

An Overview of the Functional Linear Model

Jim Ramsay, McGill University

McGill Psychology Department
Montreal, February 23, 2009

Overview

- Four functional linear models for the daily weather data.
- A functional ANOVA for precipitation.
- Predicting total annual precipitation from the temperature profile.
- Predicting today's precipitation from today's temperature.
- Predicting the entire year's precipitation from the year's temperature profile.
- A short-term feed-forward model for precipitation.
- A more general perspective.
- Predicting precipitation dynamics: a differential equation
- The idea of a linear model reviewed.

Outline

- 1 **A Quick Introduction to Functional Linear Models**
- 2 Modelling functional responses with multivariate covariates
- 3 Modelling functional responses with multivariate covariates
- 4 Modelling scalar responses with functional covariates
- 5 Functional responses, functional covariates and the concurrent model
- 6 Functional responses and functional covariates: The general case

The average Canadian weather data

- 35 Canadian weather stations selected to cover the country.
- Daily temperatures (0.1 degrees Celsius) and precipitations (0.1 mm) averaged over the years 1960 to 1994. (Feb 29th combined with Feb. 28th).
- Canada divided into Atlantic, Continental, Pacific and Arctic weather zones.

A functional analysis of variance

- Does the precipitation profile vary from one weather zone to another?
- We have a number N_g of weather stations in each climate zone $g = 1, \dots, 4$, and
- the model for the m th temperature function in the g th group, indicated by Prec_{mg} , is

$$\text{Prec}_{mg}(t) = \mu(t) + \alpha_g(t) + \epsilon_{mg}(t).$$

- $\mu(t)$ is the grand mean function, summarizing precipitation for all of Canada.
- $\alpha_g(t)$ is the functional effect of being in weather zone g
- In order to fix zone effects, we require that

$$\sum_g \alpha_g(t) = 0 \text{ for all } g$$

A scalar response and a functional independent variable

- The response is the log total annual precipitation

$$\text{PrecTot}_i = \int_0^{365} \text{Prec}_i(t) dt$$

- The model is

$$\log(\text{PrecTot}_i) = \alpha + \int_0^{365} \text{Temp}_i(s) \beta(s) ds + \epsilon_i.$$

- But here we have a real problem. How to avoid over-fitting the 35 scalar observations?
- We'll use regularization or roughness penalties on the estimated regression functions.

A functional response and a functional independent variable

This is a big topic, and breaks down into several useful special versions.

The concurrent functional model

- We might only use the temperature at the same time $s = t$ because we imagine that precipitation now depends only on the temperature now.
- Our model is

$$\text{Prec}_i(t) = \alpha(t) + \text{Temp}_i(t)\beta(t) + \epsilon_i(t)$$

- We might call this model *concurrent* or *point-wise*.
- Should we use regularization to force β to be smooth in t ?

The annual or total model

- We may prefer to allow for temperature influence on $\text{Prec}(t)$ to extend over the whole year.
- The model expands to become

$$\text{Prec}_i(t) = \alpha(t) + \int_0^{365} \text{Temp}_i(s) \beta(s, t) ds + \epsilon_i(t)$$

- The value $\beta(s, t)$ determines the impact of temperature at time s on precipitation at time t .
- We need roughness penalties for variation in both s and t

The limited-term feed-forward model

- it may be that what counts is whether the temperature has been falling rapidly up to time t . The model expands to

$$\text{Prec}_i(t) = \alpha(t) + \int_{t-\delta}^t \text{Temp}_i(s) \beta(s, t) ds + \epsilon_i(t)$$

- Here δ is the time lag over which we use temperature information.
- Now β is only defined over the somewhat complicated trapezoidal domain: $t \in [0, 365], t - \delta \leq s \leq t$.

