

Eigenstuff: Almost everything you ever wanted to know about

### **Eigenvalues & eigenvectors**

...but were too afraid to ask



## **Eigenstuff: Statistical motivations**

- Same ideas are used in low-D approximations to visualize high-D data, models in 2 (or 3) dimensions
- Fundamental for generalizing univariate (F) tests to multivariate tests (e.g., Wilk's Λ, Hotelling trace, ...) in MReg & MANOVA
  - Eigenvalues of **H E**<sup>-1</sup> → multivariate test statistics



# **Eigenstuff: Statistical motivations**

- Usually used in analysis of a covariance (S) or correlation matrix (R)
  - PCA: principal components are just eigenvectors of S or R; eigenvalues give the % of variance accounted for
  - FA: similar, but based on a "reduced" **S** or **R**, accounting for only common variance



## Eigenstuff: Everyday applications



## Eigenstuff: Everyday applications

- Google Page Rank algorithm (run often)
  - Importance of a page (r<sub>i</sub>): a function of
    - # of other pages  $(L_i)$  linking to it
    - the importance of each linking page

Network diagram of page links





paths thru

network

5

But: the WWW has > 30 billion pages! – Do this bit by bit, many robots, many servers

## **Eigenstuff: Everyday applications**

- Start with  $r^{(0)} = 1$ . Condition for a stable solution: H r = r
- $\rightarrow$  **r** = eigenvector of H assoc. with  $\lambda$ =1 (largest).
- Solution: power method:  $\mathbf{r}^{(k+1)} = H \mathbf{r}^{(k)}$

	0.0600
	0.0675
	0.0300
r =	0.0675
	0.0975
	0.2025
	0.1800
	0.2950
	L .





See: http://www.ams.org/featurecolumn/archive/pagerank.html

Lighter = more important

## Emerging applications: Eigenfaces

Computer vision & automatic face recognition

Digitize, align & normalize a large # of pics

- Each  $m \ge n$  pic  $\rightarrow mn$  vector of pixel values
- Calculate covariance matrix, S, of images
- Eigenvectors of S are "eigenfaces"

Any new face can be approximated by a linear combination of eigen faces





#### This is one reason why you can't smile in passport photos!

7

# **Eigenstuff: Definitions**

For a square matrix, A<sub>(n x n)</sub>, a vector v is an eigenvector of A, if there exists a scalar, λ, for which

$$\mathbf{A}_{n \times n} \mathbf{v}_{n \times 1} = \lambda \mathbf{v}_{n \times 1} \quad \text{or,} \quad \begin{pmatrix} \mathbf{a}_{11} & \dots & \mathbf{a}_{1n} \\ \vdots & \ddots & \vdots \\ \mathbf{a}_{n1} & \cdots & \mathbf{a}_{nn} \end{pmatrix} \begin{pmatrix} \mathbf{v}_{1} \\ \vdots \\ \mathbf{v}_{n} \end{pmatrix} = \lambda \begin{pmatrix} \mathbf{v}_{1} \\ \vdots \\ \mathbf{v}_{n} \end{pmatrix}$$

- The value λ is the corresponding *eigenvalue*
- Other terms: latent vectors (v), latent roots (λ), characteristic vectors/roots.

### **Eigenstuff: Geometric motivation**

Linear transformation of Mona Lisa



Eigenvectors arise from what happens to a vector  ${\bf v}$  under transformation by a matrix  ${\bf A}$ 

Here, Mona is subjected to a horizontal shear transform, given by

$$\mathbf{A} = \begin{bmatrix} 1 & k \\ 0 & 1 \end{bmatrix} \longrightarrow \text{Mona x A}$$

The red vector (0,1) is **unchanged** in direction– an **eigenvector** of A

$$\mathbf{A} \mathbf{v} = \lambda \mathbf{v}$$

The blue vector (1,1) is changed in direction– not an eigenvector

**Αν** ≠ λ **ν** 

9

## **Eigenstuff: Geometric motivation**



## **Eigenstuff: Geometric motivation**

Thus,  $\mathbf{v}_1$  and  $\mathbf{v}_2$  are special vectors with respect to the matrix **A**, which are mapped to vectors in the same direction, i.e., for which

 $\mathbf{A} \mathbf{v} = \lambda \mathbf{v} \qquad \lambda = \text{scalar}$ 



When this relation holds:

