CSE512 :: 21 Jan 2014

Multi-Dimensional Vis

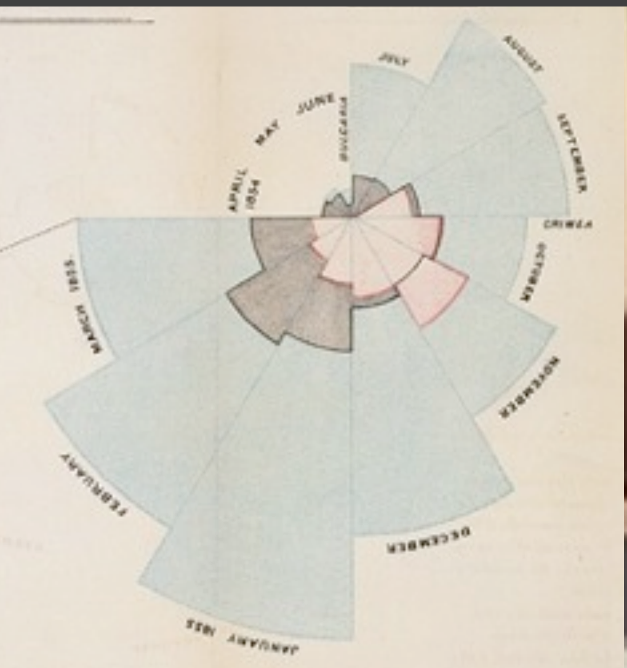

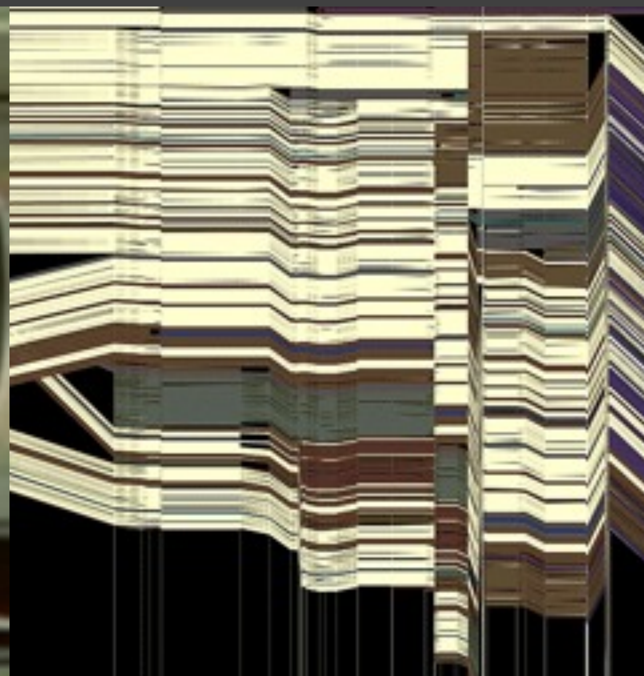

Jeffrey Heer University of Washington

Last Time:
Exploratory Data Analysis

Exposure, the effective laying open of the data to display the unanticipated, is to us a major portion of data analysis. Formal statistics has given almost no guidance to exposure; indeed, it is not clear how the informality and flexibility appropriate to the exploratory character of exposure can be fitted into any of the structures of formal statistics so far proposed.

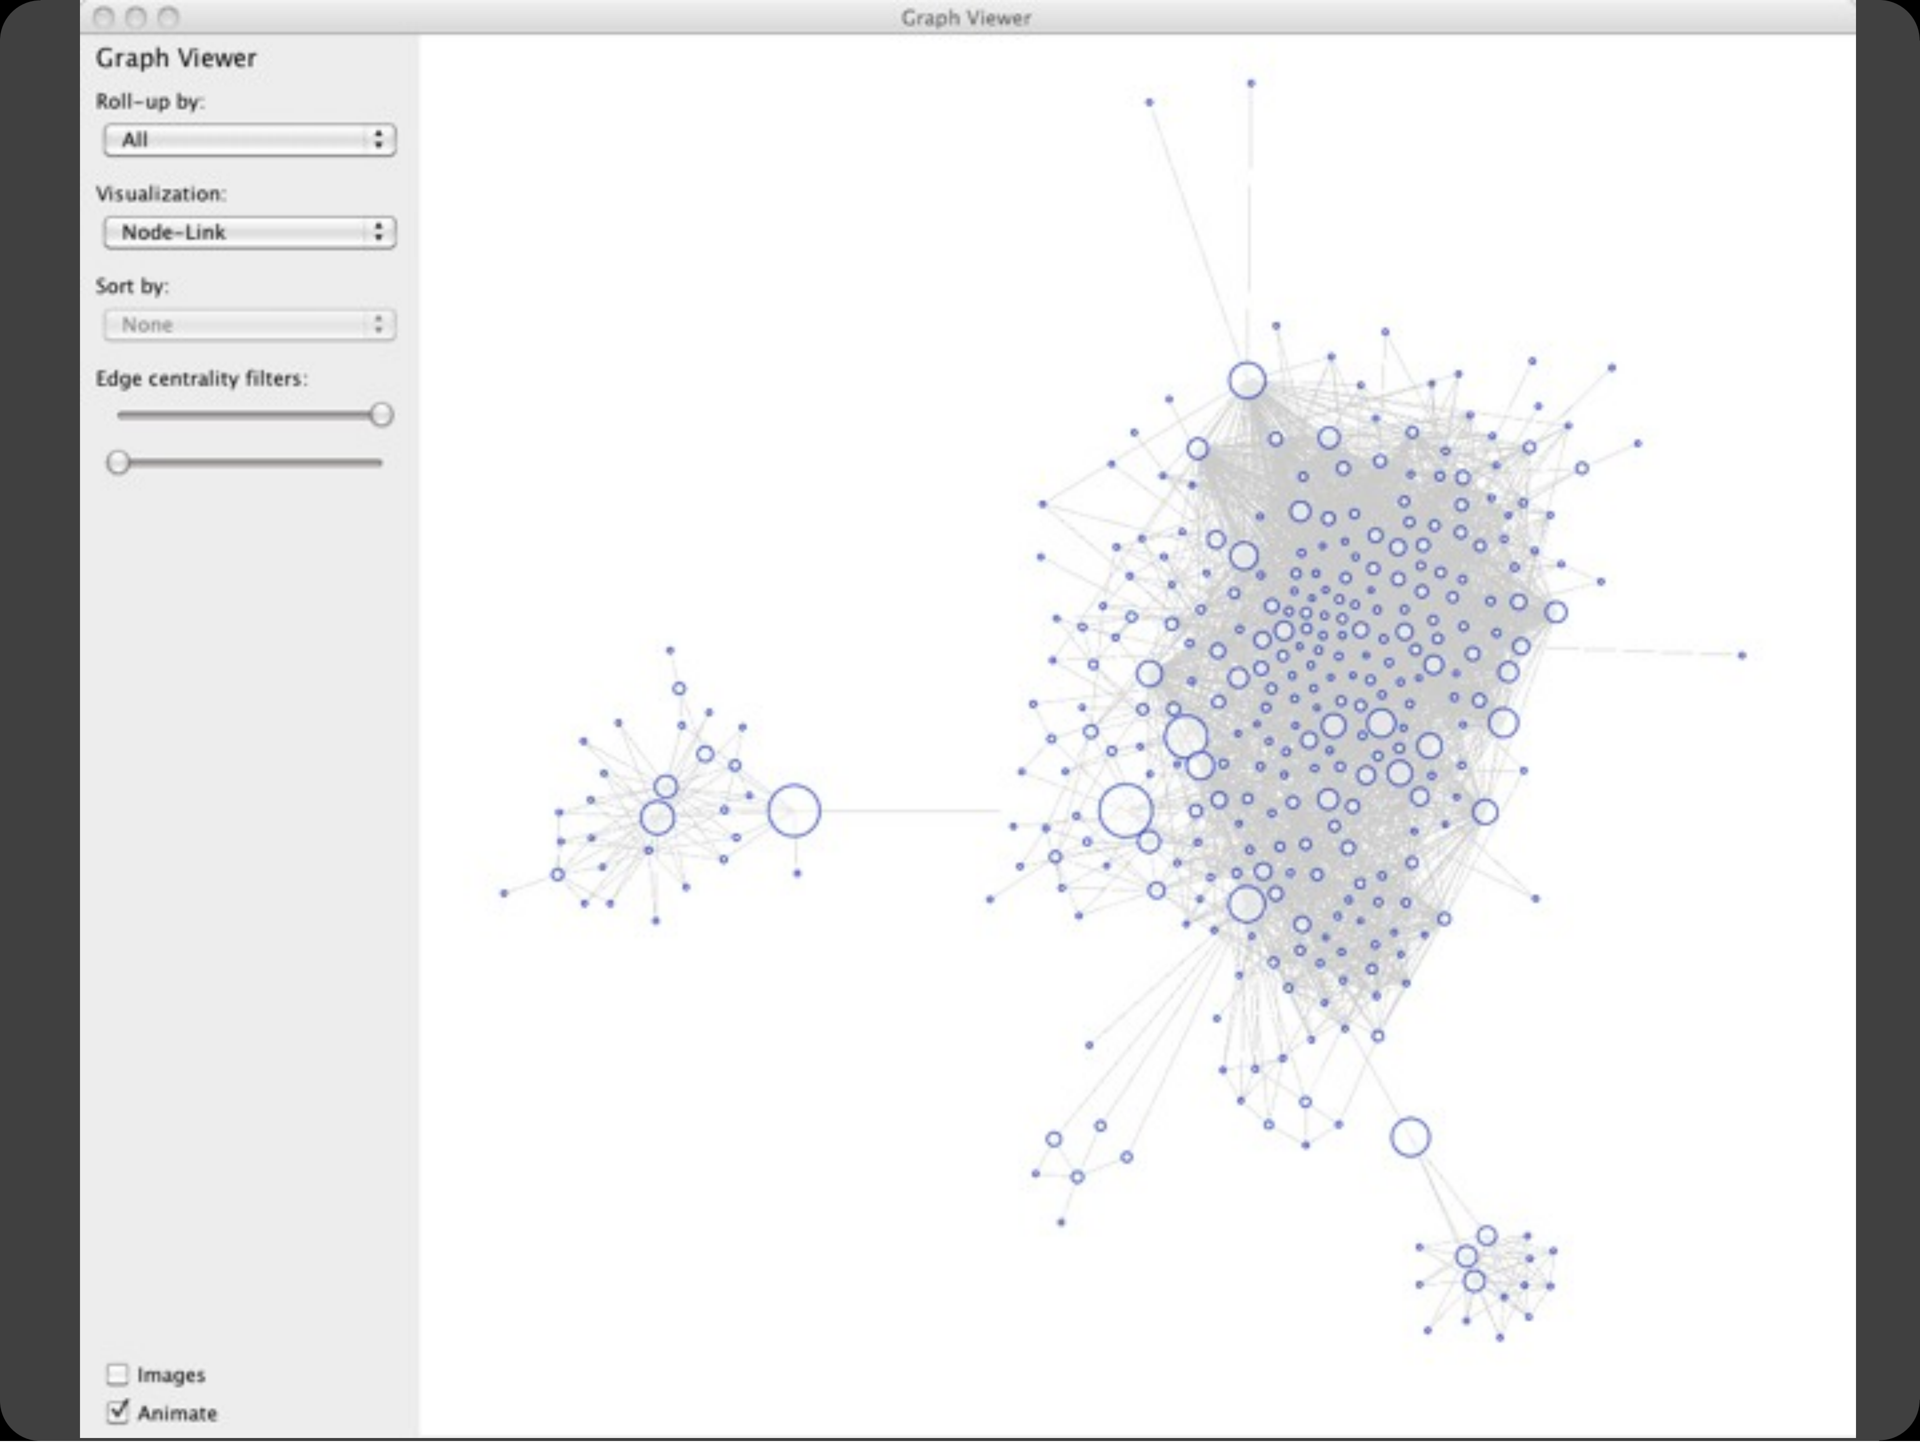
Graph Viewer
Graph Viewer
Roll-up by:
All
Visualization:
Node-Link
Sort by:
None
Edge centrality filters:
Images
Animate

