Soap film smoothing

Simon Wood, Mark Bravington & Sharon Hedley

with example by Nicole Augustin & Verena Trenkel

A problem for smoothers…

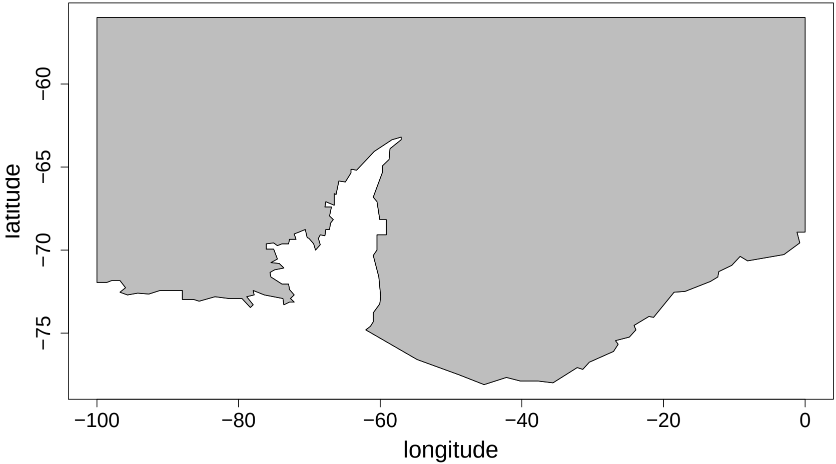

Ramsay's horshoe

Tim Ramsay (2002) proposed a test function like this.

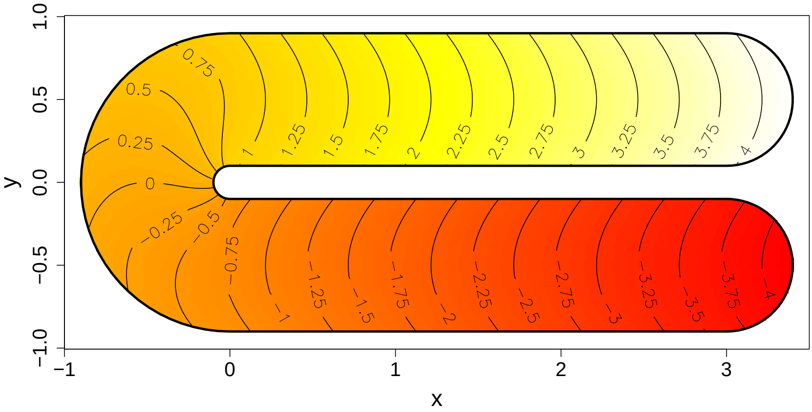

It's *very* difficult to reconstruct by smoothing samples from the function.

Stone's method

- Stone (1988) developed finite window smoothers by modifying a thin plate spline.
- Idea is to integrate thin plate spline penalty

$$\int f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2 dxdy$$

  only over the domain of interest.
- Problem is that the penalty still shrinks the smooth function towards a plane. This is innappropriate for the examples just given.

Ramsay's method

- Tim Ramsay (JRSSB, 2002), proposed a solution he called FELSPLINE.
- It is a *big* improvement on thin plate spline and other conventional smoothers.
- But it requires a computationally complex finite element approach to estimate the smoother, so that e.g. smoothness selection and variance estimation are expensive.
- It also requires the *very strong* boundary condition that contours of the smooth meet the boundary at right angles.

Our aims

1. To produce a computationally straightforward and efficient finite window smoother, representable using a penalized basis.
2. To produce a finite window smoother that can meet any smooth known boundary condition with zero wiggliness.

- Such a smoother can be easily incorporated as a component of other models (GAMs, GAMMs or something more interesting), and estimated via penalized likelihood, as a mixed model etc.
- Smoothness selection is also easy (by GCV, AIC, REML etc).

Known boundary smoothing

- Consider a region Ω of the $x - y$ plane within closed boundary B.
- Suppose that some function $g(x, y)$ is known on B, and we wish to find a smooth function over Ω which meets this boundary condition.
- Imagine the boundary condition as a loop of wire, with x, y co-ordinates given by B, and 'z' co-ordinate $g(x, y)$.
- Nature's solution to smoothly interpolating the boundary is a soap film suspended from the boundary wire.

