Topic Notes

Multivariate Visual Representations 2

CS 7450 - Information Visualization
Sep. 15, 2011
John Stasko

Recap

- We examined a number of techniques for projecting >2 variables (modest number of dimensions) down onto the 2D plane
  - Scatterplot matrix
  - Table lens
  - Parallel coordinates
  - etc.

Varieties of Techniques

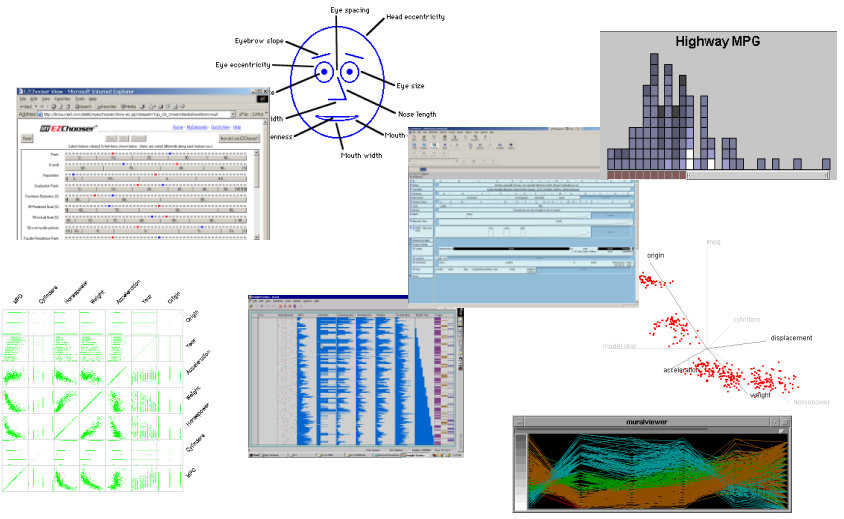

Dust & Magnet

- Altogether different metaphor
- Data cases represented as small bits of iron dust
- Different attributes given physical manifestation as magnets
- Interact with objects to explore data

Yi, Melton, Stasko & Jacko
Information Visualization '05

Interface

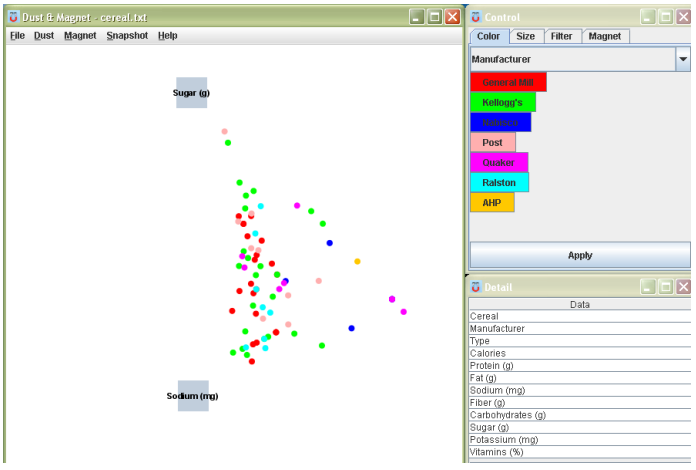

Interaction

- Iron bits (data) are drawn toward magents (attributes) proportional to that data element's value in that attribute
  - Higher values attracted more strongly
- All magnets present on display affect position of all dust
- Individual power of magnets can be changed
- Dust's color and size can connected to attributes as well

Interaction

- Moving a magnet makes all the dust move
  - Also command for shaking dust
- Different strategies for how to position magnets in order to explore the data

See It Live

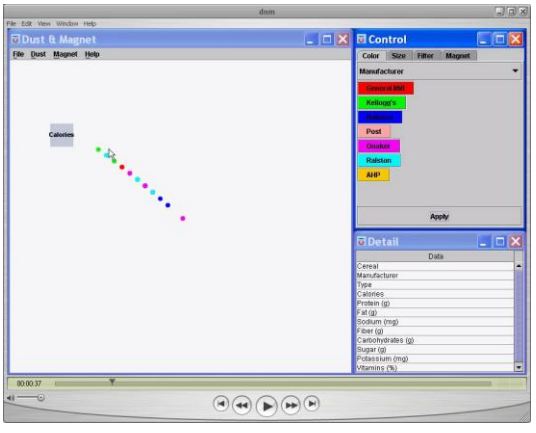

ftp://ftp.cc.gatech.edu/pub/people/stasko/movies/dnm.mov

Video & Demo

Set Operations

- Different type of problem
  - Large set of items, each can be in one or more sets
  - How do we visually represent the set membership?