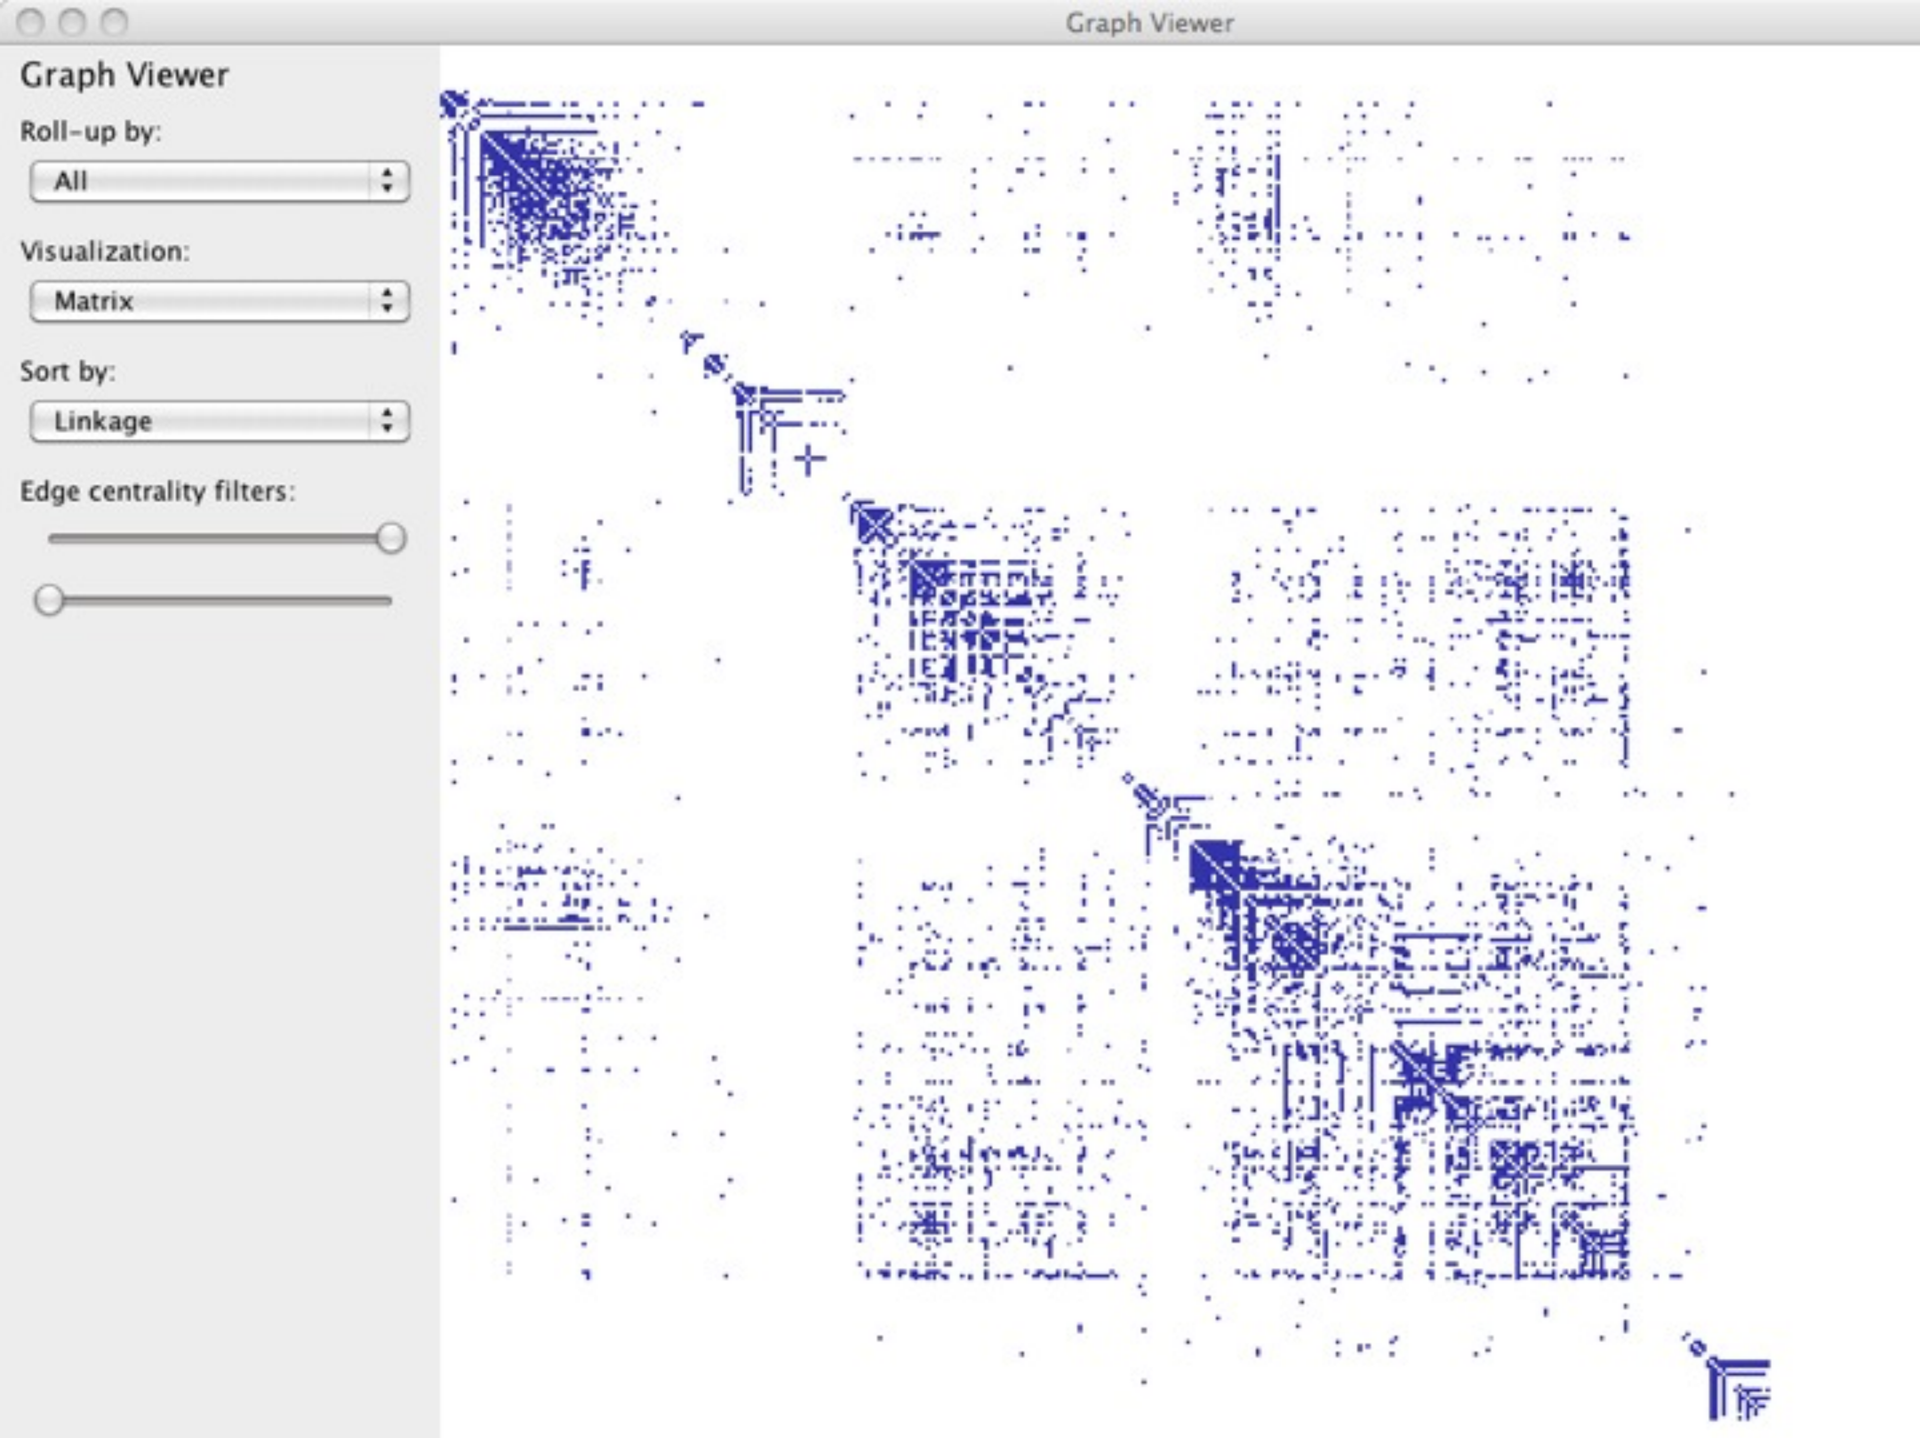
Graph Viewer
Graph Viewer
Roll-up by:
All
Visualization:
Matrix
Sort by:
Linkage
Edge centrality filters:

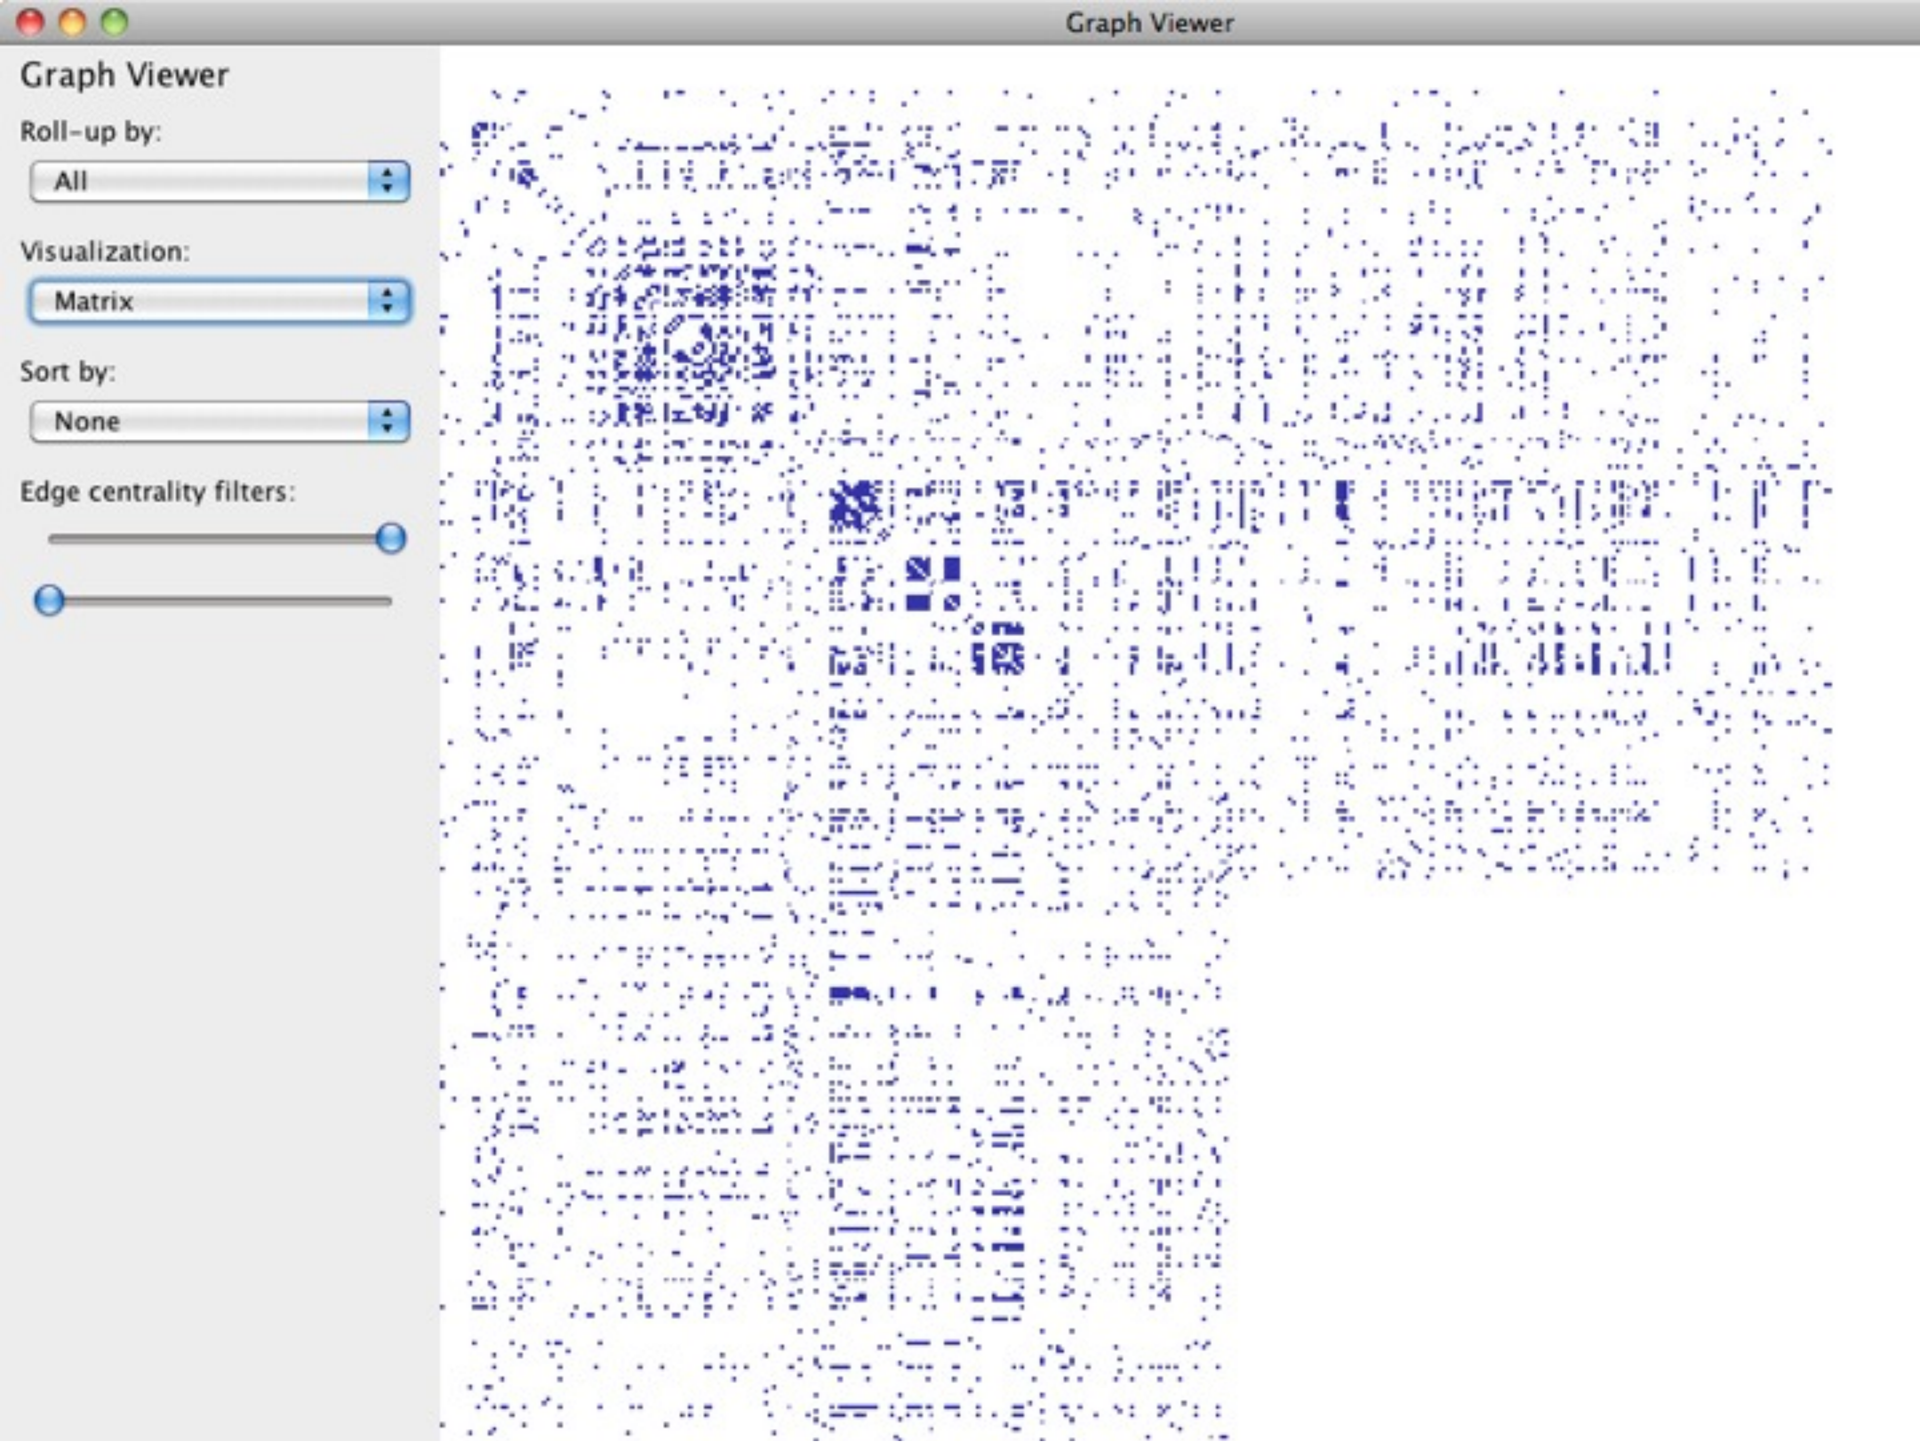
Graph Viewer
Graph Viewer
Roll-up by:
All
Visualization:
Matrix
Sort by:
None
Edge centrality filters:

Antibiotic Effectiveness

Table 1: Burtin's data.

