Module 4: Coping with Multiple Predictors

Multidimensional Splines

STAT/BIOSTAT 527, University of Washington
Emily Fox
May 9th, 2013

Nonparam. Multiple Regression

- We now consider a *d*-dimensional covariate x_i
- In its most general form, the regression equation then takes the form
- In principle, all of the methods we have discussed so far carry over to this case rather straightforwardly
- Unfortunately, the risk of the nonparametric estimator increases rapidly with covariate dimension *d*.

Curse of Dimensionality

- To maintain a fixed level of accuracy for a given nonparametric estimator, the sample size must increase exponentially in *d*
- Set MSE = δ
- Why? Using data in local nbhd
  - In high dim, few points in any nbhd
- Consider example with *n* uniformly distributed points in $[-1,1]^d$
  - d=1:
  - d=10

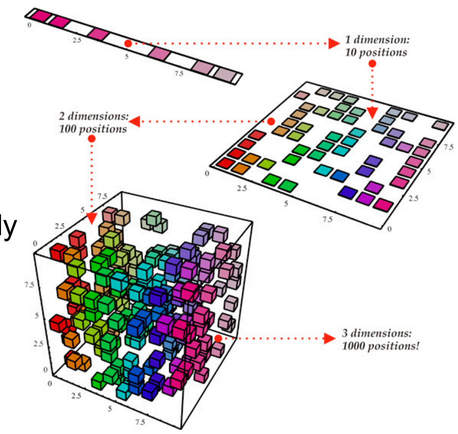

Figure from Yoshua Bengio's website

Natural Thin Plate Splines

- One-dimensional smoothing splines (obtained via regularization) can be extended to the multivariate setting as the solution to

$$\min_f \sum_{i=1}^{n} \{y_i - f(x_i)\}^2 + \lambda J(f)$$

- Recall roughness penalty in 1d

$$J(f) = \int f''(x)^2 dx$$

- The natural 2d extension to penalize rapid variation in either dim is

$$J(f) =$$

- Is the penalty affected by rotation or translation in $\mathbb{R}^2$?

Natural Thin Plate Splines

$$\min_f \sum_{i=1}^n \{y_i - f(x_i)\}^2 + \lambda J(f)$$

$$J(f) = \int\int_{\mathbb{R}^2} \left[\left(\frac{\partial^2 f(x)}{\partial x_1^2} \right)^2 + 2 \left(\frac{\partial^2 f(x)}{\partial x_1 x_2} \right)^2 + \left(\frac{\partial^2 f(x)}{\partial x_2^2} \right)^2 \right] dx_1 dx_2$$

- Solution: Unique minimizer is the ***natural thin plate spline*** with knots at the x_{ij}
- Proof: See Green and Silverman (1994) and Duchon (1977)
- Similar properties and intuition as in 1d:
  - As λ→0,
  - As λ→∞,

Natural Thin Plate Splines

$$\min_f \sum_{i=1}^n \{y_i - f(x_i)\}^2 + \lambda J(f)$$

$$J(f) = \int\int_{\mathbb{R}^2} \left[\left(\frac{\partial^2 f(x)}{\partial x_1^2} \right)^2 + 2 \left(\frac{\partial^2 f(x)}{\partial x_1 x_2} \right)^2 + \left(\frac{\partial^2 f(x)}{\partial x_2^2} \right)^2 \right] dx_1 dx_2$$

- Solution: ***natural thin plate spline*** with knots at the x_{ij}
- For general λ, solution is a linear basis expansion of the form

$$f(x) =$$

with

$$h_j(x) = ||x - x_j||^2 \log ||x - x_j||$$

- Interpretation: We take an elastic flat plate that interpolates points (x_i,y_i) and penalize its "bending energy"

Natural Thin Plate Splines

$$f(x) = \beta_0 + \beta^T x + \sum_{j=1}^{n} b_j h_j(x)$$

- Coefficients are found via standard penalized LS

$$\min_{\beta,b} (y - X\beta - Eb)^T (y - X\beta - Eb) + \lambda b^T E b$$

s.t. $\sum_i b_i = \sum_i b_i x_{i1} = \sum_i b_i x_{i2} = 0$

- Interpretation: We take an elastic flat plate that interpolates points (x_i,y_i) and penalize its "bending energy"