Standard Technique

Venn Diagram

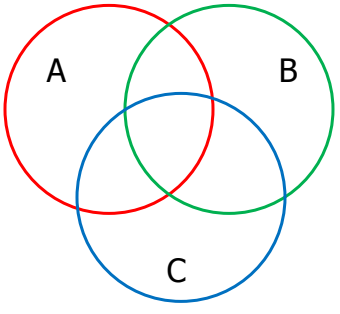

Contains all possible zones of overlap

Alternately

Euler Diagram

Does not necessarily show all possible overlap zones

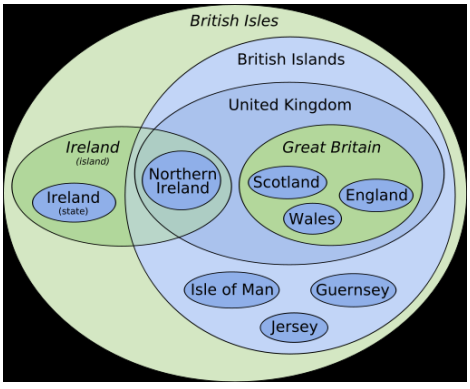

http://en.wikipedia.org/wiki/File:British_Isles_Euler_diagram_15.svg

But what's the problem?

Bubble Sets

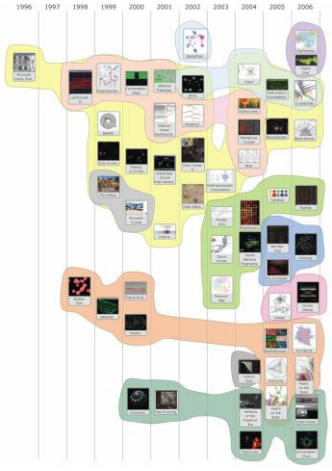

Video

Collins et al
TVCG (InfoVis) '09

ComED & DupED

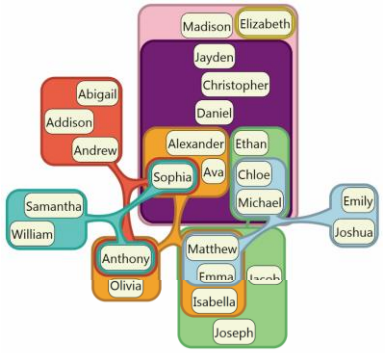

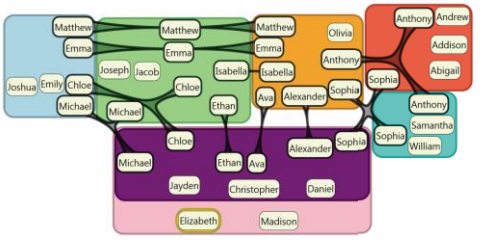

Video

Riche & Dwyer
TVCG (InfoVis) '10

Step Back

- Most of the techniques we've examined work for a modest number of data cases or variables
  - What happens when you have lots and lots of data cases and/or variables?

