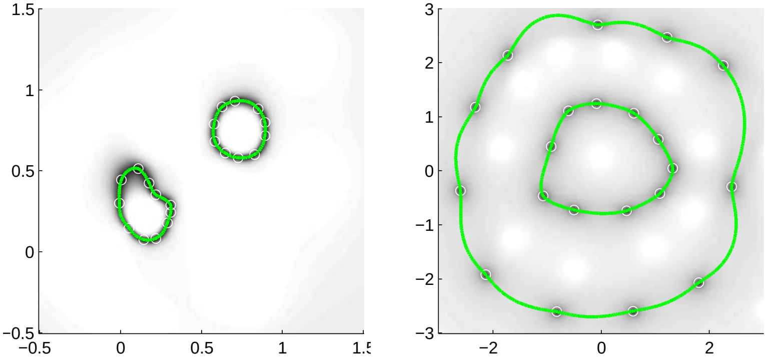

Result with squared exponential

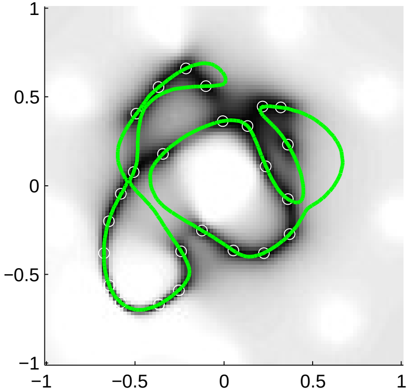

Fitting 3D surfaces

In 3D the covariance is

$$c(r) = 2|r|^3 + 3Rr^2 + R^3$$

Fitting 3D surfaces

In 3D the covariance is

$$c(r) = 2|r|^3 + 3Rr^2 + R^3$$

- Take n points on surface of object

Fitting 3D surfaces

In 3D the covariance is

$$c(r) = 2|r|^3 + 3Rr^2 + R^3$$

- Take n points on surface of object
- Define internal and external points

Fitting 3D surfaces

In 3D the covariance is

$$c(r) = 2|r|^3 + 3Rr^2 + R^3$$

- Take n points on surface of object
- Define internal and external points
- Fit Gaussian process

Fitting 3D surfaces

In 3D the covariance is

$$c(r) = 2|r|^3 + 3Rr^2 + R^3$$

- Take n points on surface of object
- Define internal and external points
- Fit Gaussian process
- Use **marching cubes** algorithm to find mean surface

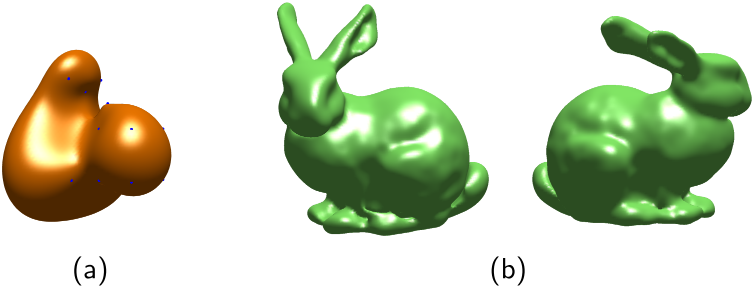

Figure: **3D surfaces.** Mean surfaces $\mu(x) = 0$ when $x \in \mathbb{R}^3$, rendered as an high resolution polygonal mesh generated by the marching cubes algorithm. *(a) A simple "blob" defined by 15 points on the surface, one interior +1 point and 8 exterior -1 points arranged as a cube; (b) Two views of the Stanford bunny defined by 800 surface points, one interior +1 point, and a sphere of 80 exterior -1 points.*

Summary

- Gaussian processes can be used to define curves and surfaces

Summary

- Gaussian processes can be used to define curves and surfaces
- Appropriate covariance must be used to obtain high quality results

Summary

- Gaussian processes can be used to define curves and surfaces
- Appropriate covariance must be used to obtain high quality results
- By using a GPIS, curves and surfaces have a meaningful probabilistic interpretation

Summary

- Gaussian processes can be used to define curves and surfaces
- Appropriate covariance must be used to obtain high quality results
- By using a GPIS, curves and surfaces have a meaningful probabilistic interpretation

Shortcomings / ideas for future work:

- Exploit probabilistic nature of GPIS in computer vision problems
- More elegant methods for constraining surface normals?
- Can this be used to learn a meaningful prior?
- Scale/smoothness control?