Complexity of Thin Plate Splines

- Natural thin plate splines place knots at every location x_{ij}
- Computational complexity scales as O(n^3)
  - Can get away with fewer knots
  - If we use K knots, then computational complexity reduces to O($nK^2 + K^3$)
- Can choose some lattice of knots

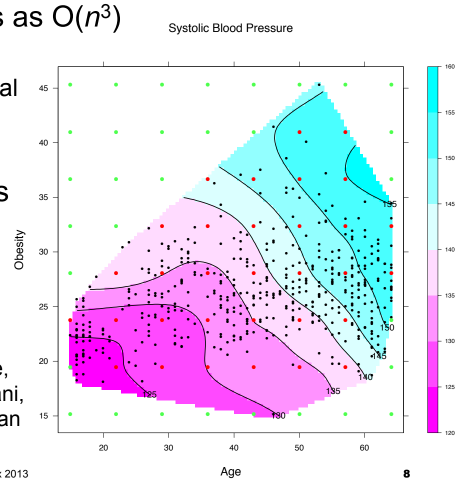

From Hastie, Tibshirani, Friedman book

Thin Plate Regression Splines

- Thin plate regression splines truncate the "wiggly" basis b_i
- Let $E = UDU^T$

- Grab out largest k eigenvalues and eigenvectors

- Define $b = U_k b_k$
- Minimize

$$\min_{\beta, b_k} (y - X\beta - U_k D_k b_k)^T (y - X\beta - U_k D_k b_k) + \lambda b_k^T D_k b_k$$
$$X^T U_k b_k = 0$$

- Optimal approximation of thin plate splines using low rank basis
- Retain advantages of (i) no choice of knots, (ii) rotation invariance
- See Wood (2006) for more details

Tensor Product Splines

- Again, assume x in $\mathbb{R}^2$
- Instead of thin plate splines, consider modeling $f(x)$ as follows
- Suppose for each dimension we have a basis of functions

- Then the M_1 x M_2 dimensional ***tensor product basis*** is

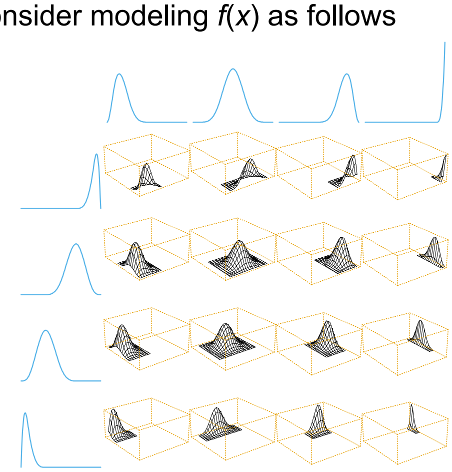

From Hastie, Tibshirani, Friedman book

Tensor Product Splines

- We use this tensor product basis

  $$g_{jk}(x) = h_{1j}(x_1)h_{2k}(x_2)$$

  to model *f*(*x*)
- This formulation extends (in theory) to any dimension *d*
- Note that as the dimension of the basis grows exponentially with the input dimension *d*

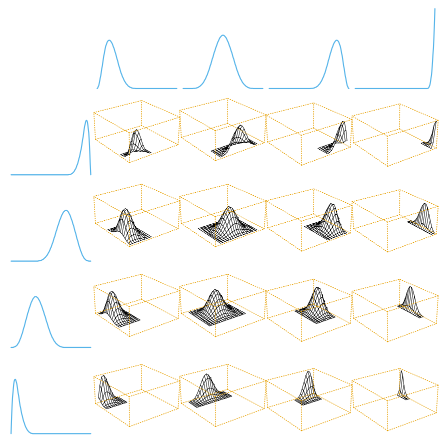

From Hastie, Tibshirani, Friedman book

Tensor Product Splines Example

- Linear spline basis with L_1 truncated lines for x_1 and L_2 for x_2

  $$1, x_1, (x_1 - \xi_{11})_+, \ldots, (x_1 - \xi_{1L_1})_+$$
  $$1, x_2, (x_2 - \xi_{21})_+, \ldots, (x_1 - \xi_{2L_2})_+$$
- Then, the tensor product expansion is

  $$f(x_1, x_2) =$$
- Number of parameters:
- Note: Captures interaction terms between x_1 and x_2