- v is called an eigenvector (latent vector) of the matrix A
- λ is the corresponding eigenvalue (latent root)

#### An n x n matrix has n (eigenvector – eigenvalue) pairs

## Finding eigen solutions

• Finding eigenvectors of a matrix naturally gives rise to a set of homogeneous equations, since:

$$\mathbf{A}\mathbf{v} = \lambda \mathbf{v} \Longrightarrow \mathbf{A}\mathbf{v} - \lambda \mathbf{v} = \mathbf{0} \Longrightarrow (\mathbf{A} - \lambda \mathbf{I})\mathbf{v} = \mathbf{0}$$

- This system has *non-trivial* solutions only for values of λ for which rank (**A** – λ**I**) < n</li>
- This implies  $(\mathbf{A} \lambda \mathbf{I})$  is singular  $\rightarrow |\mathbf{A} \lambda \mathbf{I}| = 0$
- This *characteristic equation* → one way to find eigenvalues
- (Modern software uses more efficient methods)

## Finding eigen solutions

For a 2 x 2 symmetric matrix, A,

$$|\mathbf{A} - \lambda \mathbf{I}| = \begin{vmatrix} a_{11} & a_{12} \\ a_{12} & a_{22} \end{vmatrix} - \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{vmatrix} = \begin{vmatrix} a_{11} - \lambda & a_{12} \\ a_{12} & a_{22} - \lambda \end{vmatrix}$$
$$= (a_{11} - \lambda)(a_{22} - \lambda) - a_{12}^{2}$$
$$= \lambda^{2} - (a_{11} + a_{22})\lambda + (a_{11}a_{22} - a_{12}^{2}) = 0$$

This is a quadratic equation that can (usually) be solved for two distinct values,  $\lambda_1 \& \lambda_2$ . Writing the quadratic as

$$a\lambda^{2} + b\lambda + c = 0 \Rightarrow \lambda = \frac{-b \pm \sqrt{b^{2} - 4ac}}{2a} \qquad b = (a_{11} + a_{22}) = trace(\mathbf{A})$$
$$c = (a_{11}a_{22} - a_{12}^{2}) = det(\mathbf{A})$$

#### 13

 $\frac{{\bf v}_1^T {\bf A} {\bf v}_1}{1} = 9.54$ 

15

## Finding eigen solutions

Example: 
$$A = \begin{pmatrix} 1 & -s \\ -s & 1 \end{pmatrix}, \lambda = \frac{1}{2} \begin{pmatrix} 1+1 & \pm \sqrt{(1-s)^2 + 4(ss)^2} \end{pmatrix}$$
  
 $= \frac{1}{2} \begin{pmatrix} 2 & \pm \sqrt{1} \end{pmatrix} = \begin{cases} 1.5 & = \lambda_1 \\ .5 & = \lambda_2 \end{cases}$   
To find eigenvector associated with any eigenvalue,  
solve  $(A - \lambda I) \underbrace{V} = \underline{0}$   
e.g. for  $\lambda_1 = 1.5$ ,  $A - \lambda I = \begin{pmatrix} 1 & -s \\ -s & 1 \end{pmatrix} - \begin{pmatrix} 1.5 & 0 \\ 0 & 1.5 \end{pmatrix}$   
 $\Rightarrow \begin{pmatrix} -s & -s \\ -s & -s \end{pmatrix} \begin{pmatrix} N_{11} \\ N_{21} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \Rightarrow \underbrace{N_1} = \begin{pmatrix} N_{11} \\ N_{21} \end{pmatrix} = \begin{pmatrix} -N_{21}^2 \\ N_{21}^2 \end{pmatrix} e.q \begin{pmatrix} -1 \\ 1 \end{pmatrix}$   
 $frank = 1$   
so non-trivial sol'n exists  $\begin{pmatrix} N_{12} \\ N_{22} \end{pmatrix} = \begin{pmatrix} N_{12} \\ N_{22} \end{pmatrix} = \begin{pmatrix} N_{12} \\ N_{22} \end{pmatrix}$ 