Bacteria	Antibiotic			Gram Staining
	Penicillin	Streptomycin	Neomycin	
Aerobacter *aerogenes*	870	1	1.6	negative
Brucella *abortus*	1	2	0.02	negative
Brucella *anthracis*	0.001	0.01	0.007	positive
Diplococcus *pneumoniae*	0.005	11	10	positive
Escherichia *coli*	100	0.4	0.1	negative
Klebsiella *pneumoniae*	850	1.2	1	negative
Mycobacterium *tuberculosis*	800	5	2	negative
Proteus *vulgaris*	3	0.1	0.1	negative
Pseudomonas *aeruginosa*	850	2	0.4	negative
Salmonella (Eberthella) *typhosa*	1	0.4	0.008	negative
Salmonella *schottmuelleri*	10	0.8	0.09	negative
Staphylococcus *albus*	0.007	0.1	0.001	positive
Staphylococcus *aureus*	0.03	0.03	0.001	positive
Streptococcus *fecalis*	1	1	0.1	positive
Streptococcus *hemolyticus*	0.001	14	10	positive
Streptococcus *viridans*	0.005	10	40	positive

Will Burtin, 1951

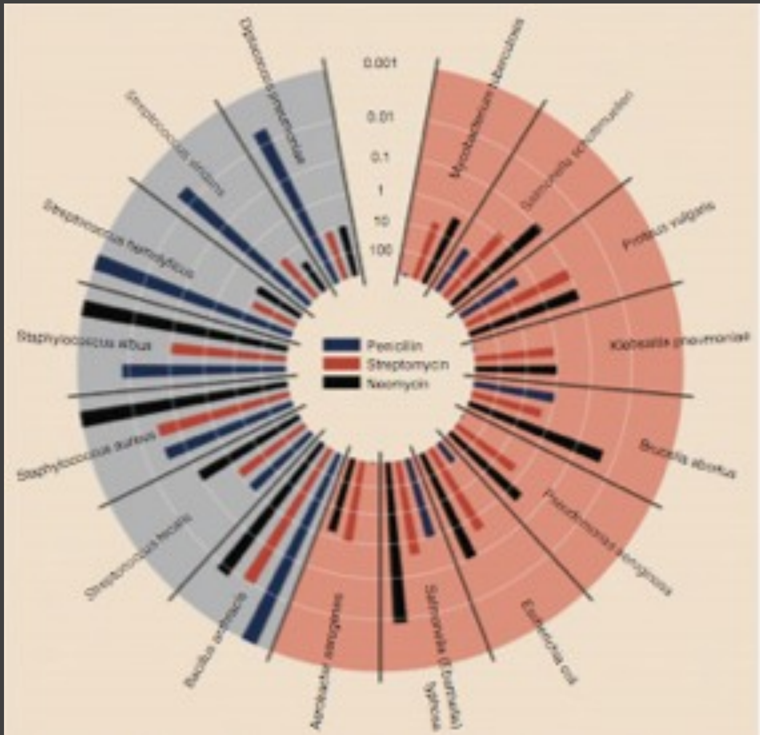

Bacteria	Penicillin	Antibiotic Streptomycin	Neomycin	Gram stain
Aerobacter aerogenes	870	1	1.6	–
Brucella abortus	1	2	0.02	–
Bacillus anthracis	0.001	0.01	0.007	+
Diplococcus pneumoniae	0.005	11	10	+
Escherichia coli	100	0.4	0.1	–
Klebsiella pneumoniae	850	1.2	1	–
Mycobacterium tuberculosis	800	5	2	–
Proteus vulgaris	3	0.1	0.1	–
Pseudomonas aeruginosa	850	2	0.4	–
Salmonella (Eberthella) typhosa	1	0.4	0.008	–
Salmonella schottmuelleri	10	0.8	0.09	–
Staphylococcus albus	0.007	0.1	0.001	+
Staphylococcus aureus	0.03	0.03	0.001	+
Streptococcus fecalis	1	1	0.1	+
Streptococcus hemolyticus	0.001	14	10	+
Streptococcus viridans	0.005	10	40	+

How do the drugs compare?

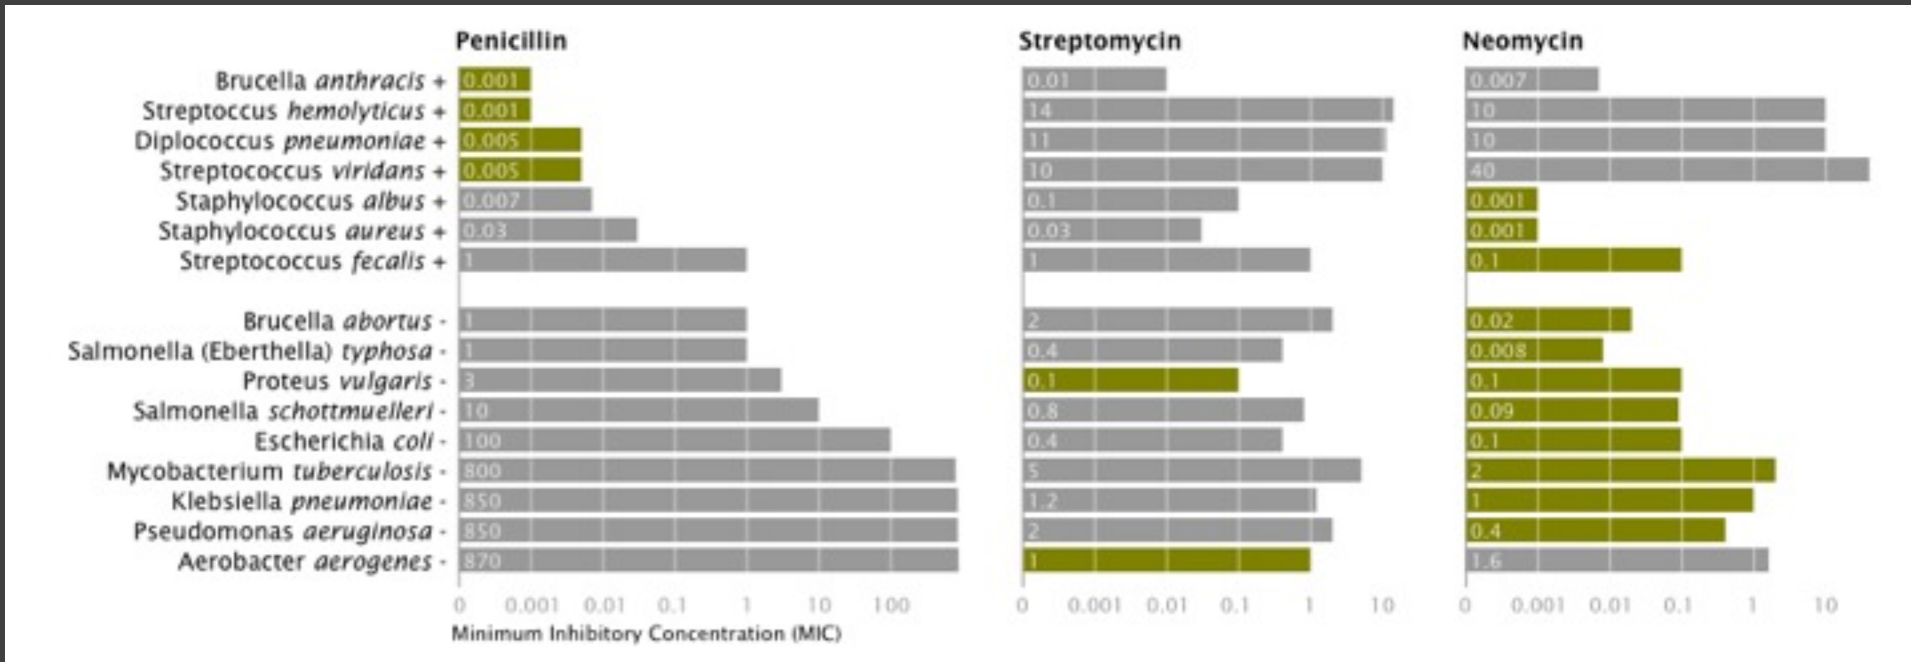

Mike Bostock, CS448B Winter 2009

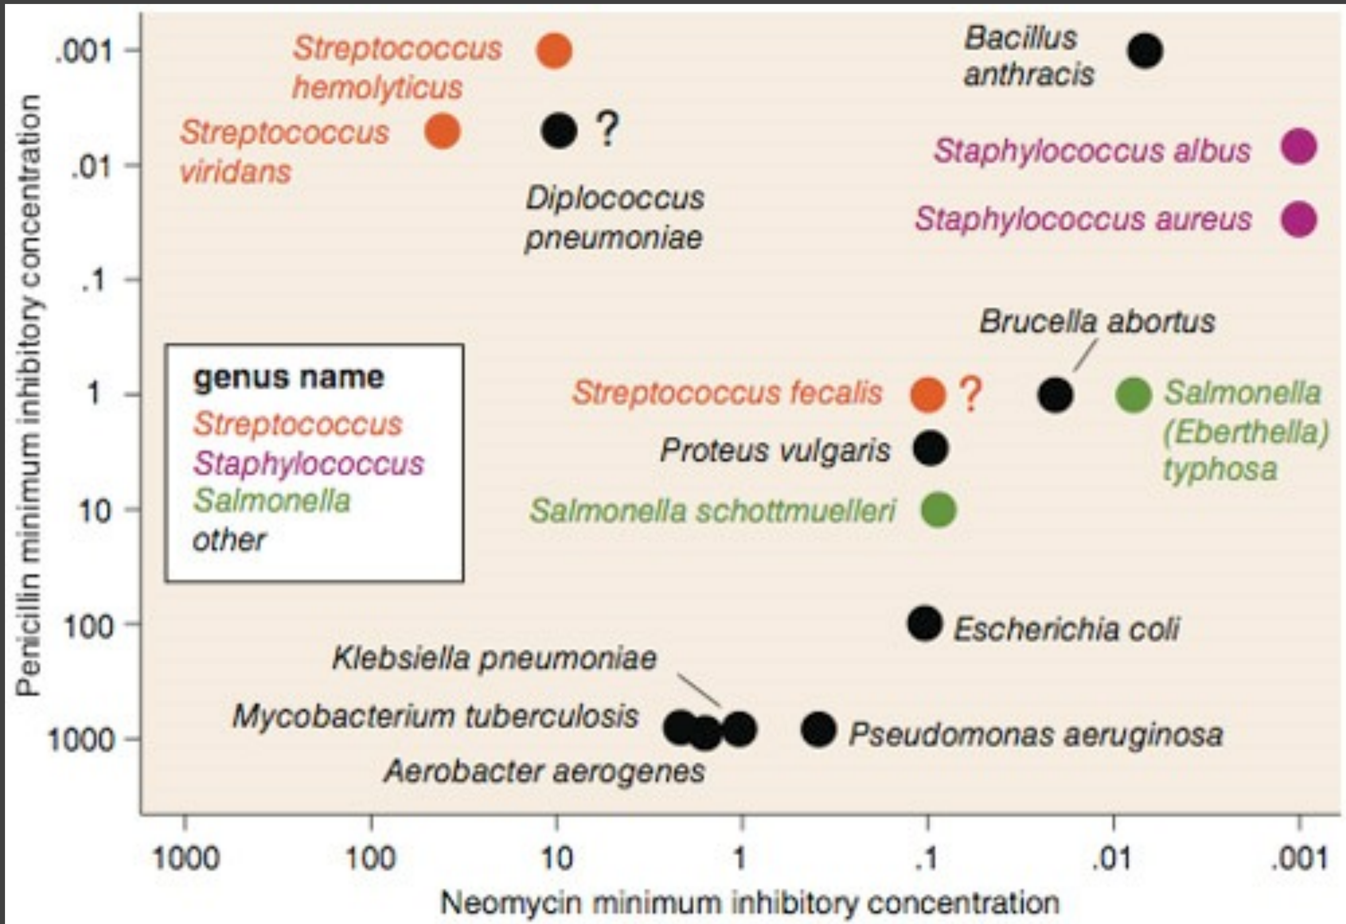

How do the bacteria group w.r.t. resistance?

Do different drugs correlate?

Wainer & Lysen
American Scientist, 2009

Transforming data

How well does the curve fit data?

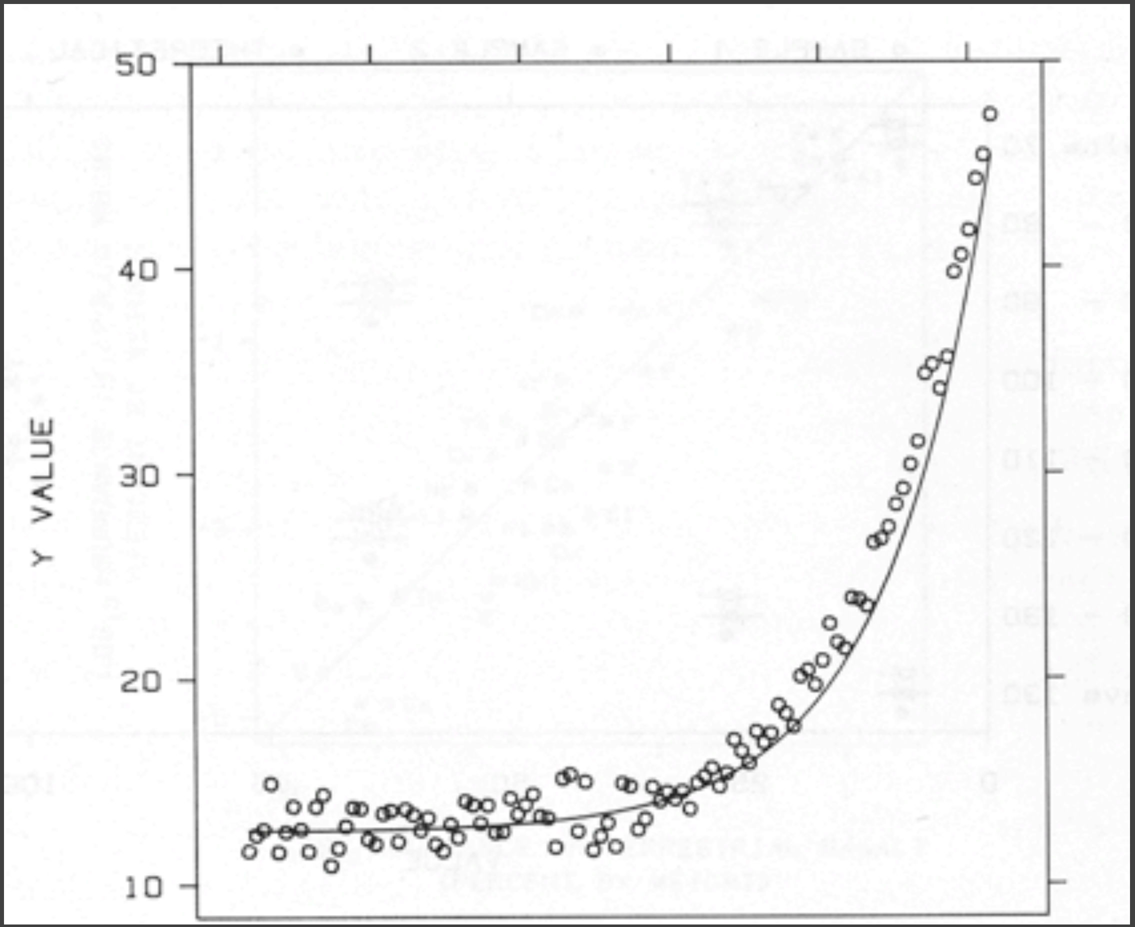

[Cleveland 85]

Plot the Residuals

Plot vertical distance from best fit curve
Residual graph shows accuracy of fit

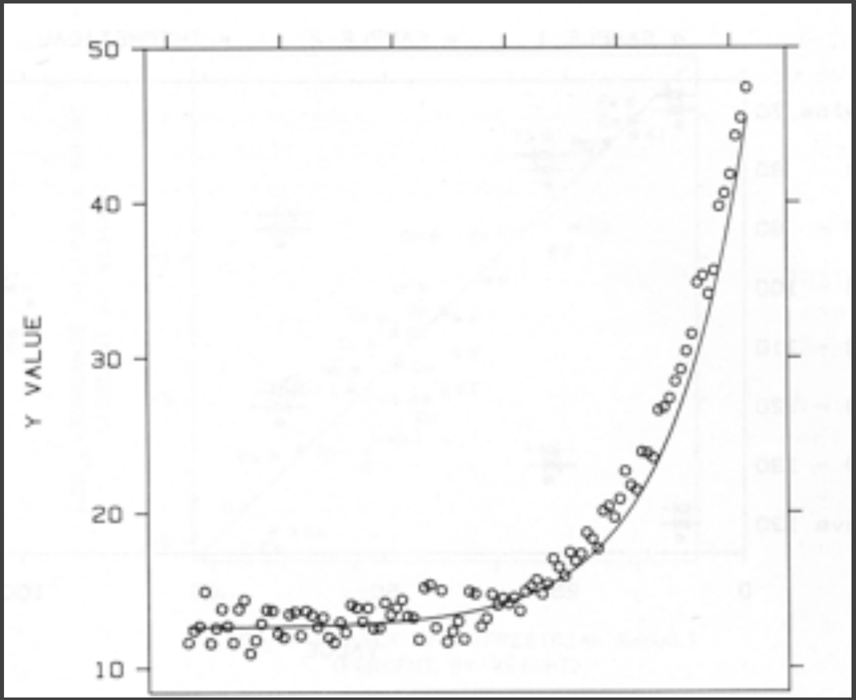

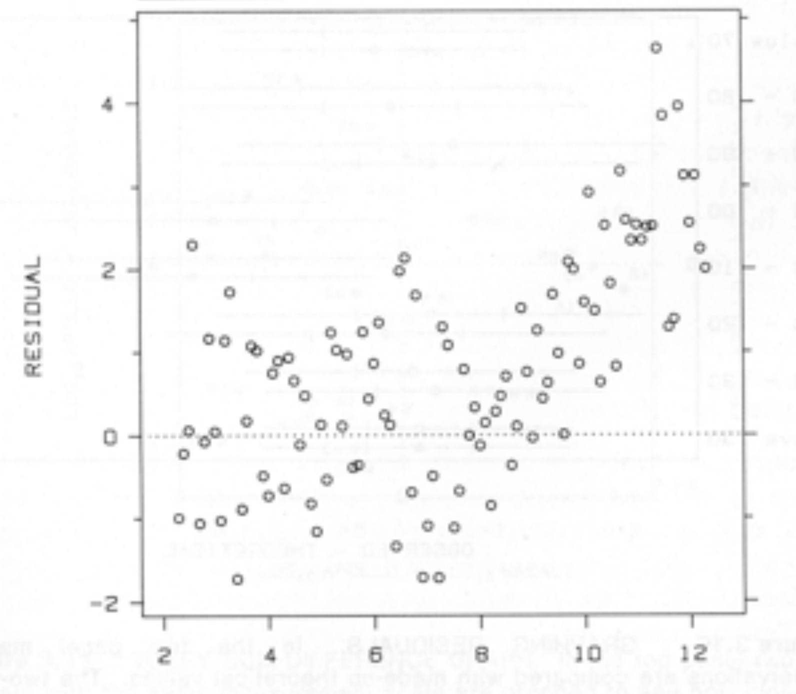

[Cleveland 85]