Soap film smoother idea

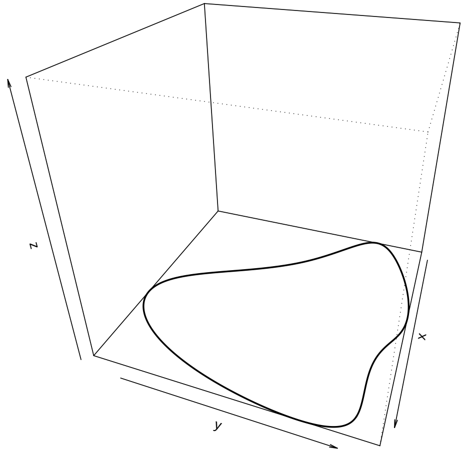

Soap film smoother idea

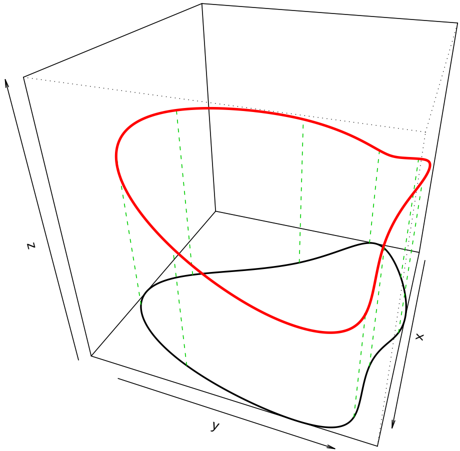

Soap film smoother idea

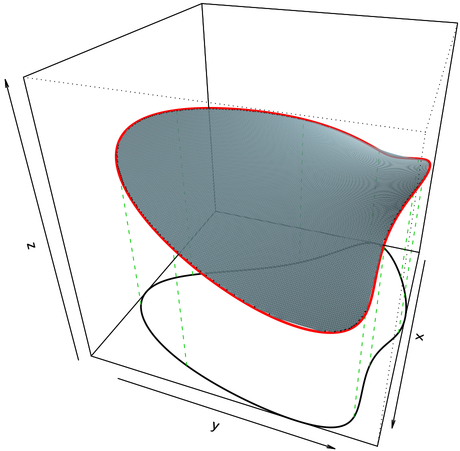

Soap film smoothness

- Making a ‘small displacement’ (and zero gravity) assumption, the boundary interpolating soap film f obeys the Laplace equation

$$f_{xx} + f_{yy} = 0$$

- This suggests that if f were to be allowed to distort (e.g. to represent some function over Ω), then

$$\int_\Omega (f_{xx} + f_{yy})^2 dxdy$$

  might be a reasonable measure of departure from smoothness.

Soap film smoother idea

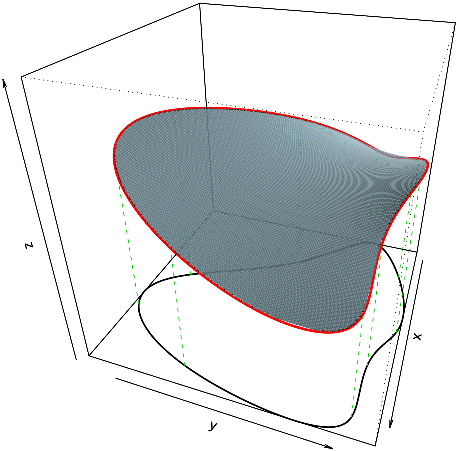

Soap film smoother idea

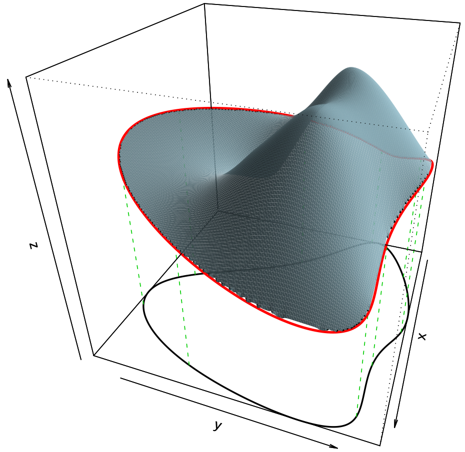

Soap film smoothing

- Consider noisy observations, z_k of some function $g(x_k, y_k)$, where $x_k, y_k \in \Omega$, and g is known on Ω's boundary, B.
- To estimate g we seek to find the function, f, minimizing