Many Cases

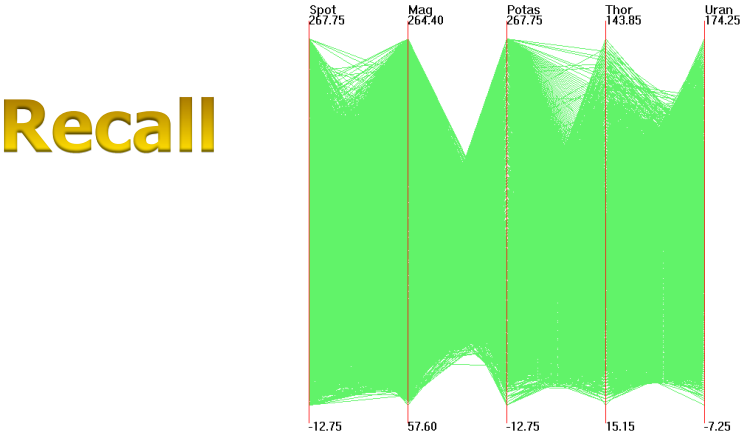

Out5d dataset (5 dimensions, 16384 data items)

(courtesy of J. Yang)

Many Variables

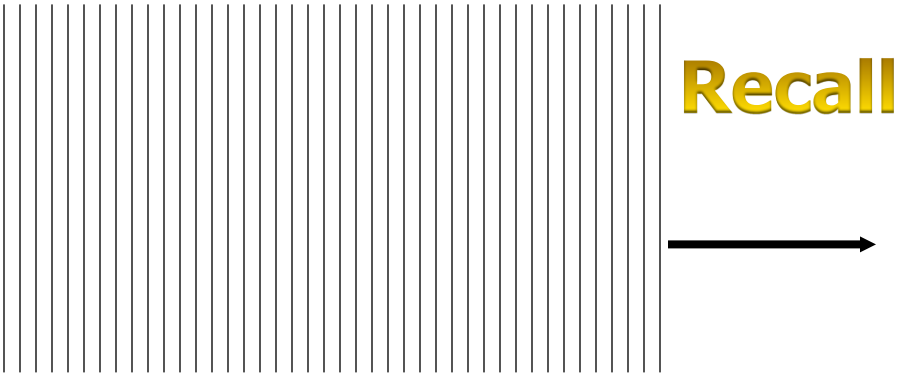

Strategies

- How are we going to deal with such big datasets with so many variables per case?
- Ideas?

General Notion

- Data that is similar in most dimensions ought to be drawn together
  - Cluster at high dimensions
- Need to project the data down into the plane and give it some ultra-simplified representation
- Or perhaps only look at certain aspects of the data at any one time

Mathematical Assistance 1

- There exist many techniques for clustering high-dimensional data with respect to all those dimensions
  - Affinity propagation
  - k-means
  - Expectation maximization
  - Hierarchical clustering

Mathematical Assistance 2

- There exist many techniques for projecting n-dimensions down to 2-D (dimensionality reduction)
  - Multi-dimensional scaling (MDS)
  - Principal component analysis
  - Linear discriminant analysis
  - Factor analysis

Comput Sci & Eng courses
Visual Analytics, Prof. Lebanon

Data mining
Knowledge discovery

Other Techniques

- Other techniques exist to manage scale
  - Sampling – We only include every so many data cases or variables
  - Aggregation – We combine many data cases or variables
  - Interaction (later)
    - Employ user interaction rather than special renderings to help manage scale

Our Focus

- Visual techniques
- Many are simply graphic transformations from N-D down to 2-D

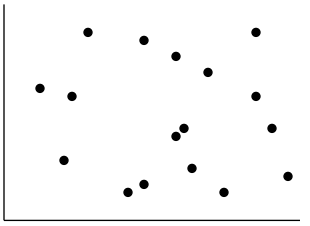

Use?

- What kinds of questions/tasks would you want such techniques to address?
  - Clusters of similar data cases
  - Useless dimensions
  - Dimensions similar to each other
  - Outlier data cases
  - …
- Think back to our "cognitive tasks" discussion

Now

- We'll examine a number of other visual techniques intended for larger, higher-dimensional data sets

Can We Make a Taxonomy?

- D. Keim proposes a taxonomy of techniques
  - Standard 2D/3D display
    - Bar charts, scatterplots
  - Geometrically transformed display
    - Parallel coordinates
  - Iconic display
    - Needle icons, Chernoff faces
  - Dense pixel display
    - What we're about to see…
  - Stacked display
    - Treemaps, dimensional stacking

TVCG '02

Minimum Possible?