The local influence model

- Finally, we may open up the model to allow integration over s within a t -dependent set Ω_t .
- The model may therefore be

$$\text{Prec}_i(t) = \alpha(t) + \int_{\Omega_t} \text{Temp}_i(\mathbf{s}) \beta(\mathbf{s}, t) d\mathbf{s} + \epsilon_i(t)$$

Predicting derivatives

- When the response is a derivative, then there is the potential for the function itself to be a useful covariate.
- The concurrent linear model

$$D\text{Prec}_i(t) = \text{Prec}_i(t)\beta(t) + \epsilon_i(t)$$

is a *homogeneous first order linear differential equation* in precipitation.

- If we also include an influence of temperature,

$$D\text{Prec}_i(t) = \text{Prec}_i(t)\beta_0(t) + \text{Temp}_i(t)\beta_1(t) + \epsilon_i(t),$$

the equation is said to be *forced* or *nonhomogeneous*.

What exactly makes a model linear?

- We see that the functional linear model has a lot more variants than it's poor multivariate cousin.
- We will need to look at a definition of a linear model that encompasses these models and many others.


Outline

- 1 A Quick Introduction to Functional Linear Models
- 2 Modelling functional responses with multivariate covariates**
- 3 Modelling functional responses with multivariate covariates
- 4 Modelling scalar responses with functional covariates
- 5 Functional responses, functional covariates and the concurrent model
- 6 Functional responses and functional covariates: The general case

Predicting temperature curves from climate zones

- We have 35 weather stations distributed across four climate zones:
 - Atlantic (16)
 - Pacific (6)
 - Continental (13)
 - Arctic (4)
- The dependent variable is $\text{Temp}(t)$, a function representing daily temperatures averaged over 1960–1994.
- The temperature functions were obtained by expanding the original 365 discrete daily averages in terms of 65 Fourier basis functions.

Montreal's temperature profile



The functional ANOVA model

- The model is

$$\text{Temp}_{mg}(t) = \mu(t) + \alpha_g(t) + \epsilon_{mg}(t).$$

- μ is the grand mean function
- α_g are the specific effects on temperature of being in climate zone g . To be able to identify them uniquely, we require that they satisfy the constraint

$$\sum_g \alpha_g(t) = 0 \text{ for all } t. \quad (1)$$

- ϵ_{mg} is the residual function showing unexplained variation specific to the k th weather station within climate group g .

Setting up the model

- Set up a 35 by 5 matrix \mathbf{Z} . Column 1 contains all 1's, and columns $g + 1, g = 1, \dots, 4$ contain zeros except for 1's in rows corresponding to stations in climate zone g .
- Append a final row with 0 in column 1, and 1's in the remaining columns.
- Let the functional response vector $\mathbf{Temp}(t)$ contain the 35 temperature profiles *plus* a final function that is zero for all t .
- Let functional regression coefficient vector $\beta(t)$ contain the functions $(\mu, \alpha_1, \dots, \alpha_4)$.
- The model in matrix notation,, including the zero sum constraint,is

$$\mathbf{Temp}(t) = \mathbf{Z}\beta(t) + \epsilon(t),$$

Outline

- 1 A Quick Introduction to Functional Linear Models
- 2 Modelling functional responses with multivariate covariates
- 3 Modelling functional responses with multivariate covariates**
- 4 Modelling scalar responses with functional covariates
- 5 Functional responses, functional covariates and the concurrent model
- 6 Functional responses and functional covariates: The general case

Fitting the model





- The residual $\text{Temp}_i(t) - \mathbf{Z}_i\beta(t)$ is now a function.
- The least squares fitting criterion becomes

$$\text{LMSSE}(\beta) = \sum_g^4 \sum_m^{N_g} \int [\text{Temp}_{mg}(t) - \sum_j^q z_{(mg),j} \beta_j(t)]^2 dt.$$





- This is minimized with respect to the regression functions by

$$\hat{\beta}(t) = (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'\mathbf{Temp}(t)$$


The region effects $\alpha_g(t)$



The mean plus region effects $\mu(t) + \alpha_g(t)$



The standard error of measurement function $\sigma(t)$



Assessing fit

- Is there significant variation in temperature over climate zones? Of course there is! This does not seem like an interesting question.
- On the other hand, whether the Atlantic, Pacific and Continental stations are significantly different in the summer might be.
- Interesting summaries of fit, of effects, and inferences are likely to be *local* in nature.