$$\sum_i \{z_i - f(x_i, y_i)\}^2 + \lambda \int_\Omega (f_{xx} + f_{yy})^2 dxdy$$

  where λ is a smoothing parameter.
- How can this be computed?

Soap film interpolation theorem

- Let f^* be a function known on the boundary, B of Ω, and z_k, x_k, y_k be data such that $x_k, y_k \in \Omega$.
- The function $f(x, y)$ which interpolates f^* on B, satisfies $z_k = f(x_k, y_k)$ and minimizes

$$\int_\Omega (f_{xx} + f_{yy})^2 dxdy$$

must satisfy

$$f_{xx} + f_{yy} = \rho(x, y), \quad \text{where} \quad \rho_{xx} + \rho_{yy} = 0$$

except at x_k, y_k (also need assumption that $\rho = 0$ on B).

Soap film smoothing lemma

- Setup is the same as for the SFIT.
- The function $f(x, y)$ which interpolates f^* on B and minimizes

$$\sum_i \{z_i - f(x_i, y_i)\}^2 + \lambda \int_\Omega (f_{xx} + f_{yy})^2 dxdy$$

must satisfy

$$f_{xx} + f_{yy} = \rho(x, y), \quad \text{where} \quad \rho_{xx} + \rho_{yy} = 0$$

except at x_k, y_k.
- This is the key to computational practicality.

Computational preliminaries

- Numerical solution of the defining PDEs is a very well studied problem! (They are the Laplace and Poisson equations.)
- For actual computation, define $\rho_i(x, y)$ as the solution of the PDE

$$\rho_{xx} + \rho_{yy} = 0,$$

  except at a singularity point x_i, y_i, subject to $\rho = 0$ on B and $\int_\Omega \rho(x, y) = 1$.
- Successive Over Relaxation (SOR) is a simple solution method for obtaining the ρ_i, but multigrid is faster. Currently we are also investigating sparse matrix direct methods.

Computation of the penalty

- Define parameters γ_k such that $\rho(x,y) = \sum_k \gamma_k \rho_k(x,y)$.
- It is easy to show that

$$\int_\Omega (f_{xx} + f_{yy})^2 dxdy = \gamma^\mathsf{T} \mathbf{S} \gamma$$

  where

$$\mathbf{S}_{ij} = \int_\Omega \rho_i(x,y) \rho_j(x,y) dxdy$$

- Given gridded numerical evaluations of the $\rho_i(x,y)$, numerical evaluation of the integrals is easy.

A computable basis for f

- Let $a(x, y)$ be the solution of $f_{xx} + f_{yy} = 0$ subject to known boundaries conditions on B.
- Let $g_i(x, y)$ be the solution of

$$f_{xx} + f_{yy} = \rho_i(x, y)$$

  with $f = 0$ on B
- The soap film smoother can be written as

$$f(x, y) = a(x, y) + \sum_k \gamma_k g_k(x, y).$$

The basis components

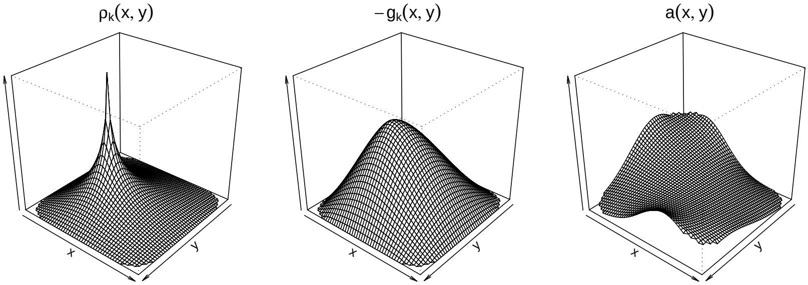

Estimation

- Given a basis and quadratic penalty, it's easy to incorporate soap film smooths as components of e.g. GAM(M)s, (G)LMMs, non-linear mixed models, etc.
- e.g. the original simple smoothing problem of finding f to minimize