- We have data cases with variables
- What's the smallest representation we can use?
  - How?

Dense Pixel Display

- Represent data case or a variable as a pixel
- Million or more per display
- Seems to rely on use of color
- Can pack lots in

- Challenge: What's the layout?

One Representation

- Grouping arrangement
- One pixel per variable
- Each data case has its own small rectangular icon
- Plot out variables for data point in that icon using a grid or spiral layout

Uses color scale

Illustration

schematic representation of 6-dim. data

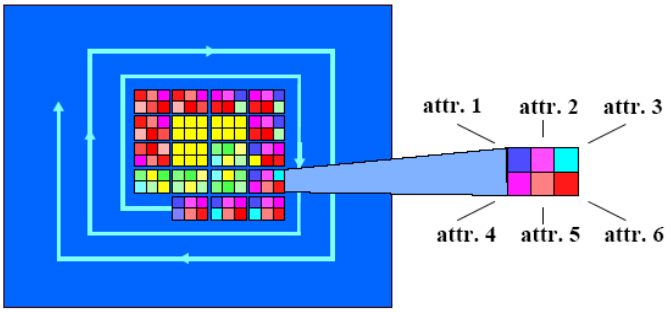

Levkowitz
Vis '91

Spiral Technique

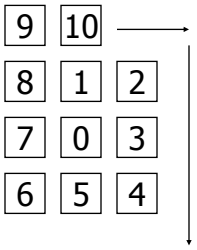

Dimensions

Related Idea

- Pixel Bar Chart
- Overload typical bar chart with more information about individual elements

Keim et al
Information Visualization `02

Idea 1

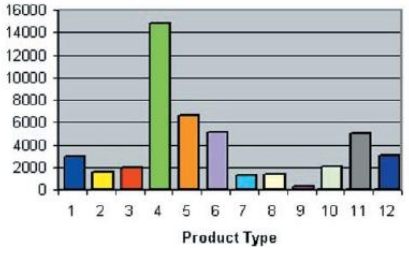

Height encodes quantity

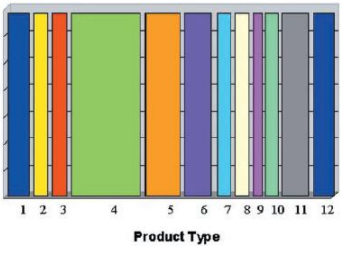

Width encodes quantity

Idea 2

- Make each pixel within a bar correspond to a data point in that group represented by the bar
  - Can do millions that way
- Color the pixel to represent the value of one of the data point's variables

Idea 3

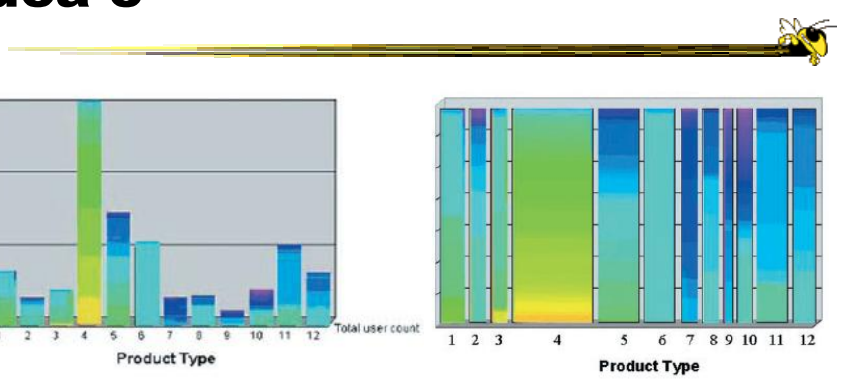

Each pixel is a customer
Color encodes amount spent by that person
High-bright, Low-dark
Ordered by that color attribute too
Right one shows more customers