Power method

- The power method is an easy way to find the eigenvector,  $\mathbf{v}_1$ , • corresponding to the *largest* eigenvalue,  $\lambda_1$ , among  $\lambda_1 \ge \lambda_2 \ge ... \ge \lambda_n$ .
  - Initialize:  $\mathbf{v}_1^{(0)} = 1$  (any non-zero vector)
  - Iterate:  $\mathbf{v}_1^{(k+1)} = A \mathbf{v}_1^{(k)} / ||A \mathbf{v}_1^{(k)}||$
  - Until:  $v_1^{(k+1)} \approx v_1^{(k)}$

• Then, 
$$\lambda_1 = \frac{\mathbf{v}_1^T \mathbf{A} \mathbf{v}_2}{\mathbf{v}_1^T \mathbf{v}_2}$$

 $\mathbf{A} = \begin{pmatrix} 7 & 3 \\ 3 & 6 \end{pmatrix}$ 

This is the idea behind the computation for the Google page rank algorithm

Properties of eigenvalues and eigenvectors

For an *n* x *n* matrix, **A**:

- The characteristic equation  $|\mathbf{A} \lambda \mathbf{I}| = 0$ 
  - has *n* roots,  $\lambda_1 \ge \lambda_2 \ge \ldots \ge \lambda_n$
  - if A is symmetric, roots are real (otherwise: complex)

•  $\sum_{i=1}^{n} \lambda_i = trace(\mathbf{A})$ •  $\prod_{i=1}^{n} \lambda_i = \lambda_1 \times \lambda_2 \times \cdots \times \lambda_n = det(\mathbf{A})$ 

Two measures of "size" of A in n dimensions. In MANOVA:

- Hotelling trace criterion
- Wilks' Lambda

16

#### Properties of eigenvalues and eigenvectors

- If any  $\lambda_i = 0$ , then **A** is singular
  - # non-zero λ<sub>i</sub> = rank(A) = # dimensions of A space
  - # zero  $\lambda_i$  = # linear dependencies
- If  $\lambda_i \neq \lambda_i$ , latent vectors are orthogonal & unique
  - i.e.,  $\mathbf{v}_i^{\mathsf{T}} \mathbf{v}_i = 0$
  - If  $\lambda_i = \lambda_i$ , vectors not unique, but chosen orthogonal
- Latent vectors are determined in direction, but not length
  - Usually, standardized so that length, ||v<sub>i</sub>||=1
  - Any scalar multiple, k v, also a latent vector
  - → PCA, FA: signs/scaling of loadings are arbitrary

## **Example:** R



The characteristic equation,  $|\mathbf{A} - \lambda \mathbf{I}| = 0$  gives:  $\lambda^3 - 6\lambda^2 + 9\lambda - 4 = 0$ The roots,  $\lambda$ , are: 4, 1, 1

Let's try this in R:

17

unique

19

library(matlib) (A <- J(3, 3) + diag(3)) tr(A) # sum of eigenvalues det(A) # product of eigenvalues

(ev <- eigen(A))

```
# eigenvectors are orthogonal
V <- ev$vectors
zapsmall( V %*% t(V) )
```

18

eigen() function

## **Example:** R output

> library(matlib)	> (ev <- eigen(A))
> (A <- J(3, 3) + diag(3))	\$values No 0 eigenvalues; r(A)=3
[,1] [,2] [,3]	[1] 4 1 1
[1,]       2       1       1         [2,]       1       2       1	\$vectors [1] [2] [2] Last 2 eigenvalues are equal, so eigenvectors are not unique
[3,] 1 1 2	[,'] [,2] [,3] [1,] -0.5774 0.0000 0.8165
<pre>&gt; tr(A) # sum of eigenvalues</pre>	[2,] -0.5774 -0.7071 -0.4082
[1] 6	[3,] -0.5774 0.7071 -0.4082
<pre>&gt; det(A) # product of eigenvalues [1] 4</pre>	<pre>&gt; # eigenvectors are orthogonal &gt; V &lt;- ev\$vectors &gt; zapsmall( V %*% t(V) )</pre>
	[,1] [,2] [,3] [1,] 1 0 0 [1,] 0 1 0 [1,] 0 1 0
	[3,] 0 0 1

## **Implications**

Let 
$$V_{(n\times n)} = (\underbrace{N_1}, \underbrace{N_2}, ..., \underbrace{N_n})$$
 be matrix of all n  
eigenvectors conversionding to  $\lambda_1 \ge \lambda_2 \ge \lambda_3 ... \ge \lambda_n$   
Let  $\Lambda = \begin{pmatrix} \lambda_1 & \lambda_2 & 0 \\ 0 & \ddots & \lambda_n \end{pmatrix}$  - diagonal matrix of eigenvalues