  $$\sum_i \{z_i - f(x_i, y_i)\}^2 + \lambda \int_\Omega (f_{xx} + f_{yy})^2 dxdy.$$

  becomes the standard generalized ridge regression problem of finding γ to minimize

  $$\|\mathbf{z} - \mathbf{a} - \mathbf{X}\gamma\|^2 + \lambda \gamma^\mathsf{T} \mathbf{S} \gamma.$$

Alternative estimation

- Alternatively, the simple smoothing problem can be viewed from a mixed modelling or Bayesian perspective ...
- The model becomes

  $$\mathbb{E}(\mathbf{z}|\gamma) = \mathbf{a} + \mathbf{X}\gamma$$

  where

  $$\gamma \sim N(\mathbf{0}, \alpha\mathbf{S}^{-1}).$$

  α is a variance component to be estimated ($\propto \lambda^{-1}$).
- As a linear mixed model this can be estimated by likelihood based methods, or stochastic simulation can be used.

Penalized regression bases

- There is one γ_k per datum z_k.
- Usually this is computationally wasteful.
- It is better to choose a ‘nicely distributed’, but relatively small set of x_i, y_i points in Ω, and to set up a basis and penalty as if these were the data locations.
- This ‘penalized regression basis’ is then used to model the actual data.
- Usually leads to large computational savings at little cost in ‘statistical performance’.

Unknown boundaries

- Unknown boundaries can be dealt with by defining a cyclic penalized regression spline on the loop B.
- Suppose that the boundary spline has parameters $\boldsymbol{\beta}$.
- Let $a_i(x, y)$ be the solution of

$$f_{xx} + f_{yy} = 0$$

  subject to the boundary condition given by setting $\beta_k = 0 \ \forall \ k \neq i$ and $\beta_i = 1$.
- Then $a(x, y) = \sum_i \beta_i a_i(x, y)$, and we acquire an extra quadratic penalty forcing boundary smoothness.

Better than TPS

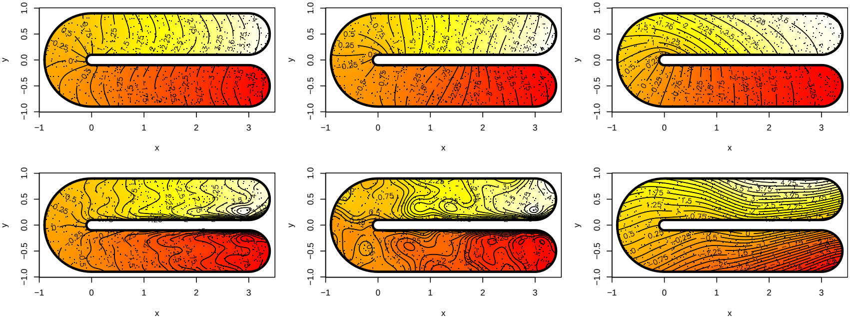

- Reconstructions of test function from scattered noisy data (noisier left to right).
- Free B soap upper. TPS lower. λ selection by GCV.

Better than FELSPLINE

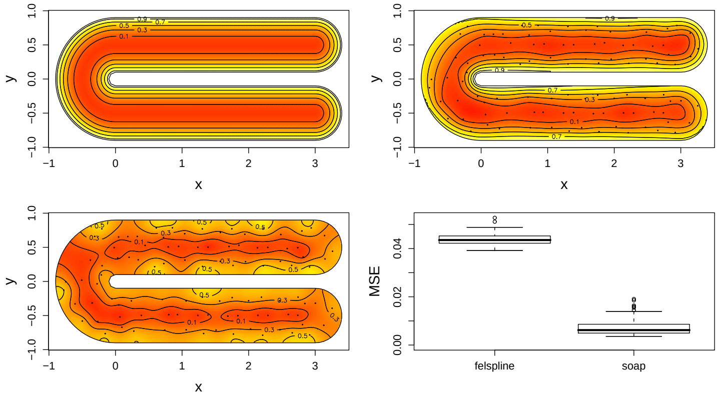

- Truth top left; FB soap top right; FELSPLINE lower left.