Idea 4

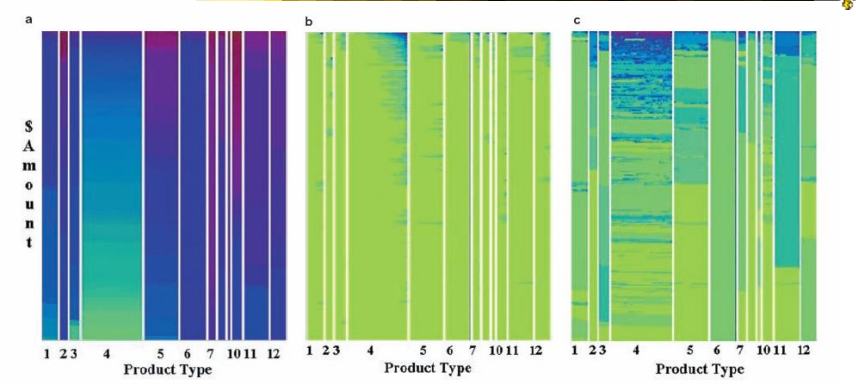

Product type is x-axis divider
Customers ordered by
- y-axis: dollar amount
- x-axis: number of visits

Color is (a) dollar amount spent, (b) number of visits, (c) sales quantity

Idea 5

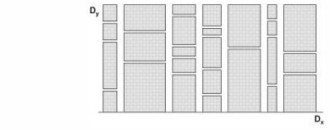

Figure 7 Dividing attributes on x- and y-axis (e.g., D_x = *Product Type*, D_y = *Region*).

Figure 8 Ordering attributes on x- and y-axis (e.g., O_x = *Dollar Amount*, O_y = *Quantity*).

Can divide on two different attributes on x and y

Order items on both x and y

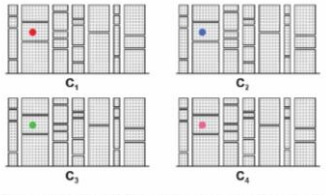

Figure 9 Multiple coloring attributes (e.g., C_1 = *dollar amount*, C_2 = *no. of visits*, C_3 = *quantity*, C_4 = *region*).

Color maps to some attribute
(Same item always at same x,y position)

Idea 6

Mapping specified by 5 tuple <D_x, D_y, O_x, O_y, C>

D_x – Attribute partitions x axis
D_y – Attribute partitions y axis
O_x – Attribute specifies x ordering
O_y – Attribute specifies y ordering
C – Attribute specifies color mapping

Example Application

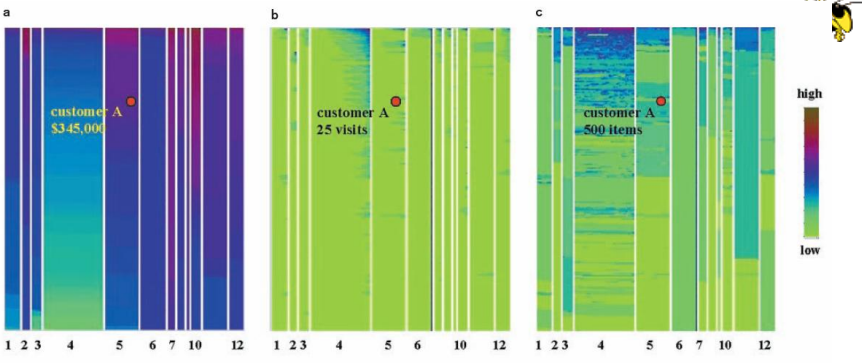

Figure 13 Multi-pixel bar chart for mining 405,000 sales transaction records. (D_x = *Product Type*, D_y = ⊥, O_x = *no. of visits*, O_y = *dollar amount*, C). (a) Color: dollar amount. (b) Color: no. of visits. (c) Color: quantity.

1. Product type 7 and product type 10 have the top dollar amount customers (dark colors of bar 7 and 10 in Figure 13a)
2. The dollar amount spent and the number of visits are clearly correlated, especially for product type 4 (linear increase of dark colors at the top of bar 4 in Figure 13b)
3. Product types 4 and 11 have the highest quantities sold (dark colors of bar 4 and 11 in Figure 13c)
4. Clicking on pixel A shows details for that customer

Thoughts?