- It is useful to use the error sum of squares function

$$SSE(t) = \sum_{mg} [\text{Temp}_{mg}(t) - \mathbf{Z}_{mg}\hat{\beta}(t)]^2.$$

to assess fit at or near time t .

- As in ordinary regression, we can compare this to the variation of the response about its mean

$$SSY(t) = \sum_{mg} [\text{Temp}_{mg}(t) - \hat{\mu}(t)]^2$$

- The corresponding mean squared error functions are

$$MSE(t) = SSE(t)/df(\text{error})$$

$$MSR(t) = \frac{SSY(t) - SSE(t)}{df(\text{model})}$$

Multiple correlation and F-ratio functions


- The squared multiple correlation function is

$$RSQ(t) = [SSY(t) - SSE(t)]/SSY(t).$$

- and the F-ratio function is

$$FRATIO(t) = \frac{MSR(t)}{MSE(t)}.$$

R^2 and F -ratio plots



Estimating the regression functions $\beta_j(t)$

- We want a general framework for estimating functional parameters in this and other linear models.
- We want to be able to penalize the roughness of any parameter β_j .
- We also want the capacity to estimate confidence intervals for a parameter,
- and for functionals $\rho(\beta_j)$ of a parameter.

Some basis function expansions for $\beta_j(t)$

- Let the regression coefficient vector $\beta(t)$ have the expansion

$$\beta(t) = \mathbf{B}\theta(t)$$

where matrix \mathbf{B} is q by K_β and the K_β basis functions $\theta_\ell(t)$ are contained in vector $\theta(t)$.

- In the temperature example, it would be natural to use a certain number K_β of Fourier basis functions.

A roughness penalty for $\beta_j(t)$

- If the response curves in $\mathbf{y}(t)$ are rough, we may want to impose some smoothness on the estimated β_j 's.
- Let L be a linear differential operator, such as $L = D^2$, that defines variation $L\beta(t)$ that we wish to penalize.
- Our roughness penalty on $\beta(t)$ is

$$\text{PEN}(\beta) = \int [L\beta(s)]' [L\beta(s)] ds .$$

The penalized least squares criterion

- Let the response function vector $\mathbf{y}(t)$ have the basis function expansion in terms of K_y basis functions $\phi_k(t)$:

$$\mathbf{y}(t) = \mathbf{C}\phi(t)$$

- Then the penalized least squares function is

$$\begin{aligned} \text{PENSSE}(y|\beta) &= \int (\mathbf{C}\phi - \mathbf{ZB}\theta)' \mathbf{W} (\mathbf{C}\phi - \mathbf{ZB}\theta) \\ &\quad + \lambda \int (\mathbf{LB}\theta)' (\mathbf{LB}\theta) . \end{aligned}$$

Functional probes or contrasts

- Estimating the entire regression function $\beta_j(t)$ is fine,
- but we want to focus our attention on local or specific shape features of $\beta_j(t)$.
- Perhaps, for example, we want to examine the behavior of the temperature coefficient functions in mid-winter.
- A functional *probe* or *contrast* is of the form

$$\rho(\beta) = \int \xi(s) \beta_j(s) ds$$

- $\xi(s)$ is a weight function that we choose so as to concentrate our attention on a local region, or to look for specific patterns of variation in $\beta_j(t)$.
- There no particular need for $\xi(s)$ to integrate to 0.

- When $\beta_j(s)$ has the basis function expansion

$$\beta_j(s) = \mathbf{B}_j \boldsymbol{\theta}(s),$$

where \mathbf{B}_j is the j th row of \mathbf{B} , the contrast becomes

$$\rho(\beta) = \mathbf{B}_j \int \xi(s) \boldsymbol{\theta}(s) ds$$

Some examples

- *Point evaluation:*

$$\xi(s) = \delta(s - t)$$

This simply produces the function value $\beta(t)$.