Multiple Plotting Options

Plot model in data space

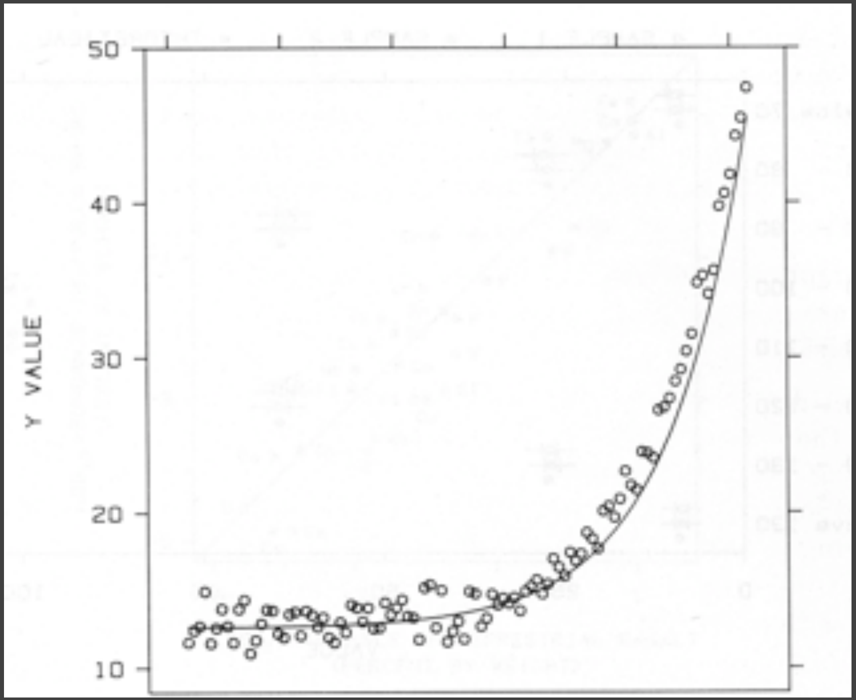

Plot data in model space

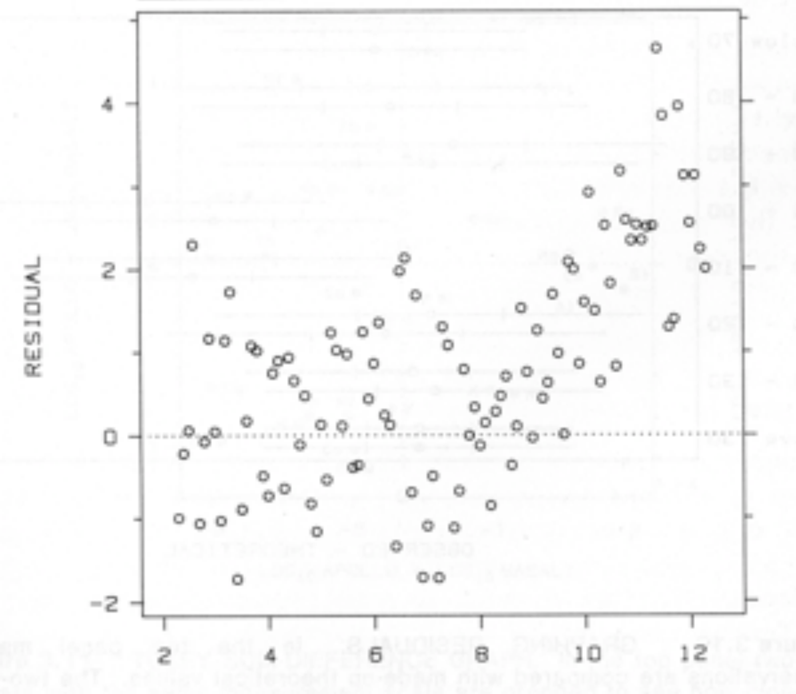

[Cleveland 85]

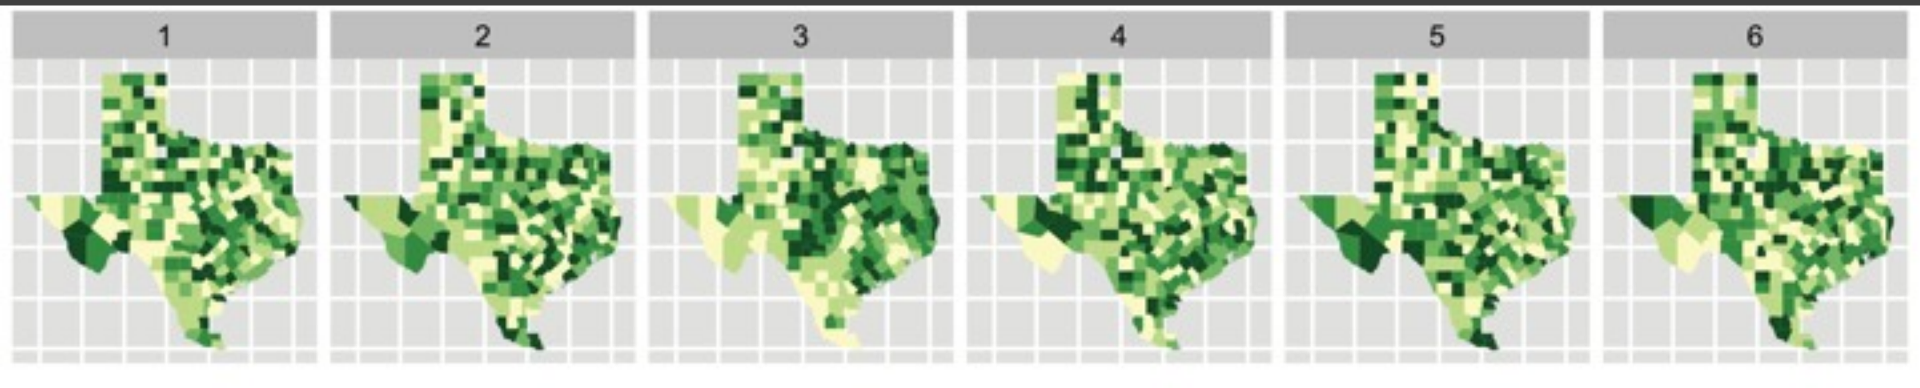

Choropleth maps of cancer deaths in Texas.

One plot shows a real data sets. The others are simulated under the null hypothesis of spatial independence.

Can you spot the real data? If so, you have some evidence of spatial dependence in the data.

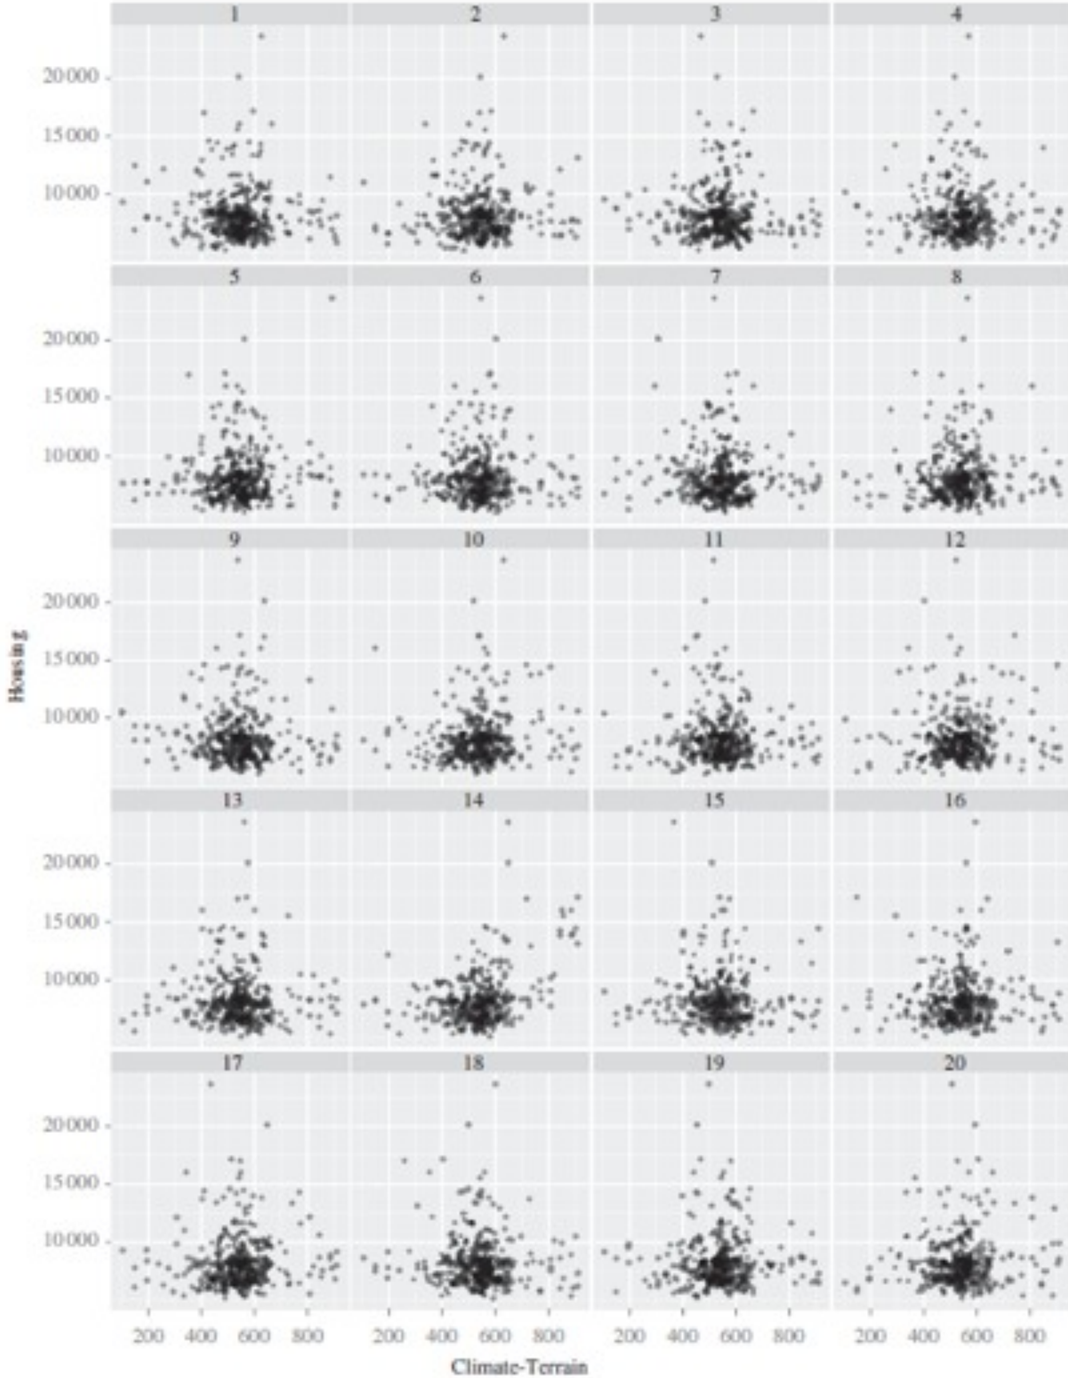
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
Housing
20000
15000
10000
200
400
600
800
Climate-Terrain

Multidimensional Visualization

Visual Encoding Variables

Position
Length
Area
Volume
Value
Texture
Color
Orientation
Shape

~8 dimensions?

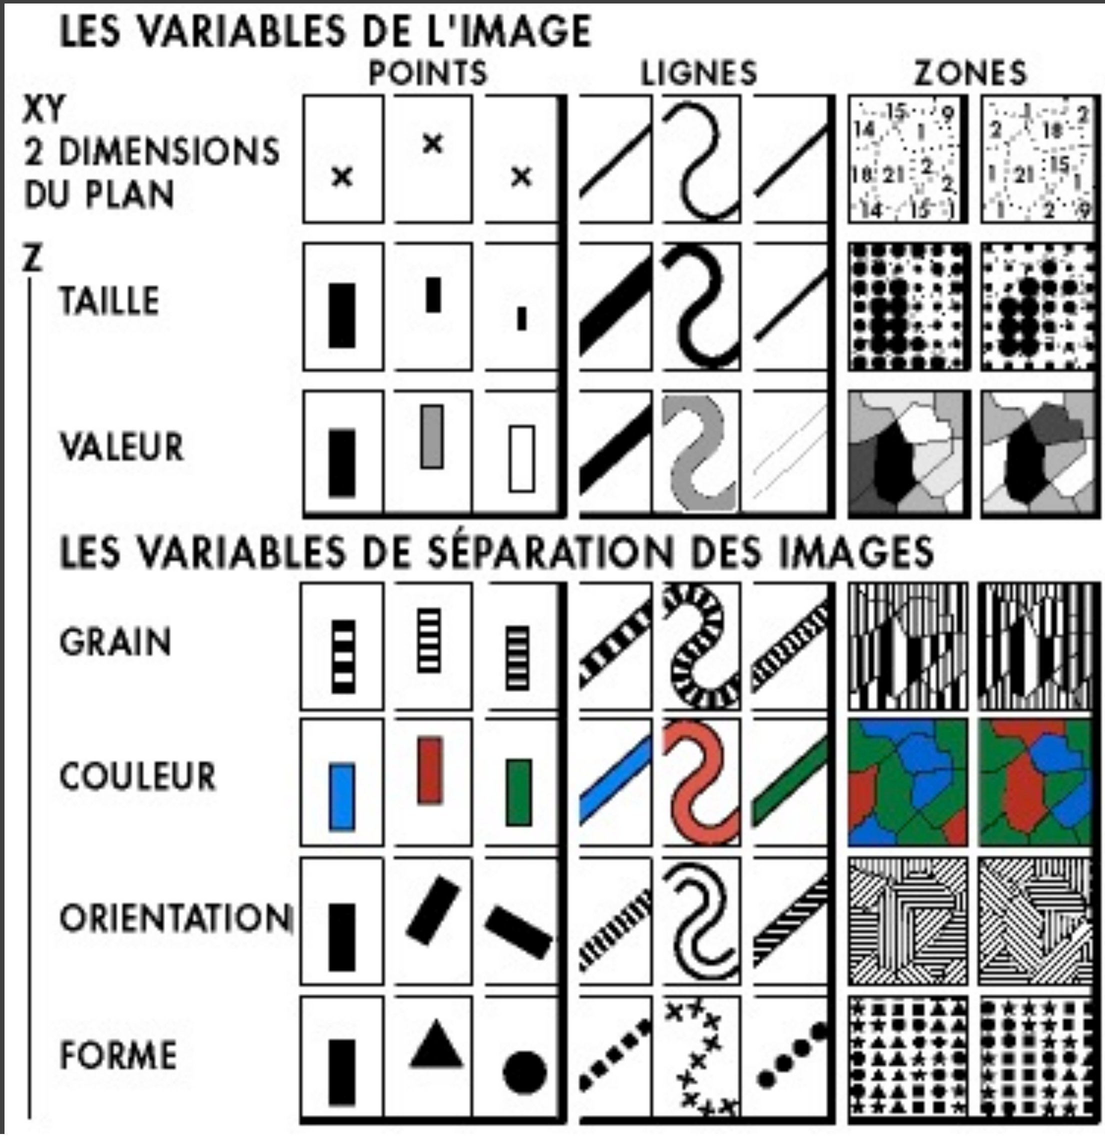

Example: Coffee Sales

Sales figures for a fictional coffee chain:

Sales	Q-Ratio	
Profit	Q-Ratio	
Marketing	Q-Ratio	
Product Type	N	{Coffee, Espresso, Herbal Tea, Tea}
Market	N	{Central, East, South, West}

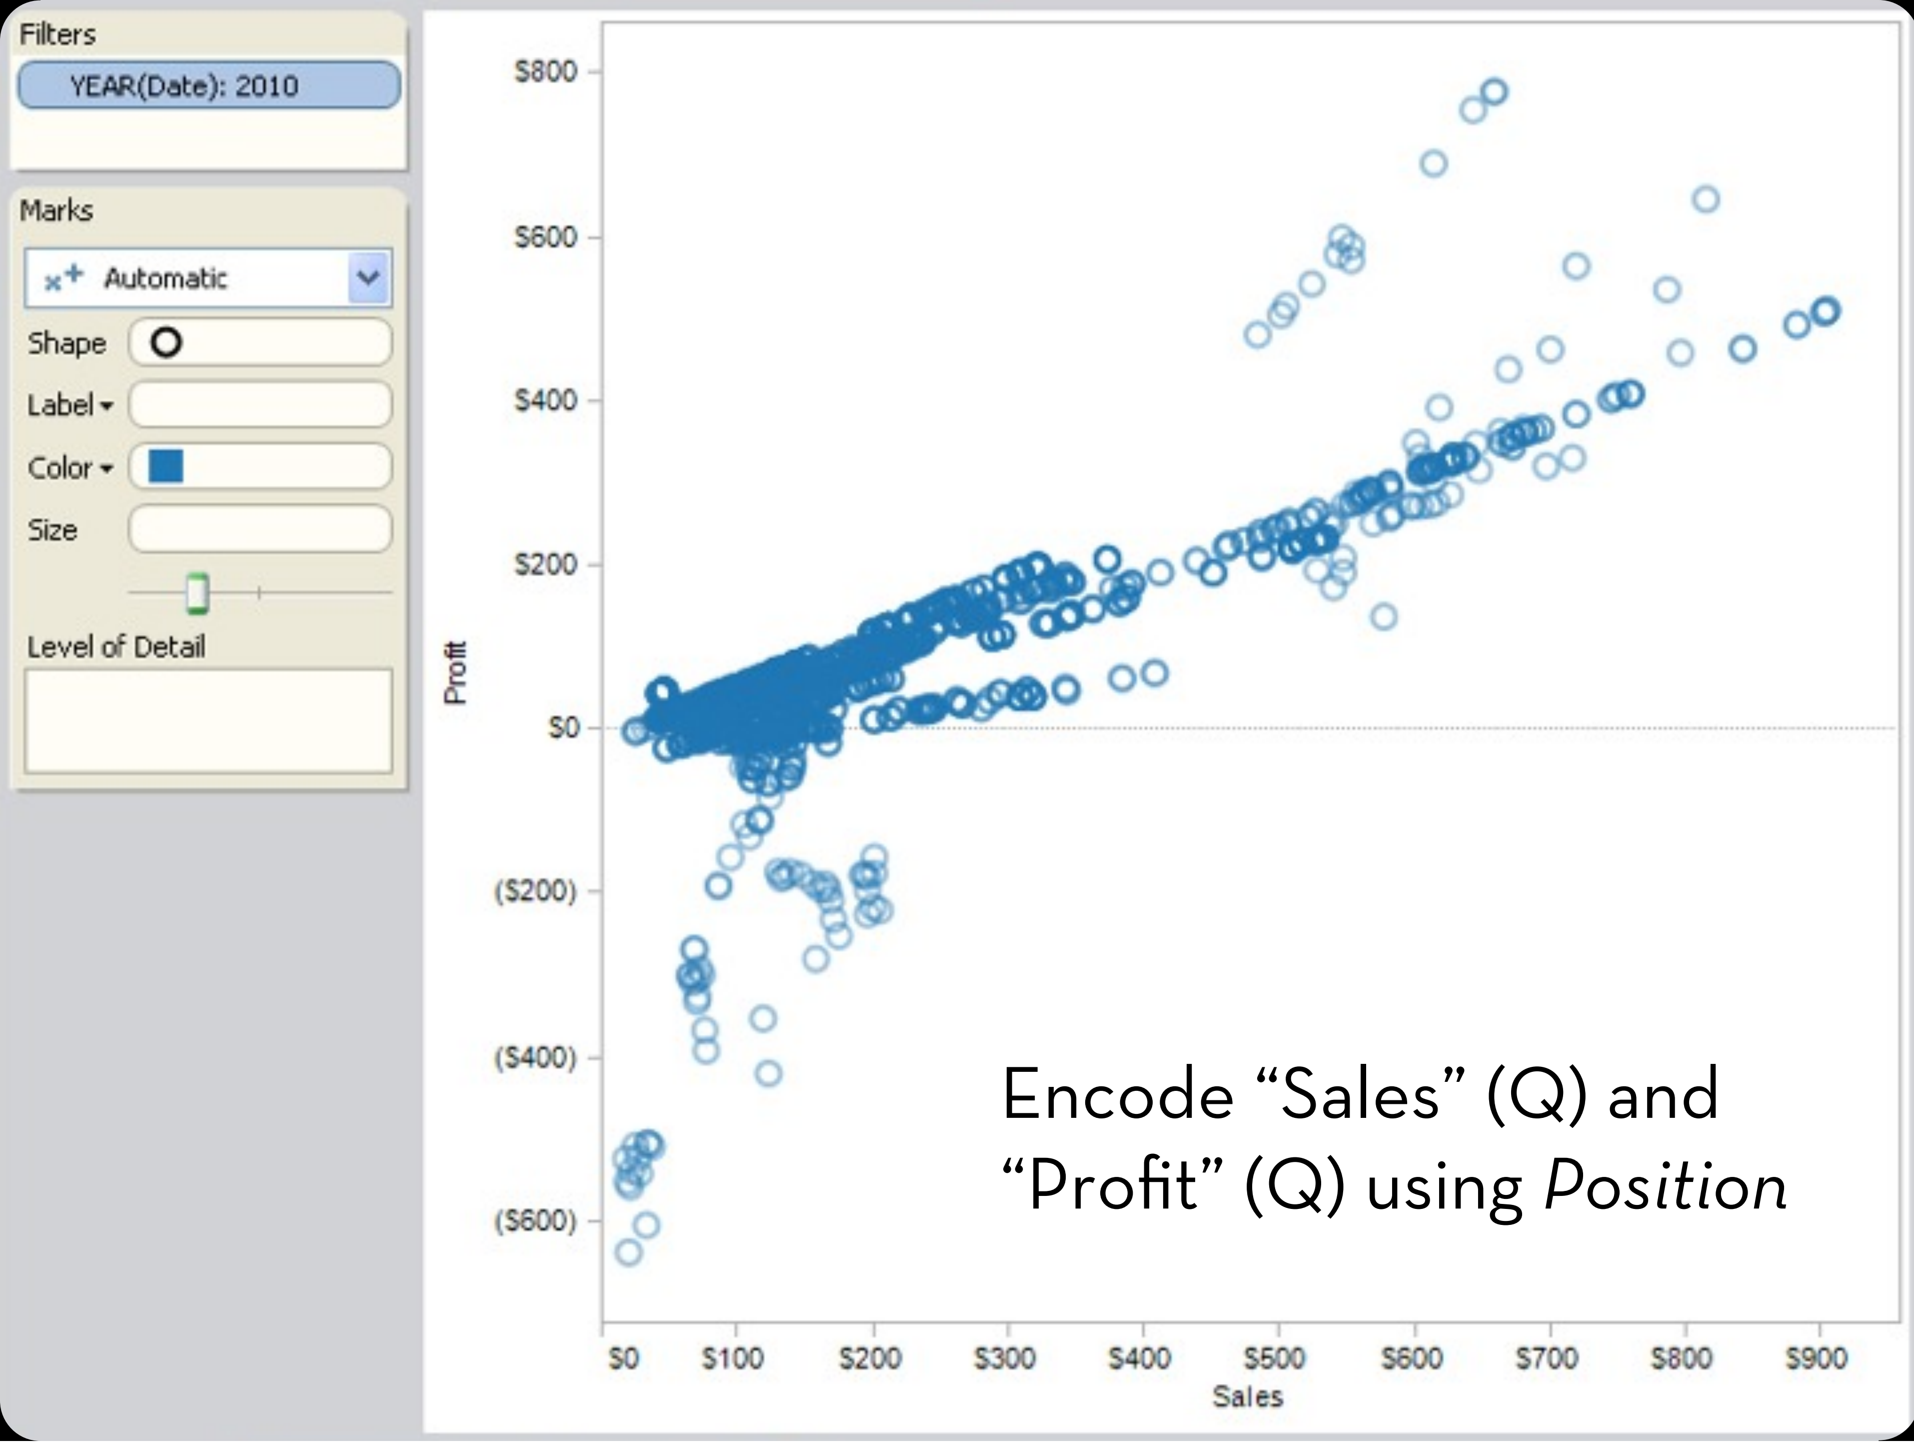

Encode "Sales" (Q) and "Profit" (Q) using *Position*

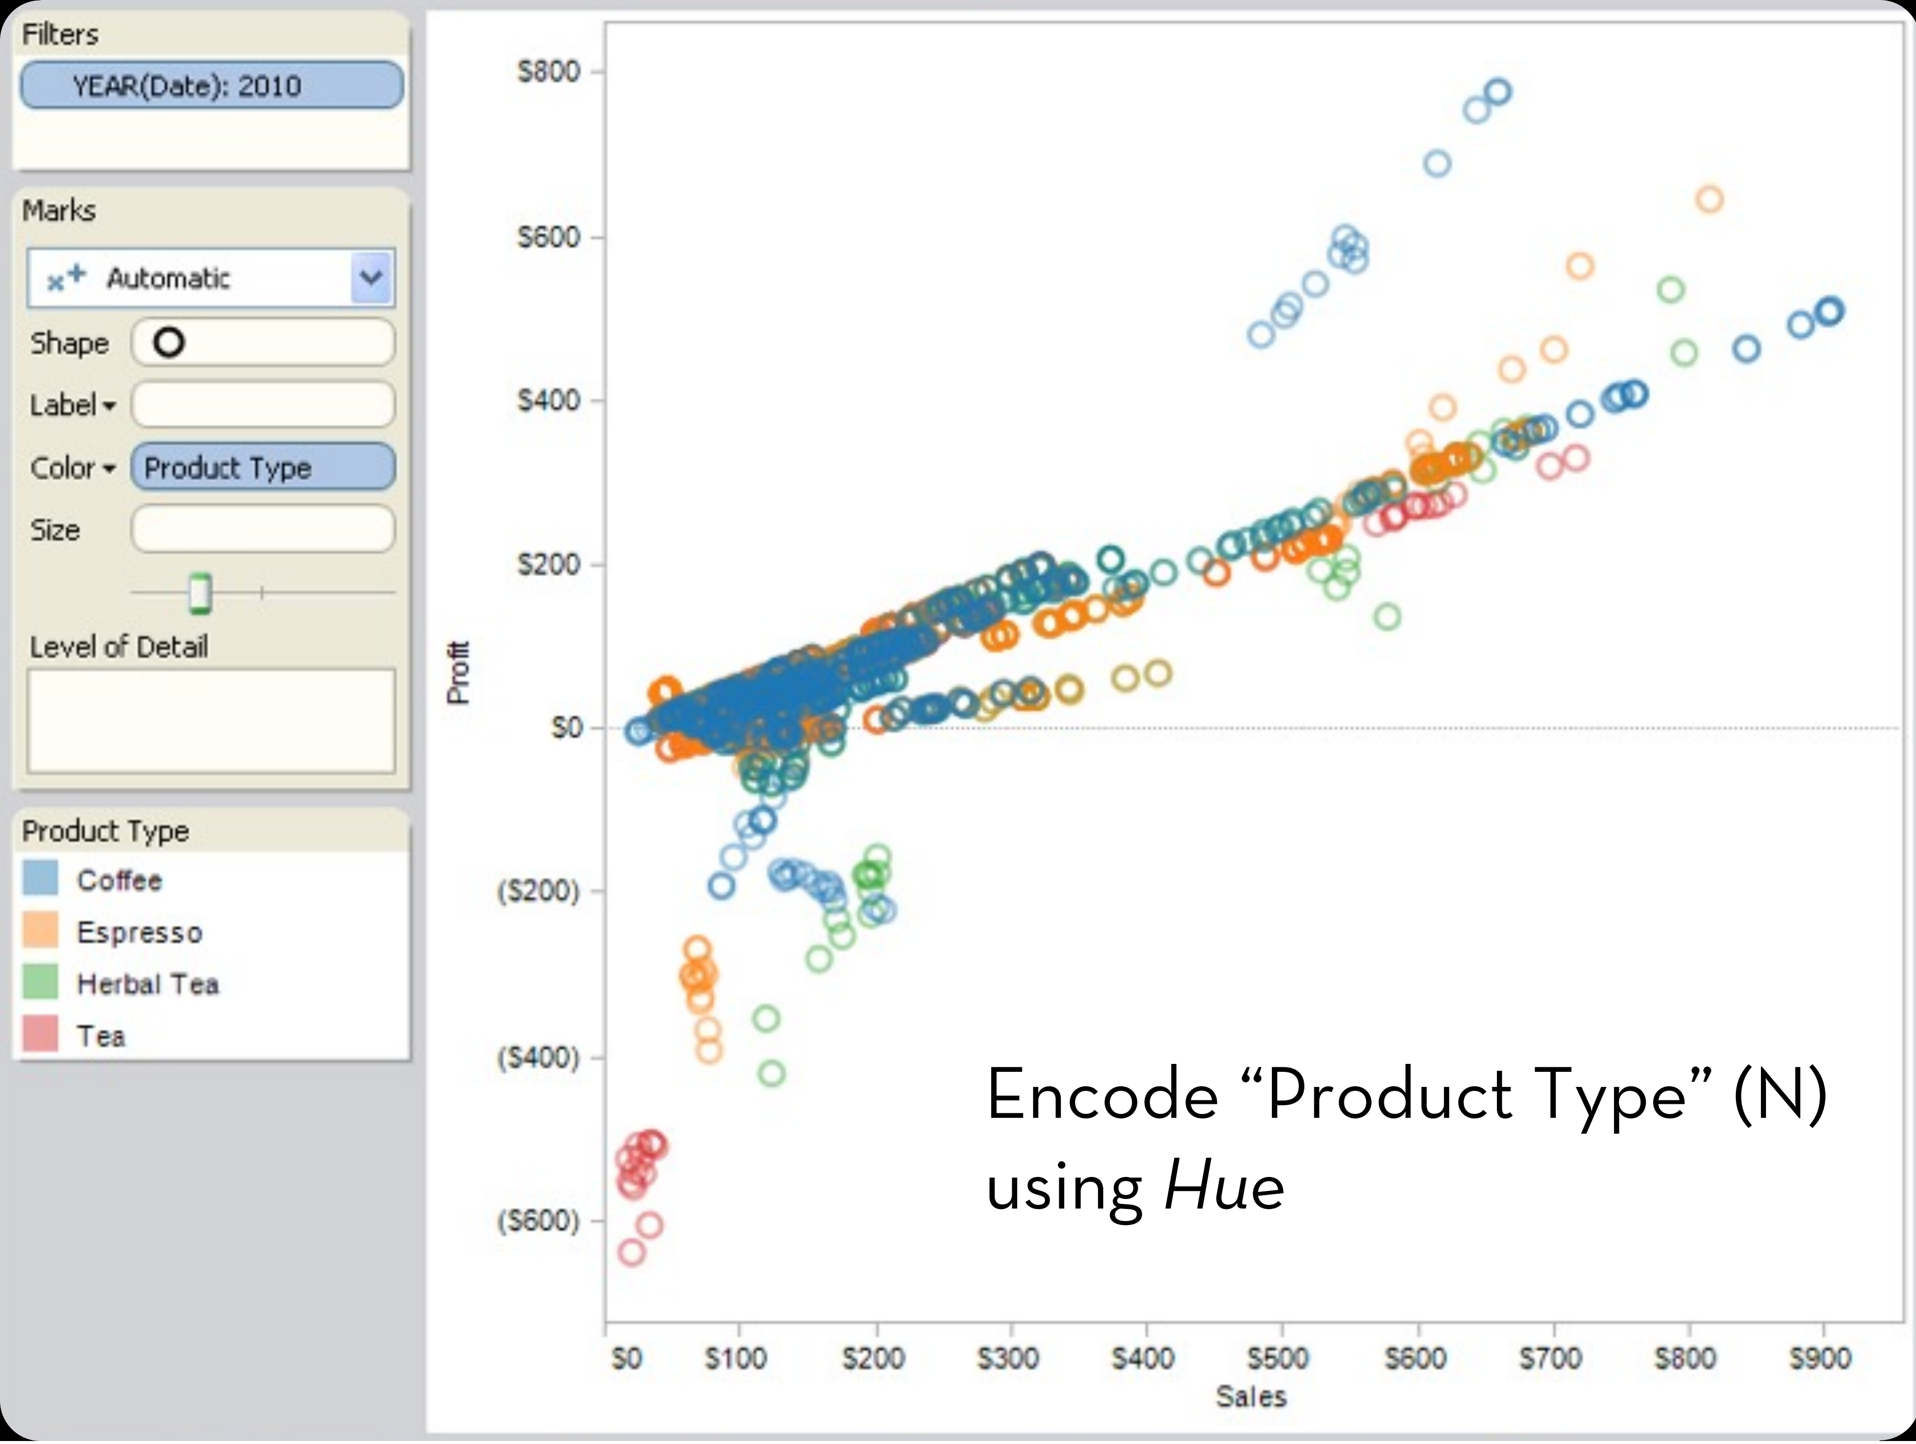

Encode "Product Type" (N) using *Hue*