- Do you think that would be a helpful exploratory tool?

High Dimensions

- Those techniques could show lots of data, but not so many dimensions at once
  - Have to pick and choose

Another Idea

- Use the dense pixel display for showing data and dimensions, but then project into 2D plane to encode more information
- VaR – Value and relation display

Yang et al
InfoVis '04

Algorithm

- Find a correlation function for comparing dimensions
- Calculate distances between dimensions (similarities)
- Make each dimension into a dense pixel glyph
- Assign position for each glyph in 2D plane using multi-dimensional scaling

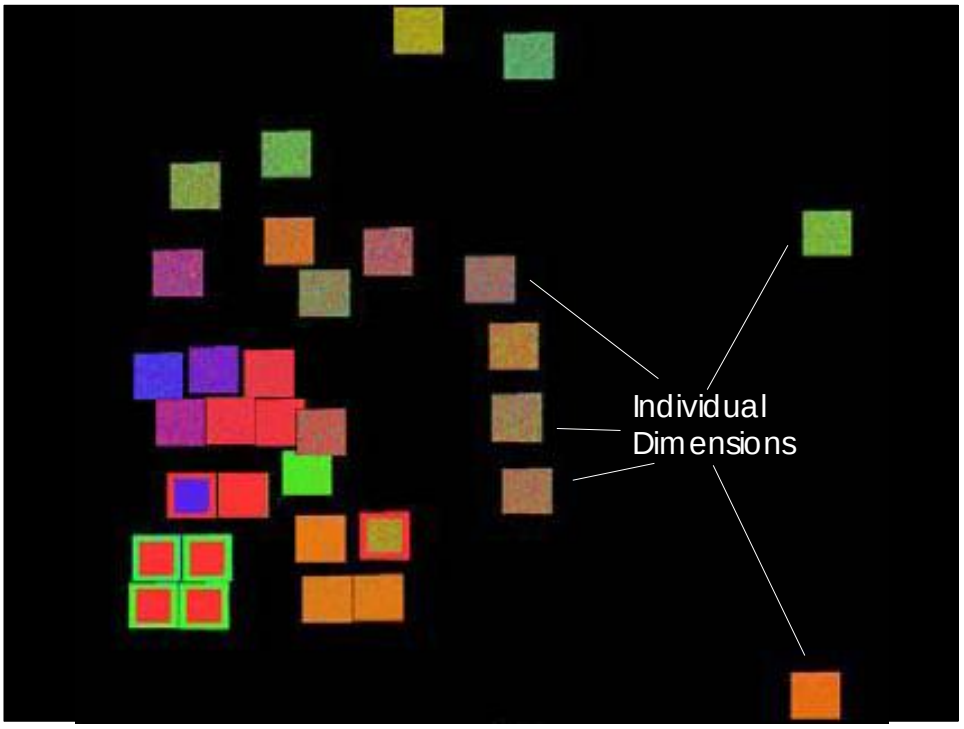

Questions

- What order are the data cases in each dimension-glyph?
  - Maybe there is a predefined order
  - Choose one dimension as "important" then order data cases by their values in that dimension
    - "Important" one may be the one in which many cases are similar

Alternative

- Instead of each glyph being a dimension, it can be a data case

Follow-on Work

- Use alternate positioning strategies other than MDS
- Use Jigsaw map idea (Wattenberg, InfoVis '05) to lay out the dimensions into a grid
  - Removes overlap
  - Limits number that can be plotted