Better than modified FELSPLINE

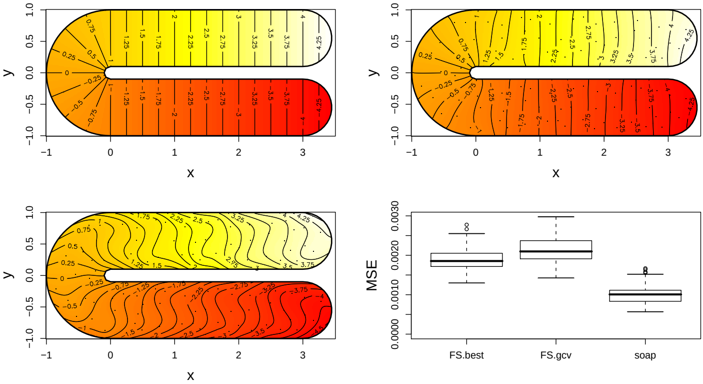

- Example from Ramsay (2002). Truth top left; FB soap top right; modified FELSPLINE lower left.

Benign comparison?

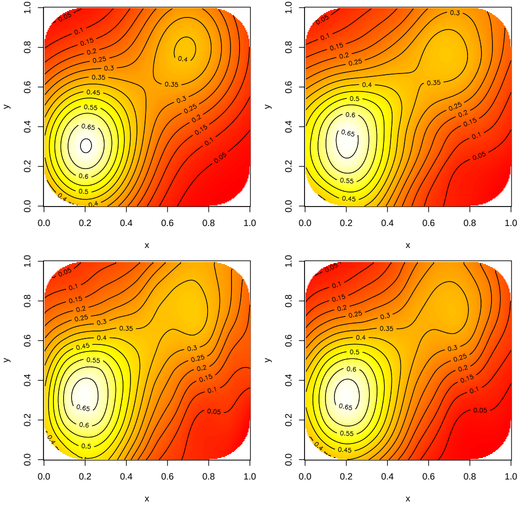

Benign MSE comparison

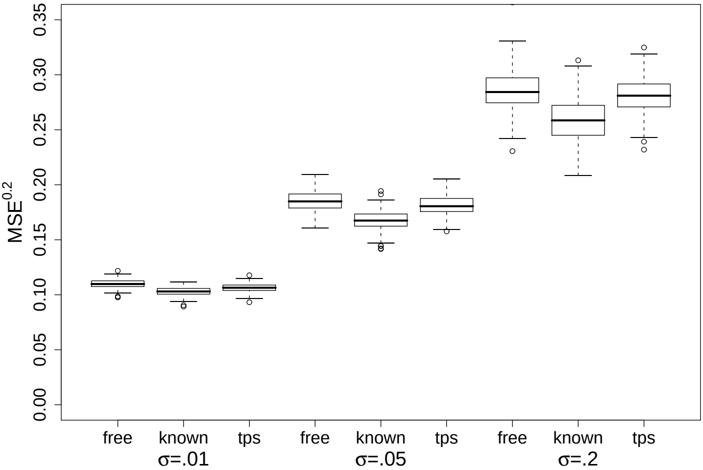

Simulation conclusions

- Soap films offer a substantial advantage over the alternatives for smoothing over complicated domains.
- The price paid is in performance on uncomplicated domains, but seems to be fairly modest.
- Is all this of any use for anything real?

How many Soles are there?