Then  
(i) 
$$A_{\underline{v}_i} = \lambda_i \underline{v}_i \implies A \vee = \sqrt{\underline{\Lambda}}$$
  
(2)  $\underline{N}_i \wedge \underline{N}_j = 0 \implies \sqrt{\underline{\Lambda}} = I$  (socied so  $\|\underline{N}_i\| = 1$ ,  
 $\sqrt{\underline{V}'} = I$ 

Here, we are just collecting the  $\mathbf{v}_i$  and  $\lambda_i$  in matrices. But this has interesting applied implications

#### Implications: Spectral decomposition



$$\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{T} = \begin{bmatrix} \mathbf{v}_{1} & \mathbf{v}_{2} & \cdots & \mathbf{v}_{n} \end{bmatrix} \begin{bmatrix} \lambda_{1} & & & \\ & \lambda_{2} & & \\ & & \ddots & \\ & & & \ddots & \\ & & & & \lambda_{n} \end{bmatrix} \begin{bmatrix} \mathbf{v}_{1}^{T} \\ \mathbf{v}_{2}^{T} \\ \vdots \\ \mathbf{v}_{n}^{T} \end{bmatrix}$$

or

 $\mathbf{A} = \lambda_1 \mathbf{v}_1 \mathbf{v}_1^T + \lambda_2 \mathbf{v}_2 \mathbf{v}_2^T + \dots + \lambda_n \mathbf{v}_n \mathbf{v}_n^T$ 

Ex:

$$\mathbf{A} = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix} = \frac{4}{-.5774} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \end{pmatrix} + \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix} \begin{pmatrix} 0 & -1 & 1 \end{pmatrix} + \frac{1}{.4082} \begin{pmatrix} 2 \\ -1 \\ -1 \end{pmatrix} \begin{pmatrix} 2 & -1 & -1 \end{pmatrix}$$

This is the basis for PCA: fewer terms can approximate A

## **Implications: Matrix powers**

(4) Eigenvalues of matrix inverse & powers:

$$\mathbf{A}^{-1} = \mathbf{V} \mathbf{\Lambda}^{-1} \mathbf{V}^{T}$$
 because, if so:  $\mathbf{A} \mathbf{A}^{-1} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{T} \mathbf{V} \mathbf{\Lambda}^{-1} \mathbf{V}^{T} = \mathbf{I}$ 

Hence,

- Eigenvectors of A-1 are the same as those of A
- Eigenvalues of  $\mathbf{A}^{-1}$  are  $1/\lambda_i$

Similarly:  $\mathbf{A}^2 = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \ \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T = \mathbf{V} \mathbf{\Lambda}^2 \mathbf{V}^T$ 

- So, eigenvalues of  $\mathbf{A}^2$  are  $\lambda_i^2$  and eigenvectors  $\mathbf{V}$  are the same
- This applies to all powers of a matrix

25

# **Implications: Orthogonality**

(5) Rearranging (3),  $\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{\mathsf{T}}$  gives:

$$\mathbf{V}^{T}\mathbf{A}\mathbf{V} = \mathbf{V}^{T}(\mathbf{V}\mathbf{\Lambda}\mathbf{V}^{T})\mathbf{V} = \mathbf{\Lambda} = \begin{bmatrix} \lambda_{1} & \lambda_{2} & \lambda_{2} & \lambda_{2} & \lambda_{2} \\ & \ddots & & \ddots & \\ 0 & & \ddots & \lambda_{n} & \lambda_{n} & \lambda_{n} \end{bmatrix}$$

Гλ

Hence, transforming **A** by **V** and  $\mathbf{V}^{\mathsf{T}}$ :

gives an orthogonal matrix – uncorrelated variables, which are linear combinations of the original ones; the weights are the eigenvectors in V.

In applications here,  ${\bm A}$  is usually a covariance matrix,  ${\bm S},$  of a matrix of variables,  ${\bm Y}.$  Then,

 $V S V' = \Lambda$  is the covariance matrix of the transformed variables, Y V. That  $\Lambda$  is diagonal means the transformed variables are uncorrelated.