Yang et al
TVCG '07

New Layout

Plot the glyphs into the grid positions

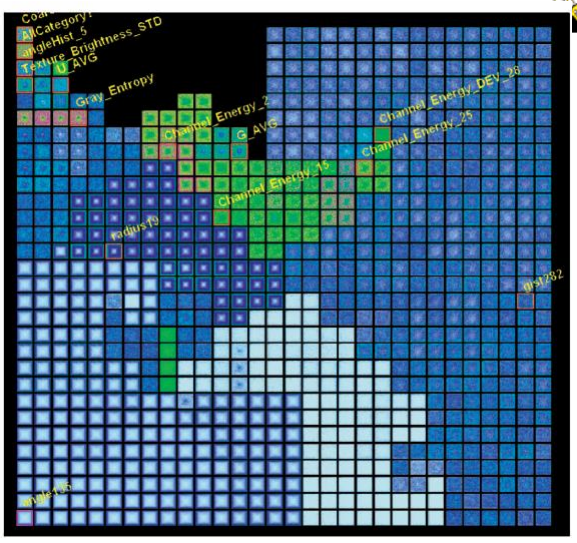

HCE

- Hierarchical Clustering Explorer
- Implements "rank by feature" framework
- Help guide user to choose 1D distributions and 2D scatterplots from various dimensions of a data set
- Combine statistical analysis with user-directed exploration

Seo & Shneiderman
Information Visualization '05

Idea

- Choose a feature detection criterion to rank 1D and 2D projections of a data set
- Use person's perceptual abilities to pick out interesting items from view

HCE UI

Some chosen distributions and scatterplots

Cases in columns, variables in rows

Group similar cases

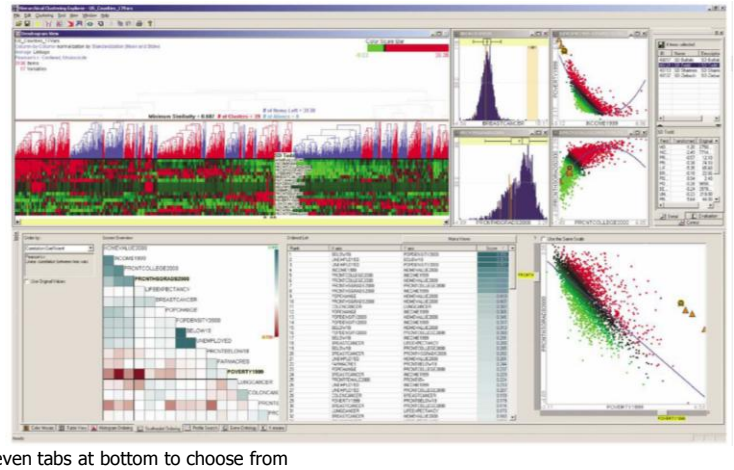

Seven tabs at bottom to choose from

Operation

- When you choose the histogram ordering or scatterplot ordering tabs at the bottom left, these give results based on various statistical measures
- You can then choose some of them to visualize

Demo

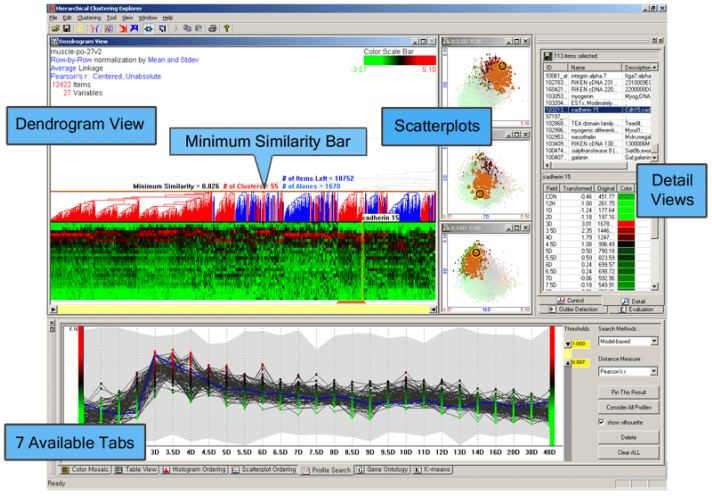

Recap

- We've seen many general techniques for multivariate data these past two days
  - Know strengths and limitations of each
  - Know which ones are good for which circumstances
  - We still haven't explored interaction much

HW 3

- Visualization design
- Due Tuesday
  - Bring two copies
- Questions?

Upcoming

- Tufte's Design Principles
  - Reading:
    Envisioning Information (if you have it)

- Few's Design Guidance
  - Reading
    Now You See It chapters 5-12