Tensor Product Splines Example

- For prostate cancer dataset, fits of log PSA as a function of log cancer volume and log weight for various models

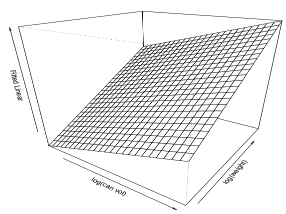

Linear fit

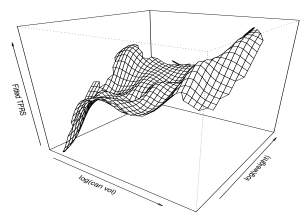

Thin plate regression spline

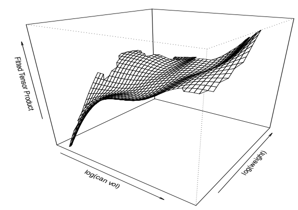

Tensor product spline

From Wakefield textbook

Generalized Additive Models

- Both for computational reasons and added interpretability, models that assume an additive structure are very popular
- Assuming a GLM framework:

$$g(\mu(x)) =$$

- Is this model identifiable?
- Can model $f_j(x_j)$ using any smoother

GAM Example

- Consider using a penalized regression spline of order p_j with L_j knots for each covariate x_j

$$g(\mu) =$$

- Penalization is applied to the spline coefficients b_j

$$\lambda_j \sum_{\ell=1}^{L_j} b_{j\ell}^2$$

Comments:

- The GAM is very interpretable
  - $f_i(x_i)$ is not influenced by the other $f_j(x_j)$
  - Can plot f_j to straightforwardly see the relationship between x_i and y
- Will see that this also leads to computational efficiencies

Backfitting

- To begin, assume a standard (non-GLM) regression setting

- For concreteness, consider

- Result is an ***additive cubic spline model*** with knots at the unique values of x_{ij}
  - For X full column rank, can show that solution is unique. Otherwise, linear part of $f_j(x_j)$ is not uniquely determined

- Here, clearly $\hat{\alpha} =$

- How do we think about fitting the other parameters??

Backfitting

- ***Backfitting*** is an iterative fitting procedure
- Since $f(x)$ is additive, if we condition on the fit of all other components $f_j(x_j)$, $j \neq i$, then we know how to fit $f_i(x_i)$
- Iterate the estimation procedure until convergence

Backfitting Algorithm

Algorithm 9.1 *The Backfitting Algorithm for Additive Models.*

1. Initialize: $\hat{\alpha} = \frac{1}{N}\sum_1^N y_i$, $\hat{f}_j \equiv 0, \forall i, j$.
2. Cycle: $j = 1, 2, \ldots, p, \ldots, 1, 2, \ldots, p, \ldots,$

$$\hat{f}_j \leftarrow \mathcal{S}_j\left[\{y_i - \hat{\alpha} - \sum_{k \neq j} \hat{f}_k(x_{ik})\}_1^N\right],$$

$$\hat{f}_j \leftarrow \hat{f}_j - \frac{1}{N}\sum_{i=1}^N \hat{f}_j(x_{ij}).$$

until the functions $\hat{f}_j$ change less than a prespecified threshold.

From Hastie, Tibshirani, Friedman book

GAMs and Logistic Regression

- A generalized additive logistic regression model has the form
- The functions $f_1, \ldots, f_d$ can be estimated using a backfitting algorithm, too
- First, recall IRLS algorithm for *parametric* logistic regression

$$z = X\beta^{\text{old}} + W^{-1}(y - p)$$

$$\beta^{\text{new}} \leftarrow \arg\min_{\beta} (z - X\beta)^T W (z - X\beta)$$

GAMs and Logistic Regression

Algorithm 9.2 *Local Scoring Algorithm for the Additive Logistic Regression Model.*

1. Compute starting values: $\hat{\alpha} = \log[\bar{y}/(1-\bar{y})]$, where $\bar{y} = \text{ave}(y_i)$, the sample proportion of ones, and set $\hat{f}_j \equiv 0 \; \forall j$.
2. Define $\hat{\eta}_i = \hat{\alpha} + \sum_j \hat{f}_j(x_{ij})$ and $\hat{p}_i = 1/[1 + \exp(-\hat{\eta}_i)]$.