## **Implications:** Maximization

(6) If v<sub>i</sub>' v<sub>i</sub> = 1(unit length), then (1) A v<sub>i</sub> = λ<sub>i</sub> v<sub>i</sub> can be written as

$$\mathbf{v}_i$$
  $\mathbf{A} \mathbf{v}_i = \lambda_i$ 

- But, **v**' **A v** is the variance of a linear combination, with weights in **v**.
- Hence, to find a v which maximizes v' A v, choose v = v<sub>1</sub>
   = eigenvector corresponding to largest eigenvalue, λ<sub>1</sub>
  - The maximum variance is then  $\lambda_1$
- In statistical applications, A = covariance matrix of a set of variables
  - Many other problems of maximizing (or minimizing) something have eigenvalue – eigenvector solutions!

A sum of rank

24

 $0^{-}$ 

1 matrices

# Maximizing variance

To see what this means, consider the relation between x & y. What is the direction along which the variance of the points is greatest?



29

31

## Maximizing variance

Imagine a vector rotating around the mean of (x,y)

The projection of points on the vector has maximum variance at one position – the eigenvector



## Maximizing variance

Now imagine that the points are connected to the vector by springs.

The vector finds its position – the eigenvector – by balancing the forces from the springs! (This is equivalent to a least squares solution for the spring lengths)



# **Geometric interpretation**

For *n* variables, we can represent data in a space with orthogonal axes

The variables are typically correlated

v' A v = constant represents an ellipse of constant distance from the mean

Usually,  $A = S^{-1}$ , the inverse of a covariance matrix



## Geometric interpretation

Latent vectors,  $\mathbf{v}_1$ ,  $\mathbf{v}_2$  are the major/minor axes of the data ellipse.

They are orthogonal to each other and at right angles to the tangents to the ellipse

Latent roots  $\lambda_1$ ,  $\lambda_2$  are the half-lengths of the major and minor axes



#### 33

## **Geometric interpretation**

Imagine a rotation of the original space to one where  $v_1$ ,  $v_2$  become the new coordinate axes

• These vectors span the same space

• They represent **uncorrelated** variables referred to orthogonal axes (principal component scores)

• Total variance of the  $x_1, x_2, ..., x_n$ variables =  $\Sigma \lambda_i$  = total variance of the  $v_1, v_2, ..., v_n$ 



## Geometric interpretation

For given data, we can see this by plotting the data and their principal component scores, together with data ellipses



## Geometric interpretation

In 3D, with animation, this can be shown more clearly



# **Application: Finding outliers**

In multivariate data, outliers are often associated with the **smallest** eigenvalue(s) of covariance matrix, **S** 



### Mathematical definition of the SVD

- Let **X** denote an  $m \ge n$  matrix of data,  $rank(X) = r, m \ge n$
- The singular value decomposition of **X** is:

 $\mathbf{X} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}^T$ 

#### where:

- **U** is an *m* x *n* matrix (observation scores),
- $\Lambda$  is an  $n \ge n$  diagonal matrix of singular values, and
- $\mathbf{V}^{\mathsf{T}}$  is also an  $n \ge n$  matrix (variable weights).

Only the first *r* values of  $\lambda_i$  are non-zero.

### Singular Value Decomposition & biplots

- Eigenvalues/vectors are usually defined just for square matrices (covariance, correlation)
- For rectangular matrices (i.e., a data matrix), an analogous result is the SVD – also called the *basic structure* of a matrix



38

### Mathematical definition of the SVD

- V: columns are the eigenvectors of X<sup>T</sup>X and form an orthonormal basis (V<sup>T</sup>V = I) for the variables
- A: diagonal, r singular values are the square roots of the eigenvalues of both XX<sup>T</sup> and X<sup>T</sup>X
- U: columns are the eigenvectors of XX<sup>T</sup> and form an orthonormal basis for the observation profiles, so that U<sup>T</sup>U = I