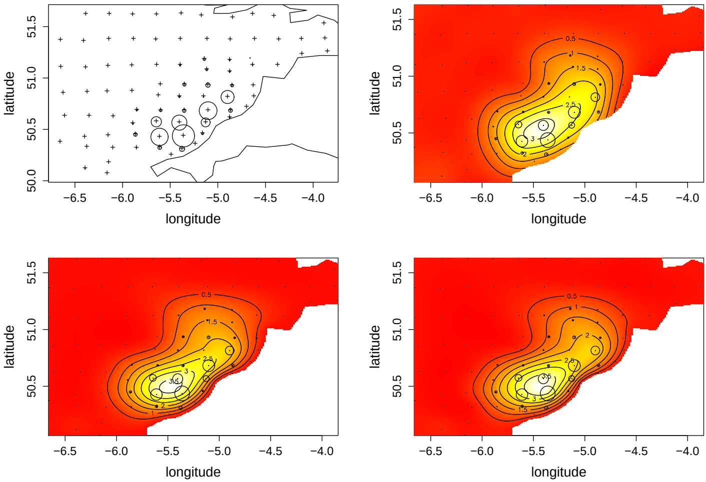

Aral Sea Chlorophyll

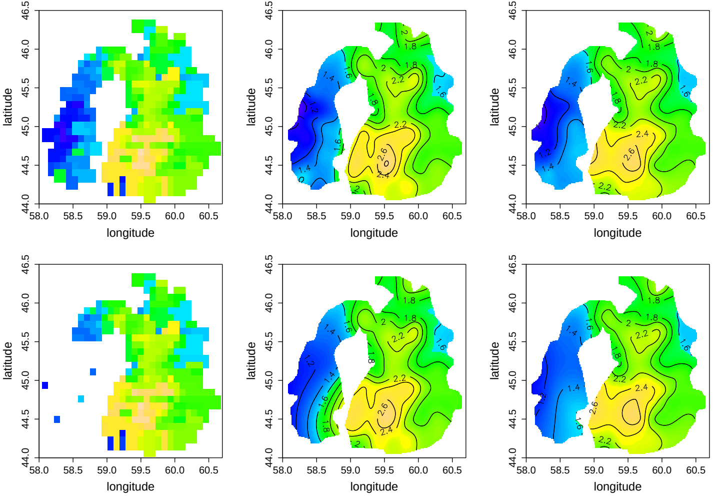

Space time modelling of Blue Ling

Blue Ling Area of Interest

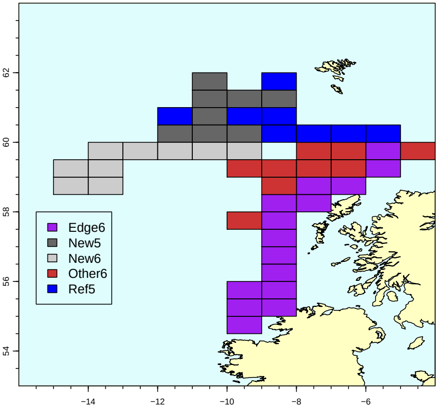

Blue Ling model (Augustin & Trenkel)

- Ling per trawl data, y, are available from commercial boats off north west Scotland.
- Interest is in understanding spatial distribution, and temporal trends in abundance.
- Basic model is

$$\begin{aligned}\log(\mu_{it}) = f_1(duration_{it}) + f_2(depth_i) + f_3(month_{it}) + \\ f_4(lat_i, lon_i, year_t) + f_5(depth_i, month_{it}) \\ + f_6(depth_i, year_{it})\end{aligned}$$

  where $E(y_{it}) = \mu_{it}$ and $y_{it} \sim \text{Tweedie}(\mu_{it}, \phi\mu_{it}^{1.5})$
- Space time term should be soap film like in space, but vary in time …

Tensor product smooths

- A time varying spatial soap film can be constructed as a (pair of) tensor product smooth(s).
- Tensor product smooths are best explained using a 2D example.
- Consider constructing a smooth of x, z.
- Start by choosing *marginal* bases and penalties, as if constructing 1-D smooths of x and z. e.g.

$$f_x(x) = \sum \alpha_i a_i(x), \quad f_z(z) = \sum \beta_j b_j(z),$$

$$J_x(f_x) = \int f_x''(x)^2 dx = \boldsymbol{\alpha}^\mathsf{T} \mathbf{S}_x \boldsymbol{\alpha} \;\; \& \;\; J_z(f_z) = \mathcal{B}^\mathsf{T} \mathbf{S}_z \mathcal{B}$$

Marginal reparameterization

- Suppose we start with $f_z(z) = \sum_{i=1}^{6} \beta_j b_j(z)$, on the left.