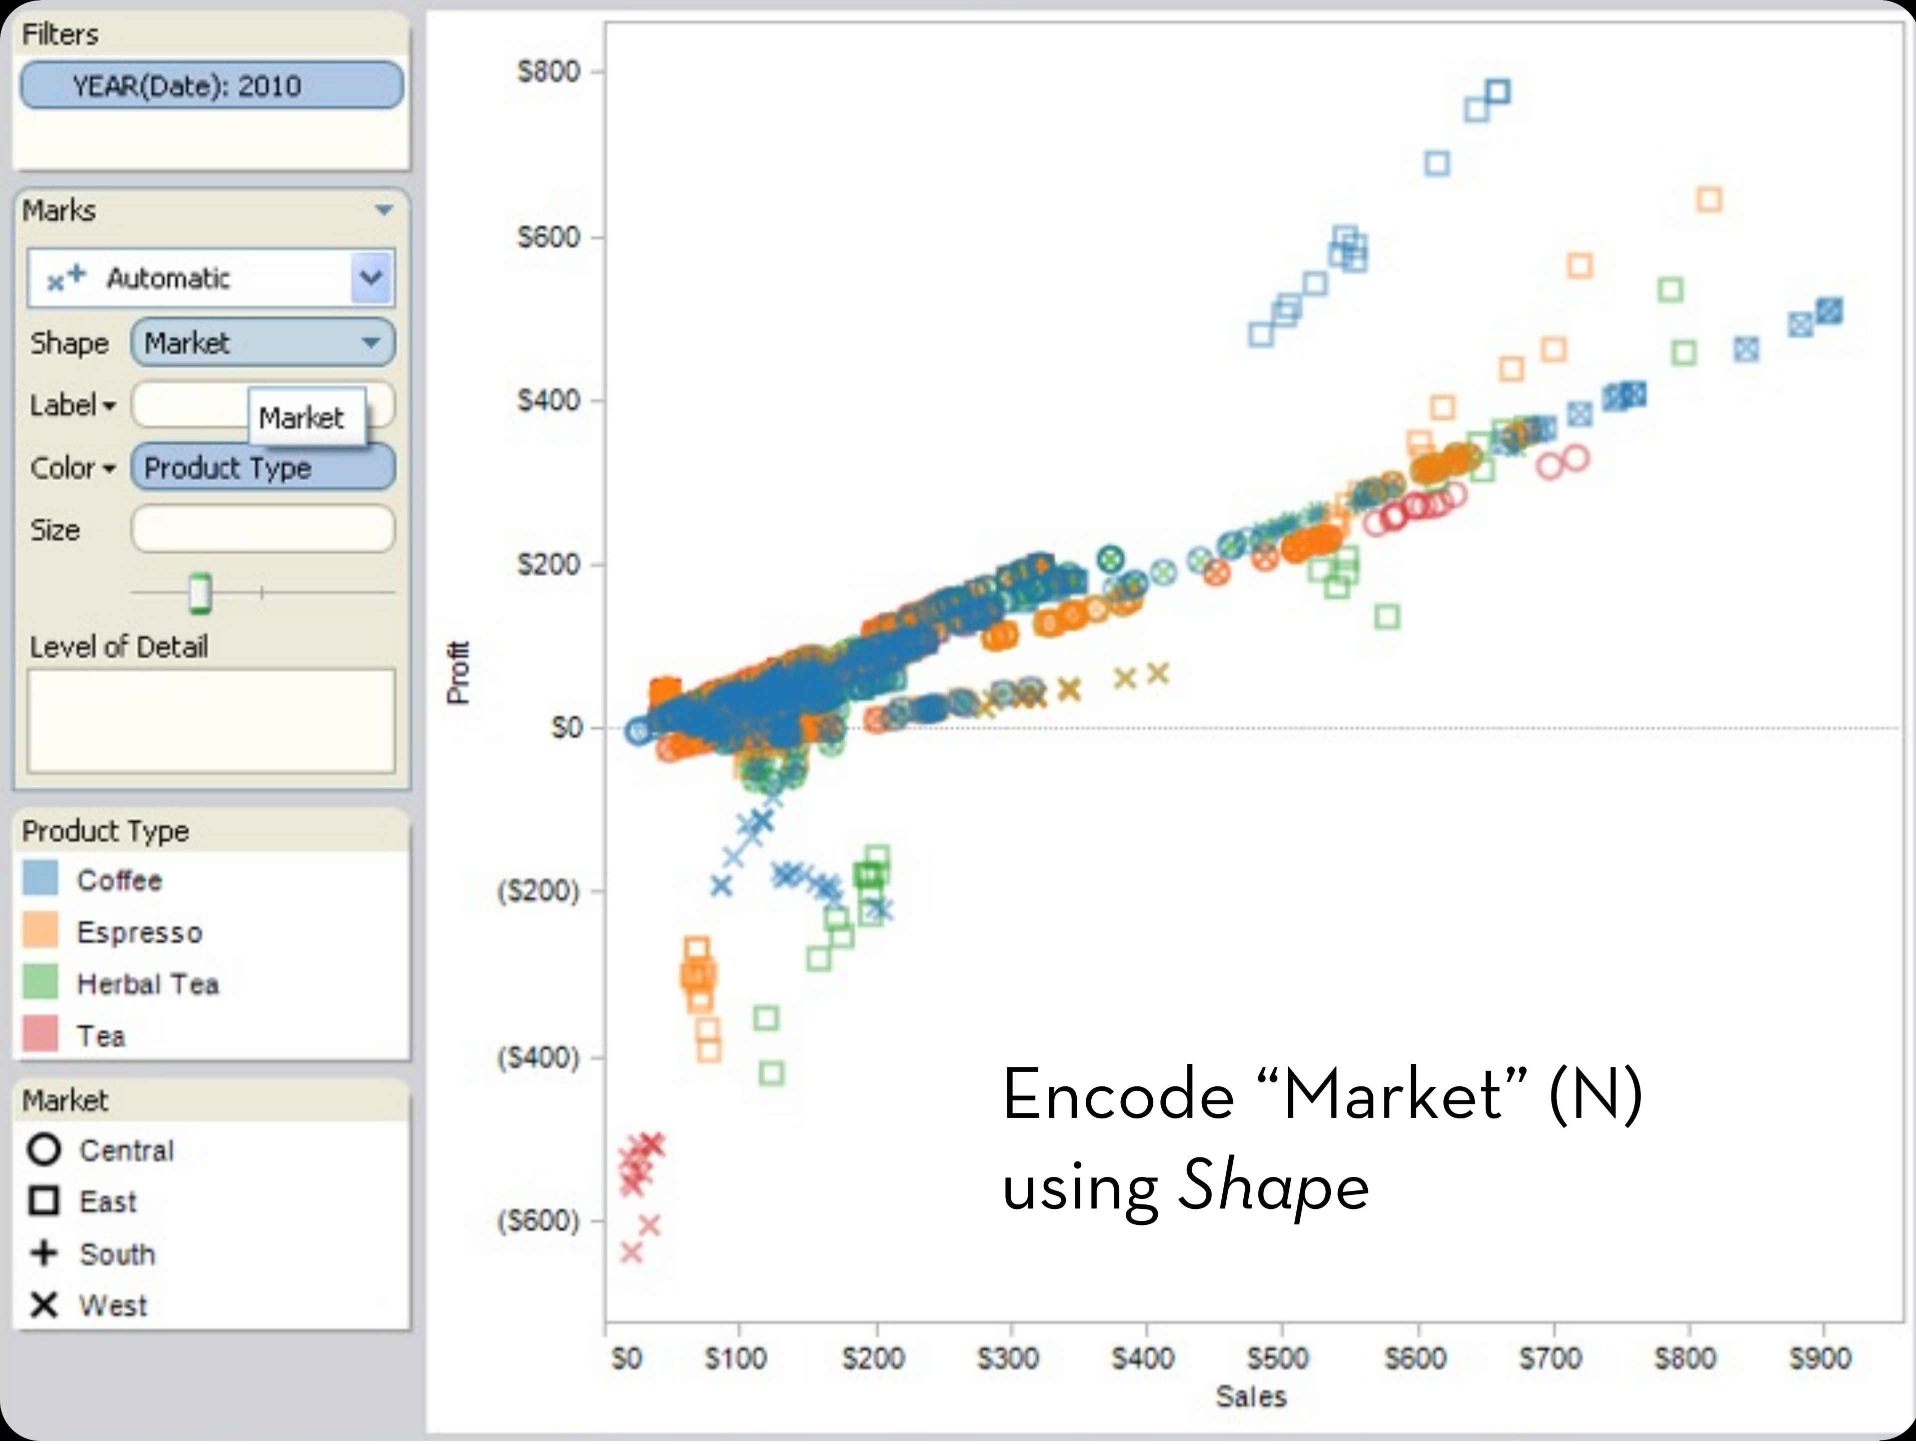

Encode "Market" (N) using *Shape*

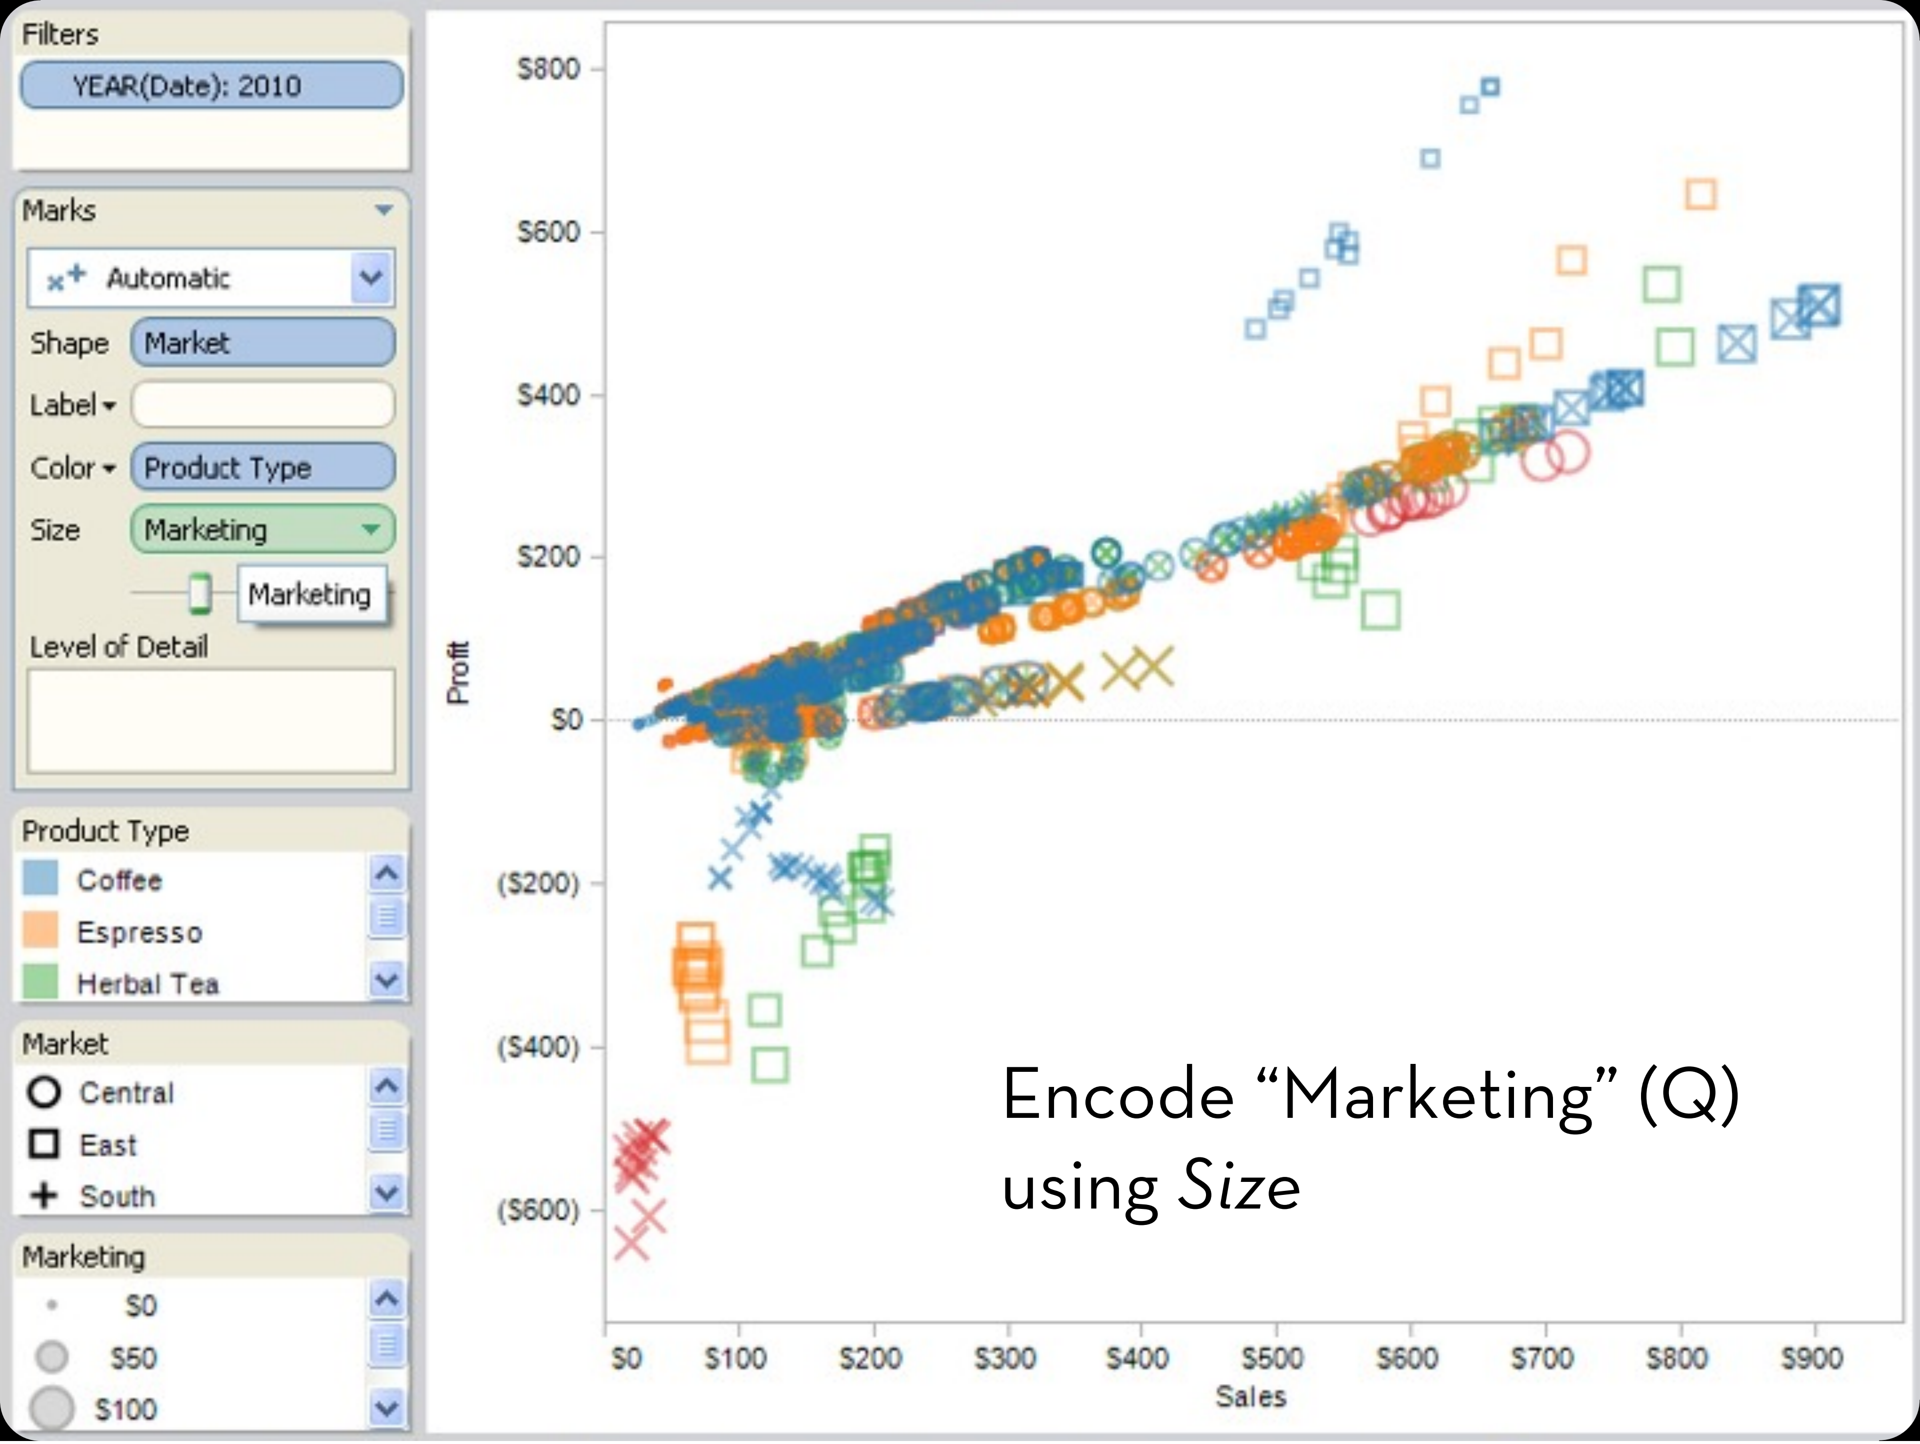

Encode "Marketing" (Q) using *Size*

Trellis Plots

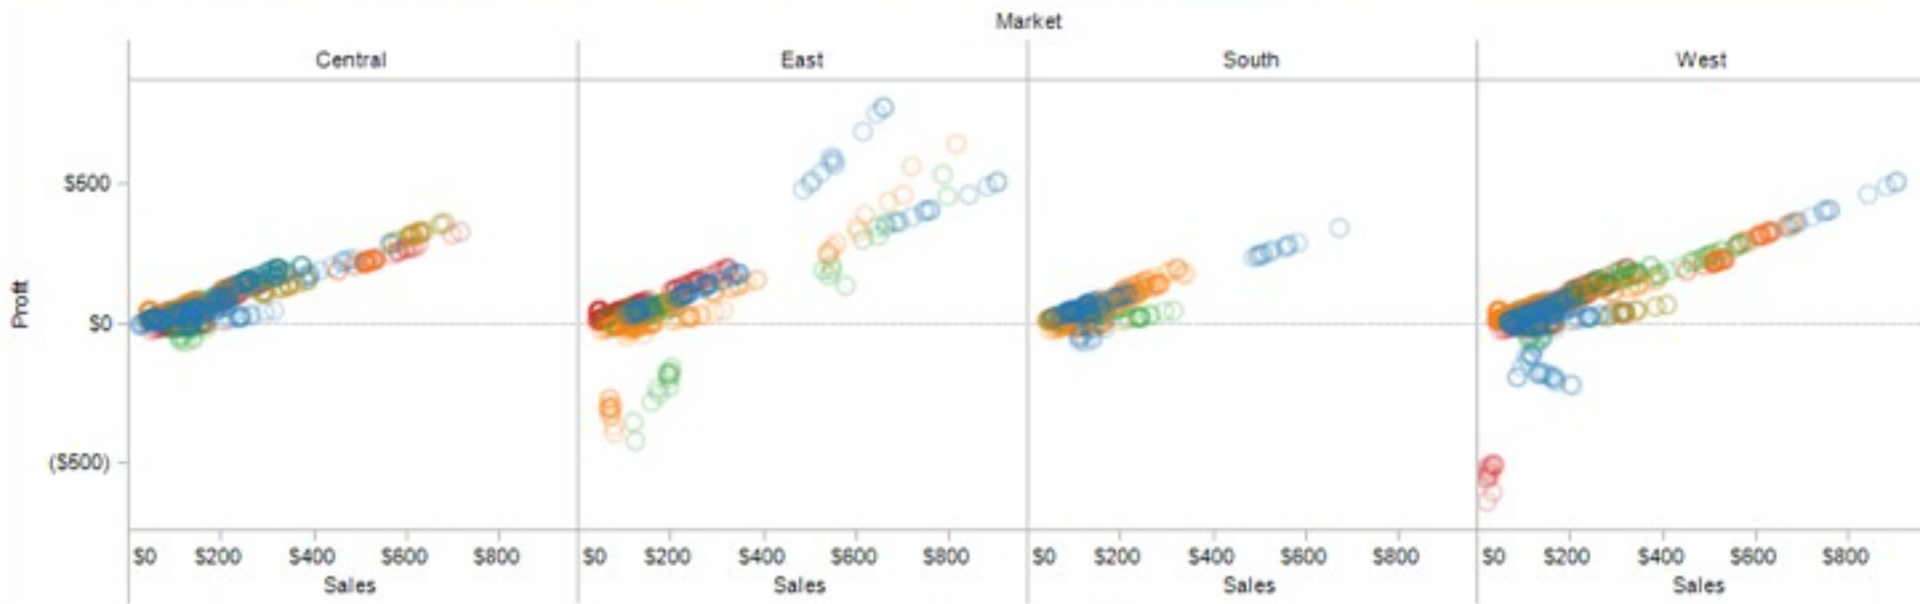

A *trellis plot* subdivides space to enable comparison across multiple plots.

Typically nominal or ordinal variables are used as dimensions for subdivision.