## If that's not clear: the SVD song

### It Had To Be U

The Singular Value Decomposition (SVD)

http://www.youtube.com/watch?v=JEYLfIVvR9I

### Matrix approximation

- Let **X** be an  $m \ge n$  matrix such that rank(**X**) = r
- If λ<sub>1</sub> ≥ λ<sub>2</sub> ≥ ... ≥ λ<sub>r</sub> are the singular values of X, then **x̂**, the rank *q* approximation of X that minimizes || X − **X̂** || , is

$$\hat{\mathbf{X}}_{m \times n} = \sum_{i=1}^{q} \lambda_i \begin{pmatrix} u_{i1} \\ \vdots \\ u_{im} \end{pmatrix} \begin{pmatrix} \mathbf{v}_{i1} & \cdots & \mathbf{v}_{in} \end{pmatrix} = \lambda_1 u_1 \mathbf{v}_1^T + \cdots + \lambda_q u_q \mathbf{v}_q^T$$
row scores

a sum of *q* rank=1 (outer) products. The variance in **X** accounted for each term is  $\lambda_1^2$ 

## Application: Image compression

- Digital images often require huge files
- For some purposes, a low-dimensional approx. based on the SVD may be sufficient
- This allows to vary the trade-off between fidelity and file size



color

600 x 465 x 3



grayscale 600 x 465 = 484k

1 dimension 600 + 465 + 1 = 1k



5(600 + 465 + 1) = 5k

43

41

# Application: Image compression

#### 1/465 dim, 83 % var, 1 Kb



Trade-off between image fidelity and file size

At what point is the image recognizable?

File size grows linearly

% variance doesn't seem to be a useful measure of visual fidelity.

Animation produced with the R package, animation

#### Netflix prize: Sparse SVD



- Netflix, Amazon & others use *recommender systems* to advertise, based on your picks and ratings of others
- In 2006, Neflix offered a \$1M prize for any algorithm that could beat their own RMSE by 10% or more
- The analysis is based on a huge table of ratings by users (~500,000) by items (~20,000), most of which (~99.9%) is missing.
- In 2009, the grand prize was awarded to the Bellcore team ("Pragmatic Chaos")
- Their algorithm used a variation of the SVD specially designed for huge, thin datasets (+ other methods)
- Today, this problem is a hotbed of research activity in statistics and computer science

45

47

## **Biplot** approximation

 A rank 2 (or 3) approximation to the data matrix can be obtained from the first 2 (or 3) singular values λ<sub>i</sub> and the corresponding u<sub>i</sub>, v<sub>i</sub>. e.g.,

$$\mathbf{X} \approx \mathbf{\hat{X}} = \lambda_1 u_1 v_1^T + \lambda_2 u_2 v_2^T$$

The goodness of fit (variance acct'd for) is

$$(\lambda_1^2 + \lambda_2^2) / \sum_{i=1}^q \lambda_i$$

 Biplot is obtained by representing the matrix X as the product

$$\hat{\mathbf{X}} = (\mathbf{U}\mathbf{\Lambda}^{1/2})(\mathbf{\Lambda}^{1/2}\mathbf{V}^T) = \mathbf{A}\mathbf{E}$$

## Biplot: US crime rates



Dim1: ~ Overall crime rate

Dim2: Property vs. personal

Note: clusters of southern, New England, western states

This 2D biplot shows 76.5% of total variance, the projection of the data into the 2D space acct'ing for maximum variance.

It gives a low-D summary of 7 variables and 50 observations.

## **Biplot: US crime rates**



PC1

A 3<sup>rd</sup> dimension accounts for an additional 10.3%, giving 86.8% explained variance.

PC2

# **Application: Collinearity**

In multiple regression, collinearity among Xs is associated with the **smallest** eigenvalue(s).



A collinearity biplot shows the two *smallest* dimensions in X space, as well as "outliers" (high-leverage points)





Image from: http://www.ggobi.org/book/2007-infovis/05-clustering.pdf

## Exploring high-D data

Dynamic, interactive graphics from ggobi, showing PCbased rotations of 6-D data

Interesting features, clusters are selected ("brushed"), painted (assigned colors)

Each view can show PC or X vectors as frame of reference

Repeat:

{spin, brush, paint} until {no more clusters}

### Summary: Why you should like eigenstuff

- Eigenvalues & vectors distill the "juice" in multivariate data
  - $\lambda_i$  = size of each *orthogonal* dimension
  - $#(\lambda_i \neq 0) = rank = # non-zero dimensions$
  - v<sub>i</sub> = weights of variables in each linear combination
  - small, non-zero  $\lambda_i$  relate to outliers & collinearity
- Many multivariate methods → eigen problems
  - PCA, FA
  - MANOVA, MMreg, discriminant analysis
- Important in visualization of high-D data