- *Local behavior:* Assuming that β is periodic, we can use

$$\xi(s) = \exp[(s - t)^2 / (2\sigma)]$$

to assess the behavior of β in a neighborhood of t of a size determined by constant σ .

How do I work out confidence limits for these probes?

- The random element in a linear model is the residual function value

$$\epsilon_i(t_j) = y_{ij} - x_i(t_j).$$

- Any linear function of the data inherits it's variance from the variance of the data.
- The variance of the data conditional on the model is the variance of the residuals.

- We have two tasks:
 - Estimate the variance of the residuals for a single response. (The mean can usually be taken to be 0.) Let's call this Σ_e .
 - Assuming independence of the observations, the variance of the whole response data matrix is

$$\text{Var}[\text{vec}(\mathbf{Y})] = \Sigma_e \otimes \mathbf{I}.$$

- Work out the linear mapping from the data to the probe $\rho(\beta_j)$ that is being estimated. Let us call this \mathbf{M}_j .
- The rest is easy:

$$\text{Var}[\rho(\beta_j)] = \mathbf{M}_j'(\Sigma_e \otimes \mathbf{I})\mathbf{M}_j$$

Some cautionary notes

- These sampling variances would only be “exact” if we knew Σ_e . The value of our confidence limit estimates depends critically on the quality of the estimate of Σ_e . There are many open questions about how to do this.
- We are assuming that the distribution of a probe is well summarized by its mean and variance.
- Our estimates are all conditioned on how many basis functions we use for both $y_i(t)$ and $\beta_j(t)$, namely K_y and K_β . Since we never know exactly how many to use, these should be regarded as random quantities, and a Bayesian treatment seems to be indicated.
- We should back up the use of these “delta method” confidence regions by bootstrapping and simulations.

Summary

- Regressing a functional response on multivariate independent variables or on a design matrix is not much different from the conventional regression analysis.
- One important difference is that we want to do *local* inference and interval estimation.
- We have, too, the capacity to smooth estimated functional parameters.
- But the number of basis functions that we use is not a fixed parameter in the traditional sense.

Outline

- 1 A Quick Introduction to Functional Linear Models
- 2 Modelling functional responses with multivariate covariates
- 3 Modelling functional responses with multivariate covariates
- 4 Modelling scalar responses with functional covariates**
- 5 Functional responses, functional covariates and the concurrent model
- 6 Functional responses and functional covariates: The general case

- With functional responses and multivariate independent variables we could estimate the regression coefficient functions without necessarily needing to use roughness penalties.
- The same with functional responses, functional independent variables and the concurrent model.
- Now we look at a scalar response predicted by a functional independent variable, and discover that a roughness penalty or *regularization* is indispensable.

A model for total annual precipitation


- Let $y_i = \text{LogPrec}_i$ be the logarithm of total annual precipitation at weather station i .
- Here is our model:

$$\text{LogPrec}_i = \alpha + \int_0^T \text{Temp}_i(s) \beta(s) ds + \epsilon_i .$$


- We can think of the function values $\text{Temp}(s)$ associated with each fixed s as a separate scalar independent variable.
- If so, we have enough fitting power at our disposal to fit any number of responses, and certainly only 35 of them.

A bad idea

- If we use the discrete daily temperature averages, we have 365 plus 1 for constant α independent variables to fit 35 responses.
- Using the Moore-Penrose generalized inverse to keep us out of trouble, we get the following estimate of β .



Using only monthly values doesn't help much.