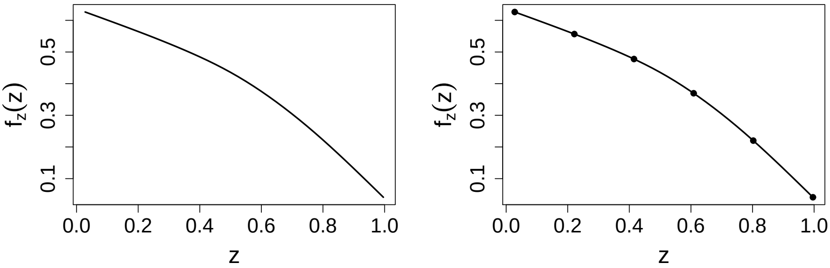

- We can always re-parameterize so that its coefficients are functions heights, at knots (right). Do same for f_x.

Making f_z depend on x

- Can make f_z a function of x by letting its coefficients vary smoothly with x

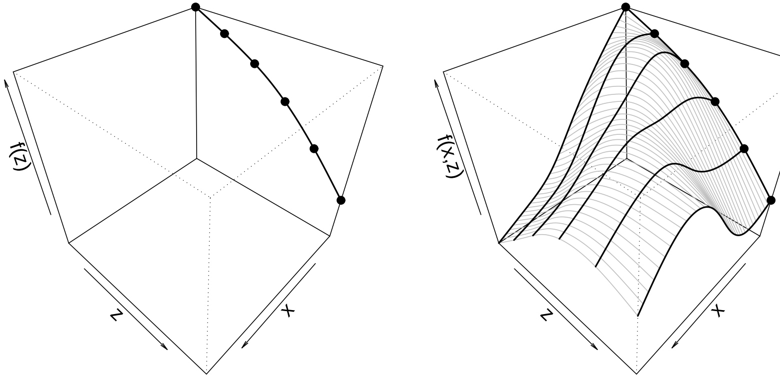

The complete tensor product smooth

- Use f_x basis to let f_z coefficients vary smoothly (left).
- Construct in symmetric (see right).

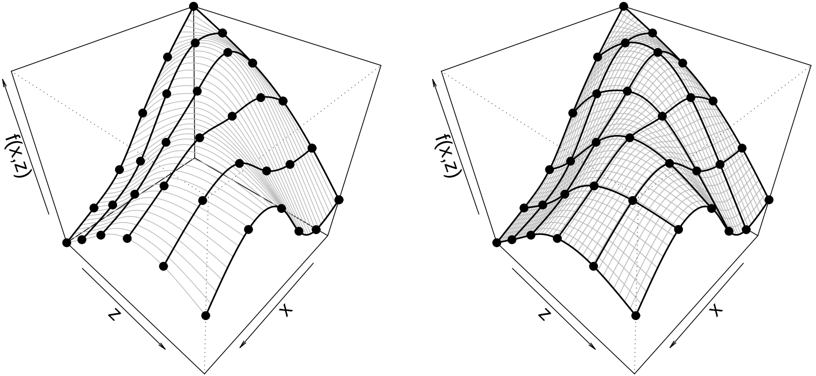

Tensor product penalties - one per margin

- *x*-wiggliness: sum marginal *x* penalties over red curves.
- *z*-wiggliness: sum marginal *z* penalties over green curves.