Separation: Small Multiples

[Figure 2.11, p. 38, MacEachren 95]

Separation: Small Multiples

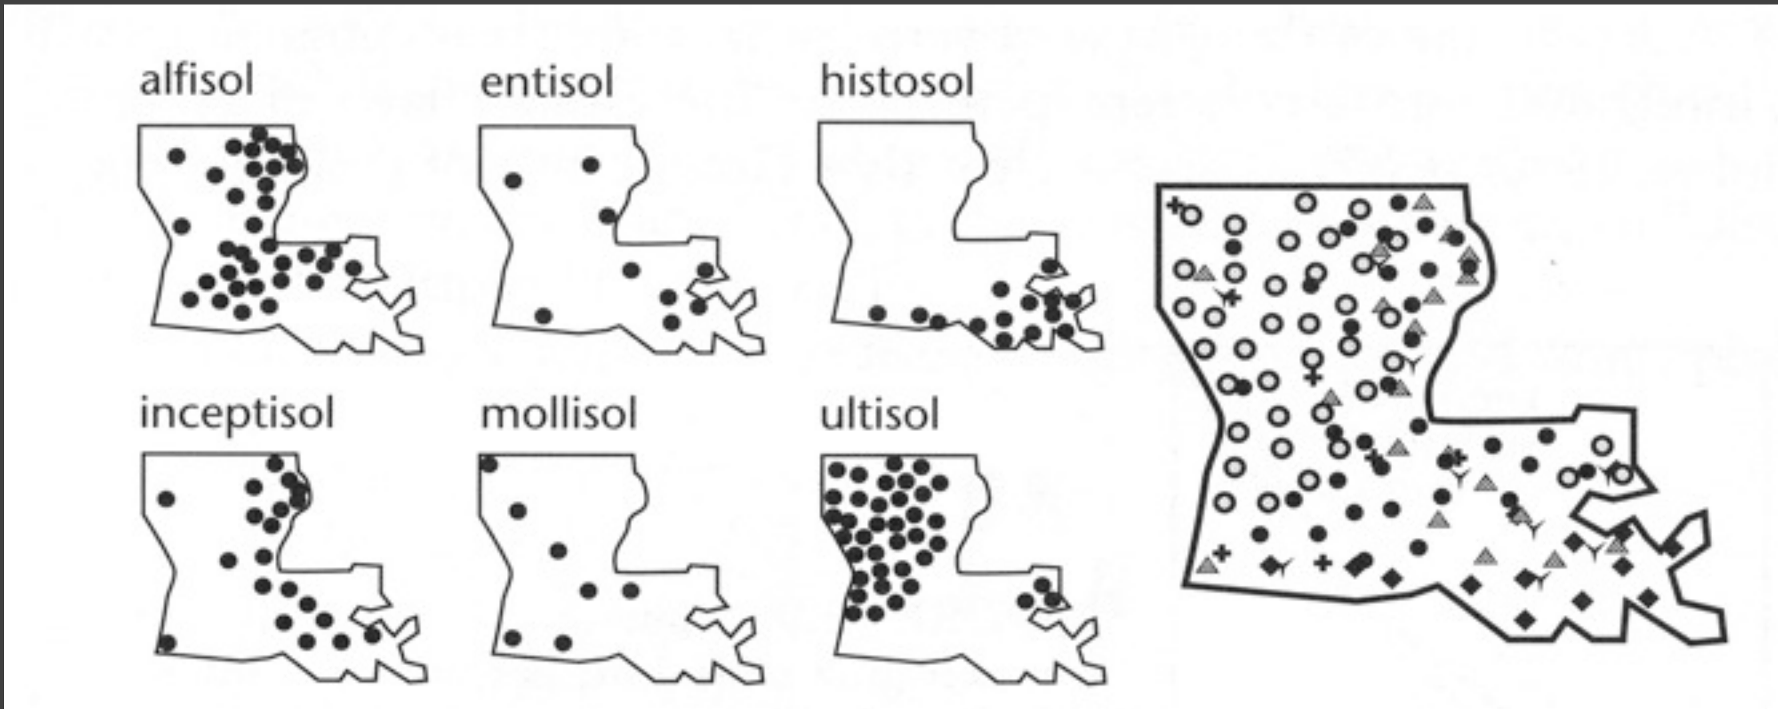

[Figure 2.11, p. 38, MacEachren 95]

Scatterplot Matrix (SPLOM)

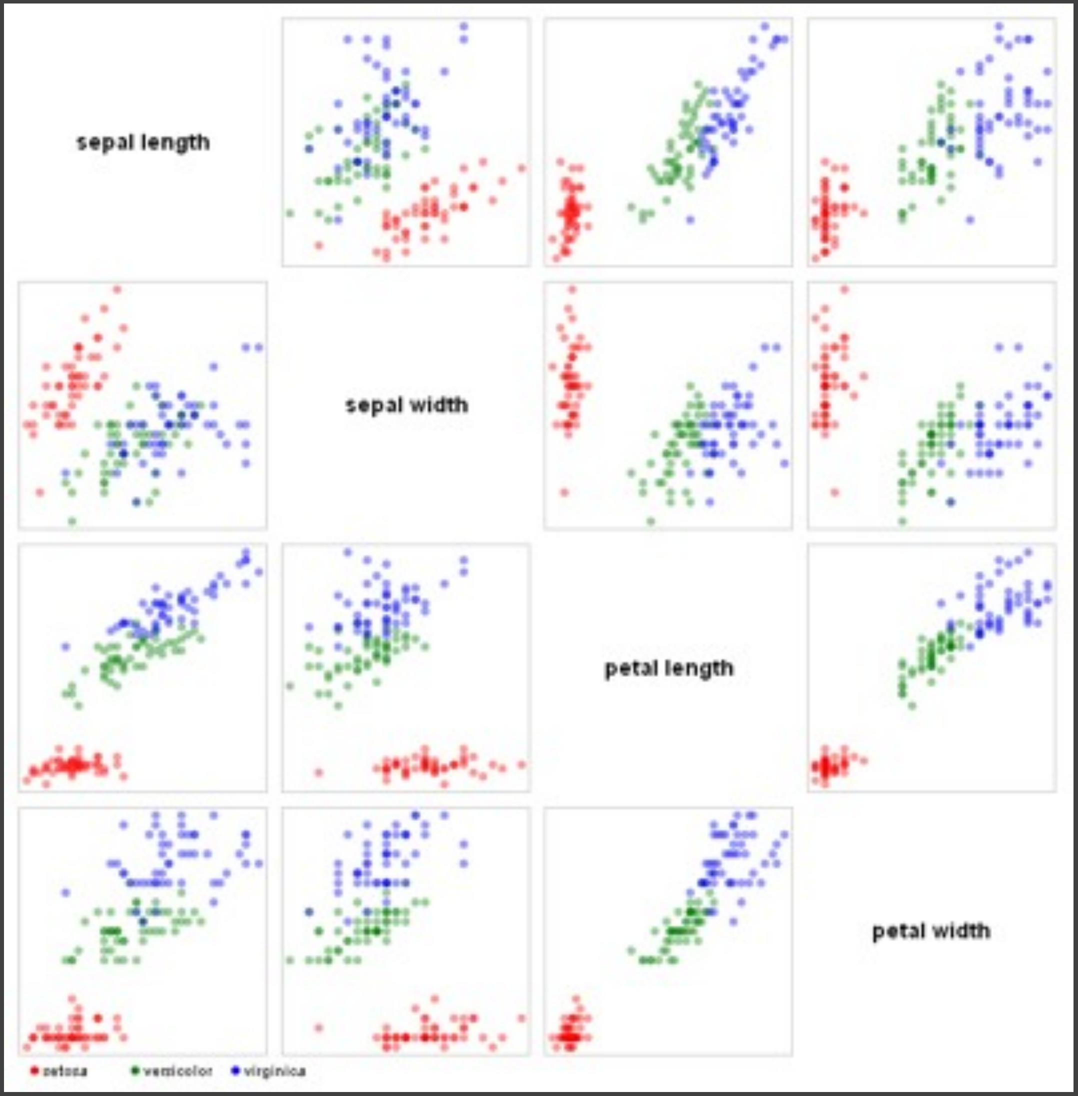

Scatter plots enabling pair-wise comparison of each data dimension.