Estimating $\beta(s)$ with a roughness penalty

- We could impose smoothness on $\beta(s)$ by expanding it in terms of a small number (< 35) of basis functions.
- Using a roughness penalty, however, gives us continuous control over smoothness and other advantages.
- Here is the penalized least squares criterion:

$$\begin{aligned} \text{PENSSSE}_{\lambda}(\alpha, \beta) &= \sum_{i=1}^N [y_i - \alpha - \int z_i(s) \beta(s) ds]^2 \\ &+ \lambda \int [L\beta(s)]^2 ds, \end{aligned}$$

Choosing a roughness penalty


- Let's penalize *harmonic acceleration* because we want $\beta(s)$ to be periodic:

$$L\beta(s) = \left(\frac{2\pi}{365}\right)^2 D\beta(s) + D^3\beta(s)$$


- We choose the smoothing parameter λ by minimizing the cross-validation criterion.
- Let $\alpha_\lambda^{(-i)}$ and $\beta_\lambda^{(-i)}$ be the estimates using all the responses except y_i .
- The criterion to be minimized is

$$CV(\lambda) = \sum_{i=1}^N [y_i - \alpha_\lambda^{(-i)} - \int z_i(s) \beta_\lambda^{(-i)}(s) ds]^2$$


A plot of $CV(\lambda)$



$\beta(s)$ for $\log_{10} \lambda = 12.5$



A plot of the fit



Summary

- Either dimension reduction or regularization is essential when the dimensionality of the covariate exceeds the dimensionality of the response.
- Functional covariates for scalar responses has been heavily researched.
 - The group STAPH at U. Toulouse uses functional PCA to reduce covariate dimensionality. Go to `www.lsp.ups-tlse.fr/FP/Ferraty/staph.html` to learn more.
 - Gareth James, the group at the U. of Granada, and Doug Clarkson at Insightful Corp. have introduced functional covariates into the generalized linear model.

Outline

- 1 A Quick Introduction to Functional Linear Models
- 2 Modelling functional responses with multivariate covariates
- 3 Modelling functional responses with multivariate covariates
- 4 Modelling scalar responses with functional covariates
- 5 Functional responses, functional covariates and the concurrent model**
- 6 Functional responses and functional covariates: The general case

Predicting precipitation profiles from temperature curves

- Precipitation is much harder to predict than temperature.
- It comes in two main forms:
 - *Drizzle*: Large low pressure zones drop moisture over many hours or days.
 - *Storms*: Convective, short violent storms with a lot of precipitation in a hurry, and spatially localized.
- Precipitation tends to be seasonal; more in the spring and fall than in the summer and winter.

A model

- We can assume that climate zone is important.
- We will predict log precipitation; logging stabilizes variance and eliminates the positivity constraint.
- We will use the difference $\text{TempRes}_{mg}(t)$ between a temperature profile and the mean for the climate zone as a function covariate.

- We can extend the functional ANOVA model to


$$\log[\text{Prec}_{mg}(t)] = \mu(t) + \alpha_g(t) + \text{TempRes}_{mg}(t)\beta(t) + \epsilon_{mg}(t)$$

- We call this model *concurrent* because it assumes that the temperature today affects today's precipitation.

The functional data

- Where precipitation was recorded as 0 mm, we changed it to 0.05 mm, half the minimum positive value.
- We used 11 Fourier series basis functions for precipitation with no roughness penalty.
- We used 21 Fourier series basis functions for temperature with no roughness penalty.

Log precipitation profiles




The fitting criterion and some results

- The fitting criterion was the unpenalized error sum of squares


$$\begin{aligned} \text{LMSSE}(\mu, \alpha_g, \beta) = \\ \int \sum_{k,g}^N [\text{LogPrec}_{kg}(t) - \mu(t) - \alpha_g(t) \\ - \text{TempRes}_{kg}(t)\beta(t)]^2 dt \end{aligned}$$

- The resulting root-mean-squared-residual was 0.19 mm.
- When we dropped $\text{TempRes}(t)$ from the model, this increased to 0.20 mm.
- As we see in the following plot, the only place where temperature appears to make a contribution is in mid-winter.

The estimated regression function $\beta(t)$



The fit to Vancouver's data



A probe for the winter effect

- The confidence limits are point-wise; we need a measure of the temperature influence accumulated over the winter months.
- Here is a probe that works:

$$\int_0^{365} \cos[2\pi(t - 64.5)/365]\beta(t) dt = 2.32 ,$$

- The estimated standard error of this probe is 0.77, giving a t-ratio of 3.0.
- It appears that elevated temperatures in mid-winter go along with increased precipitation.