   Iterate:

   (a) Construct the working target variable
   $$z_i = \hat{\eta}_i + \frac{(y_i - \hat{p}_i)}{\hat{p}_i(1-\hat{p}_i)}.$$
   (b) Construct weights $w_i = \hat{p}_i(1-\hat{p}_i)$

   (c) Fit an additive model to the targets z_i with weights w_i, using a weighted backfitting algorithm. This gives new estimates $\hat{\alpha}, \hat{f}_j, \; \forall j$
3. Continue step 2. until the change in the functions falls below a prespecified threshold.

From Hastie, Tibshirani, Friedman book

GAM Logistic Example

- Example: *predicting spam*
- Data from UCI repository
- Response variable: *email* or *spam*
- 57 predictors:
  - 48 quantitative – percentage of words in email that match a give word such as “business”, “address”, “internet”,…
  - 6 quantitative – percentage of characters in the email that match a given character (; , [! $ #)
  - The average length of uninterrupted capital letters: CAPAVE
  - The length of the longest uninterrupted sequence of capital letters: CAPMAX
  - The sum of the length of uninterrupted sequences of capital letters: CAPTOT

GAM Logistic Example

- Test set of 1536 emails
- Training set: n=3065
- Use a GAM with a cubic smoothing spline
  - Each with 4 dof
- Estimated functions for significant predictors
  - Note large discontinuity near 0 for many
- Test error of 6.6%

From Hastie, Tibshirani, Friedman book

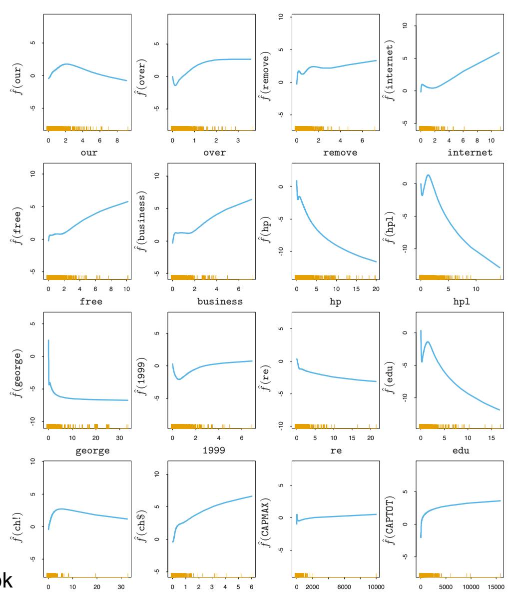

Other GAM formulations

- Semiparametric models:

$$g(\mu) =$$

- ANOVA decompositions:

$$f(x) =$$

Choice of:
  - Maximum order of interaction
  - Which terms to include
  - What representation

- Tradeoff between full model and decomposed model

Connection with Thin Plate Splines

- Recall formulation that lead to natural thin plate splines:

$$\min_f \sum_{i=1}^{n} \{y_i - f(x_i)\}^2 + \lambda J(f)$$

$$J(f) = \int\int_{\mathbb{R}^2} \left[\left(\frac{\partial^2 f(x)}{\partial x_1^2} \right)^2 + 2\left(\frac{\partial^2 f(x)}{\partial x_1 x_2} \right)^2 + \left(\frac{\partial^2 f(x)}{\partial x_2^2} \right)^2 \right] dx_1 dx_2$$

- There exists a *J*(*f*) such that the solution has the form

- However, it is more natural to just assume this form and apply

$$J(f) = J(f_1 + f_2 + \cdots + f_d) = \sum_{j=1}^{d} \int f_j^{''}(t_j)^2 dt_j$$

What you need to know

- Nothing is conceptually hard about multivariate *x*
- In practice, nonparametric methods struggle from curse of dimensionality
- Options considered:
  - Thin plate splines
  - Tensor product splines
  - Generalized additive models
  - Combinations (to model some interaction terms)

Readings

- Wakefield – 12.1-12.3
- Hastie, Tibshirani, Friedman – 5.7, 9.1
- Wasserman – 4.5, 5.12

Survey Feedback

- Lectures:
  - Useful to post reading assignments → will do!
  - Lots of material, so make clear what is expected to know → will do!

- Homeworks
  - More frequent and more in-depth
  - Less frequent/intense
  - → ???

- Recitations
  - Make same week as HW due…Was original plan and will reset to this.