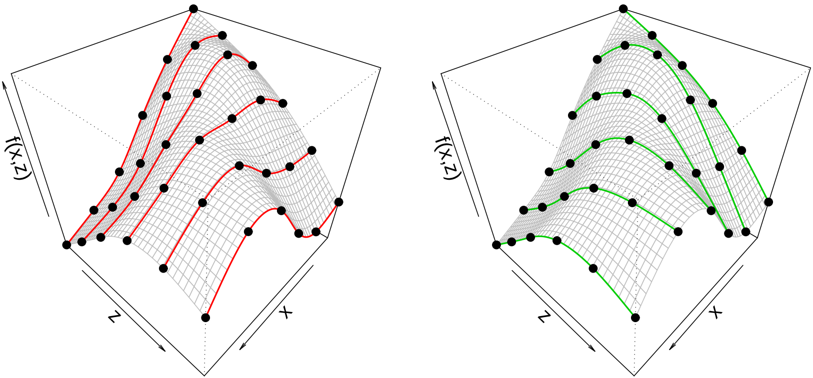

Tensor product expressions

- So the tensor product basis construction gives:

$$f(x,z) = \sum \sum \beta_{ij} b_j(z) a_i(x)$$

- With double penalties

$$J_z^*(f) = \boldsymbol{\beta}^\mathsf{T} \mathbf{I}_I \otimes \mathbf{S}_z \boldsymbol{\beta} \text{ and } J_x^*(f) = \boldsymbol{\beta}^\mathsf{T} \mathbf{S}_x \otimes \mathbf{I}_J \boldsymbol{\beta}$$

- The construction generalizes to any number of marginals and multi-dimensional marginals.
- In particular a tensor product of a soap film and a 1D smooth of time is possible.
- The soap film smoother is separated into the boundary-interpolating-film and the deviation-from-film parts, and tensor products with time are formed for each.

Blue Ling space time distribution

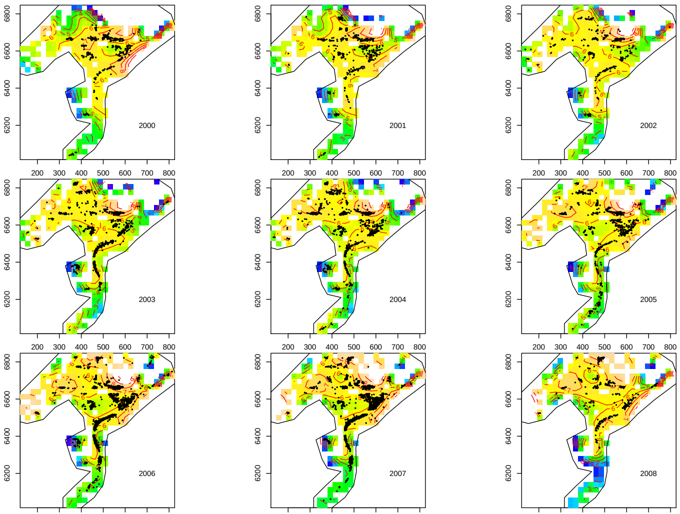

Area trends over time

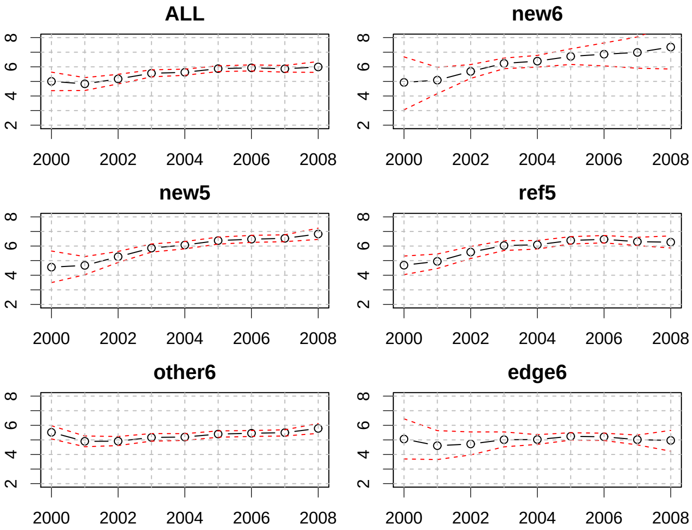

Concluding waffle

- Soap film smoothing works nicely, and the smooths are easy to incorporate as components of other models.
- R package `soap` is at `www.maths.bath.ac.uk/~sw283/`.
- Higher dimensional versions of soap are possible, but less easy to compute with.
- Theoretically the boundary smooth seems inelegant, but in practice its tendency to suppress edge effects, even for simple boundaries, seems to be one of the soap film smoothers main advantages.