Fitting the concurrent model

- Here is a general statement of the concurrent functional/functional model:

$$y_i(t) = \sum_{j=1}^q z_{ij}(t)\beta_j(t) + \epsilon_i(t) .$$


- or in matrix notation:

$$\mathbf{y}(t) = \mathbf{Z}(t)\boldsymbol{\beta}(t) + \boldsymbol{\epsilon}(t) ,$$


- We will use a penalized error sum of squares criterion:

$$\begin{aligned} \text{LMSSE}(\boldsymbol{\beta}) = & \int [\mathbf{y}(t) - \mathbf{Z}(t)\boldsymbol{\beta}(t)]' [\mathbf{y}(t) - \mathbf{Z}(t)\boldsymbol{\beta}(t)] dt \\ & + \sum_j^p \lambda_j \int [L_j \beta_j(t)]^2 dt . \end{aligned}$$

Evolution in seasonal trend for the log nondurable goods index



Four seasonal trends




- Seasonal trends are stable over a couple of years, but evolve over a longer time span.
- We can model nonseasonal trend plus an evolving seasonal trend as follows:

$$y(t) = \alpha(t) + \beta_1(t) \sin(2\pi t/365) + \beta_2(t) \cos(2\pi t/365) + \dots \\ + \beta_{p-1}(t) \sin(p\pi t/365) + \beta_p(t) \cos(p\pi t/365) + \epsilon(t)$$

- For the monthly index values from 1952 to 2000 we used $p = 10$.
- Intercept function α was modelled by B-splines with knots at each year and regularized with $\lambda = 0.01$.
- Each regression function β_j had 7 B-spline basis functions.
- A total of 121 parameters were estimated from 577 data points.

The fit to the data over seven years



Summary

- The concurrent functional linear model offers a simple way of relating a functional response to functional covariates.
- However, the influence is simultaneous, and does not permit a covariate to affect the outcome at any time other than the present.
- The model can also be fit to a single long time series provided that the number of parameters is kept small and/or regularization is used.

Outline

- 1 A Quick Introduction to Functional Linear Models
- 2 Modelling functional responses with multivariate covariates
- 3 Modelling functional responses with multivariate covariates
- 4 Modelling scalar responses with functional covariates
- 5 Functional responses, functional covariates and the concurrent model
- 6 Functional responses and functional covariates: The general case**

Overview

- The fundamental question is this: If we have a covariate $z(s)$, how much of its variation over s should we use to fit the response $y(t)$ at fixed t ?
- In the simplest case, we only use $z(t)$. We call this the *concurrent model* because it predicts y at time t by z at time t . We might call this “now-casting.”
- If we want to use the behavior of z over an interval of values s , or over all values, things are more complex because this model has, effectively, an infinite amount of fitting power.

The full model for log precipitation


- We now want to predict the log precipitation profile $\text{LogPrec}_i(t)$ at time t from the entire temperature profile $\text{Temp}_i(s)$.

- The fitting criterion is


$$\text{LogPrec}_i(t) = \alpha(t) + \int_0^{365} \text{Temp}_i(s) \beta(s, t) ds + \epsilon_i(t) .$$

- $\beta(s, t)$ indicates the influence of temperature at time s on precipitation at time t .
- We can use the whole temperature profile because the data are periodic.
- We have already learned from predicting total log precipitation that we will have to apply a roughness penalty to $\beta(s, t)$ as a function of s .
- What about its variation as a function of t ?

Log precipitation functions




Temperature functions




- We apply two harmonic acceleration roughness penalties to $\beta(s, t)$, one for its variation in s , and one for its variation in t .
- Let's see what happens with fairly light penalties on both types of variation.
- We'll look at $\beta(s, t)$ and at the fit to the log precipitation data for Vancouver.


$\beta(s, t)$ has light penalties on s and t




$\beta(s, t)$ has light penalties on s and t




$\beta(s, t)$ has heavy penalty on s and light on t




$\beta(s, t)$ has heavy penalty on s and light on t



$\beta(s, t)$ has heavy penalties on both s and t




$\beta(s, t)$ has heavy penalties on both s and t



What we see

- Penalizing the roughness of $\beta(s, t)$ as a function of s prevents over-fitting.
- Penalizing the roughness of $\beta(s, t)$ as a function of t allows us to see how the influence of temperature on precipitation varies from one time to another.
- We can now see that temperature is much more influential in the winter than in the summer.
- The rapid oscillation in s suggests that it is a derivative of temperature that really influences precipitation.

The intercept function $\alpha(t)$



The historical model and other possibilities

- We were able to use all of $z(s)$ to predict $y(t)$ in the weather example because the data were periodic.
- In nonperiodic situations, it would only be meaningful to use the values of s *up to* t .

$$y(t) = \alpha(t) + \int_{t-\delta(t)}^t z(s)\beta(s, t) ds + \epsilon(s)$$

- For the lower limit of integration $t - \delta(t)$, the width $\delta(t)$ of the interval of integration can vary over t or can be constant.
- We can call this the *historical linear model*.
- See Ramsay and Silverman (2002) for an example.

- More generally, we can integrate over a set Ω_t that varies with t ,
- and also permit the covariate function z to vary over t as well,


$$y(t) = \alpha(t) + \int_{\Omega_t} z(s, t) \beta(s, t) ds + \epsilon(s)$$

- How do we construct basis function systems for complex regions of integration?
- The triangular mesh algorithms used in finite element methods to solve partial differential equations are natural here.
- Triangular meshes adapt well to nonstandard boundaries.
- Triangular basis functions also can be regularized.
- The PDEtools toolbox in Matlab contains powerful functions for mesh generation.

Lip acceleration predicted from EMG signal

- Malfait and Ramsay (2004) studied how the acceleration of the lower lip while saying "bob" was related to the EMG signal record from the depressor lip muscle. (see also Ramsay and Silverman, 2002)
- Each of 32 replications lasted for about 0.7 seconds.
- EMG is an indirect indication of muscle activation.
- It was hoped to learn something about the brain controls speech production.

32 replications of lip position and acceleration and of EMG signal




The feed–forward model

- Only feed–forward effects of EMG on acceleration are of interest.
- The model is

$$lip(t) = \alpha(t) + \int_{t-\delta}^t EMG(s)\beta(s, t) ds + \epsilon(s)$$

- Regression function $\beta(s, t)$ is defined over the triangular region $0 \leq s \leq t$; $0 \leq t \leq 0.7$.
- How far back should we allow for an influence? How large should the lag δ be?


The correlation surface




The finite element basis for $\beta(s, t)$

- The finite element basis for functions defined over two arguments uses piece-wise linear functions defined over a triangular mesh.
- Triangular meshes can easily adapt to complex boundaries. Our problem here is particularly easy.
- The coefficient matrices become more and more sparse as the number of triangles increases.
- Sparse matrix computation, available in Matlab, makes for fast solutions.


A triangular mesh for the lip/EMG problem



The regression function surface for lag $\delta = 5$ triangles



Standard error of estimate functions for various lags



Summary

- When both response and covariates are functional, there are a lot of modelling possibilities.
- Using the full variation in a functional covariate $z(s)$ is only useful if we regularize the solution, either at the level of the regression coefficient $\beta(s, t)$, or at the level of the response $y(t)$.
- Using the full variation in the covariate usually only makes sense for periodic problems.
- A feed-forward model is more likely for nonperiodic functional regressions.
- Estimating the amount of functional history to use is an important issue.
- There is a lot more work to do in this exciting area!

Where to we go from here?

- What if we want to use derivatives, $Dy(t)$ and $Dz(t)$, in our models?
- What about functional variables that interact with each other, such as we find in input/output systems with feedback?
- Nonlinear functional models?
- We have only tickled the ear of the elephant.