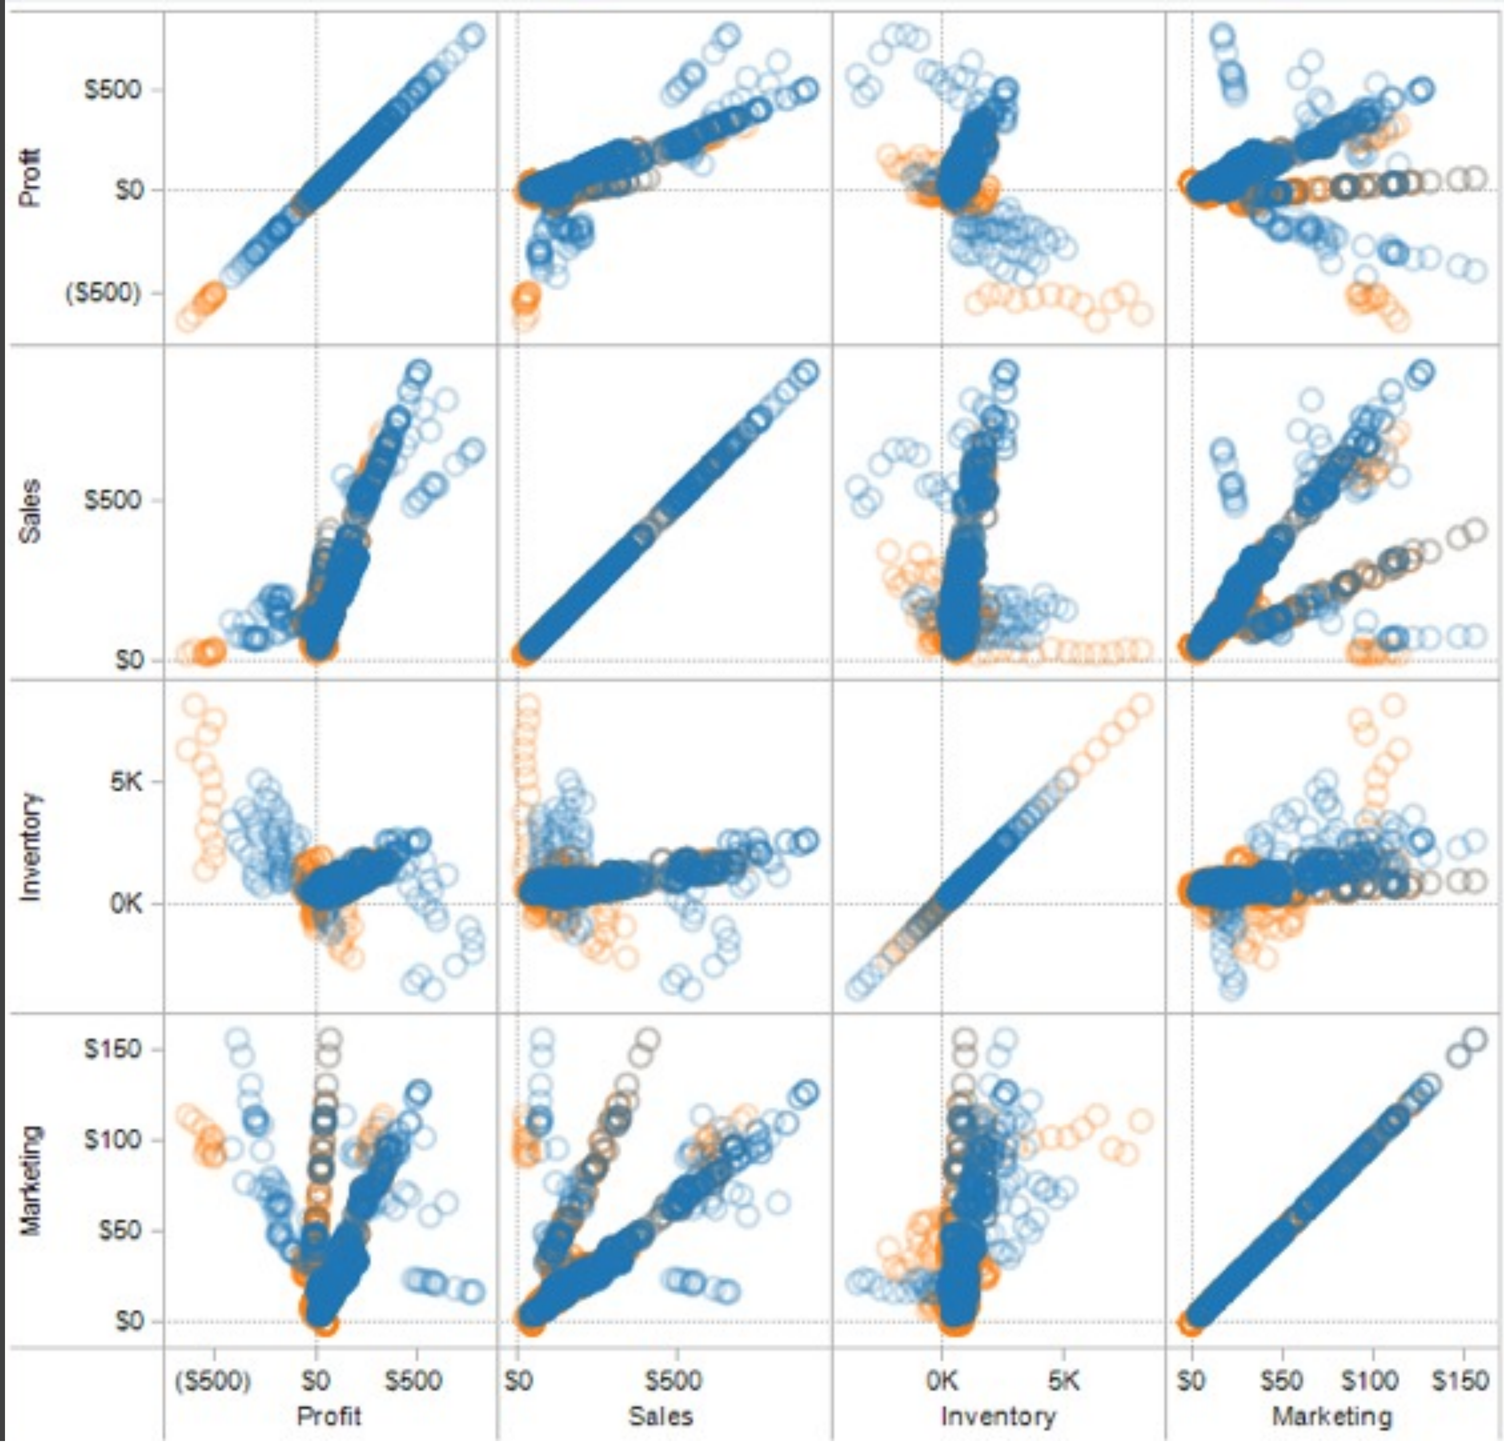
Profit
$500
$0
($500)
Sales
$500
$0
Inventory
5K
0K
Marketing
$150
$100
$50
$0
($500)
$0
$500
Profit
$0
$500
Sales
0K
5K
Inventory
$0
$50
$100
$150
Marketing

Multiple Coordinated Views

how long
in majors

select high
salaries

avg assists vs
avg putouts
(fielding ability)

avg career
HRs vs avg
career hits
(batting ability)

distribution
of positions
played

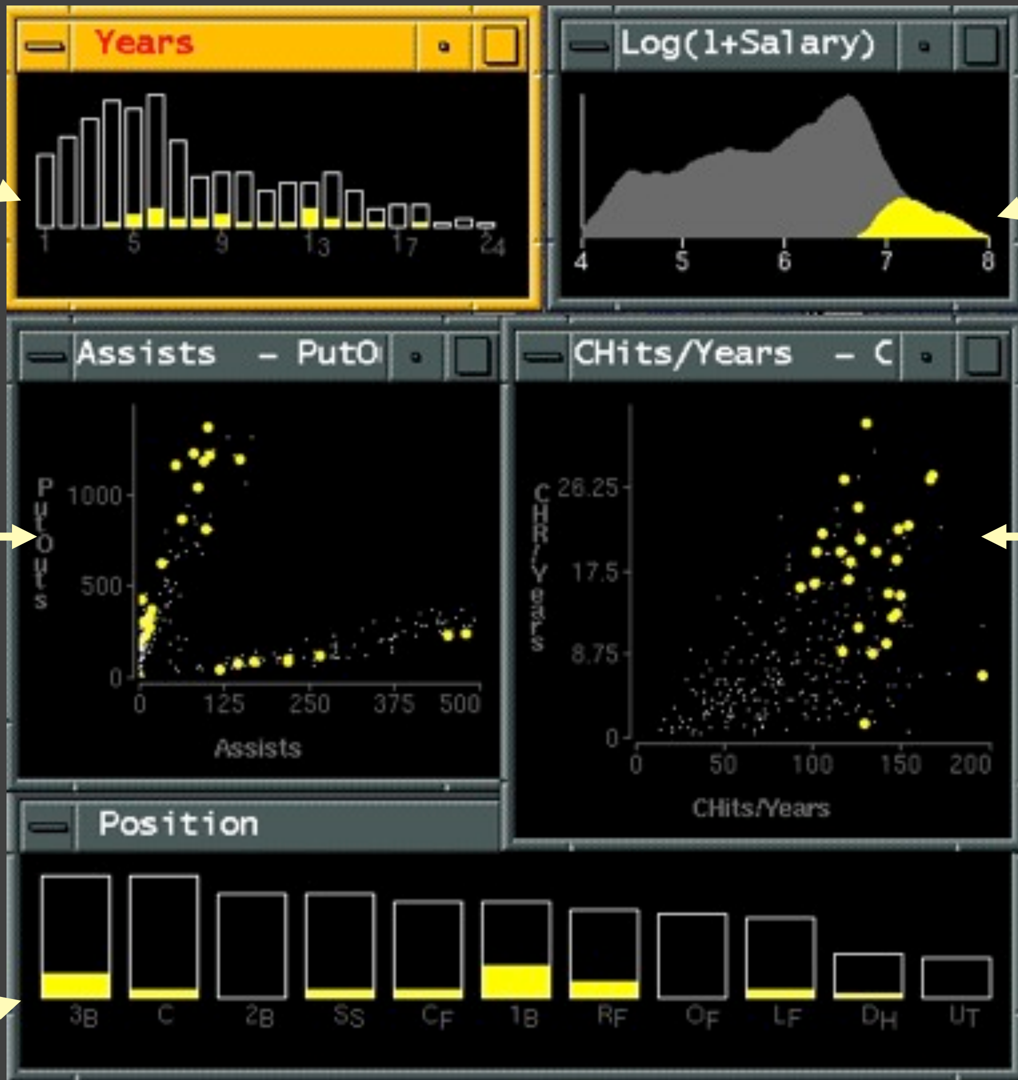

Linking Assists to Positions

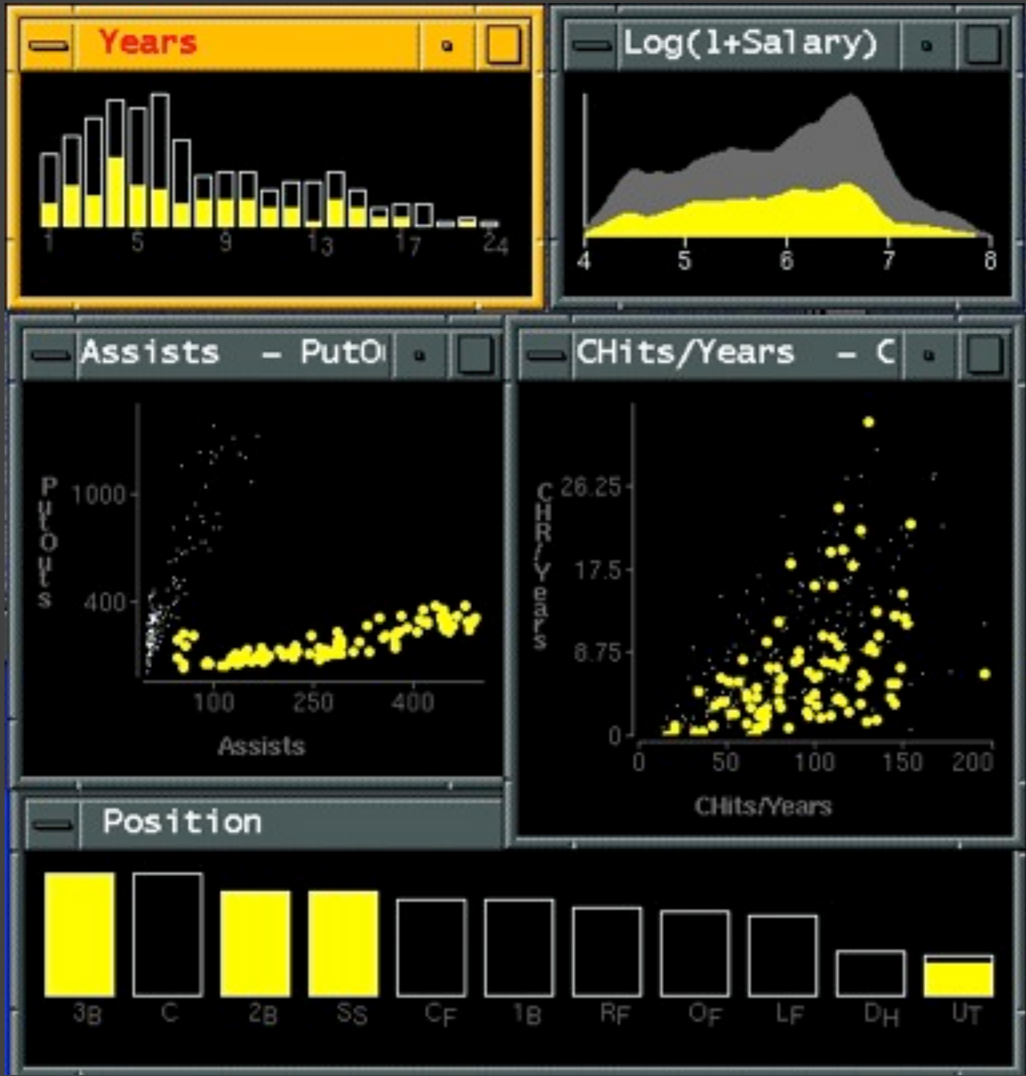

Dimensionality Reduction

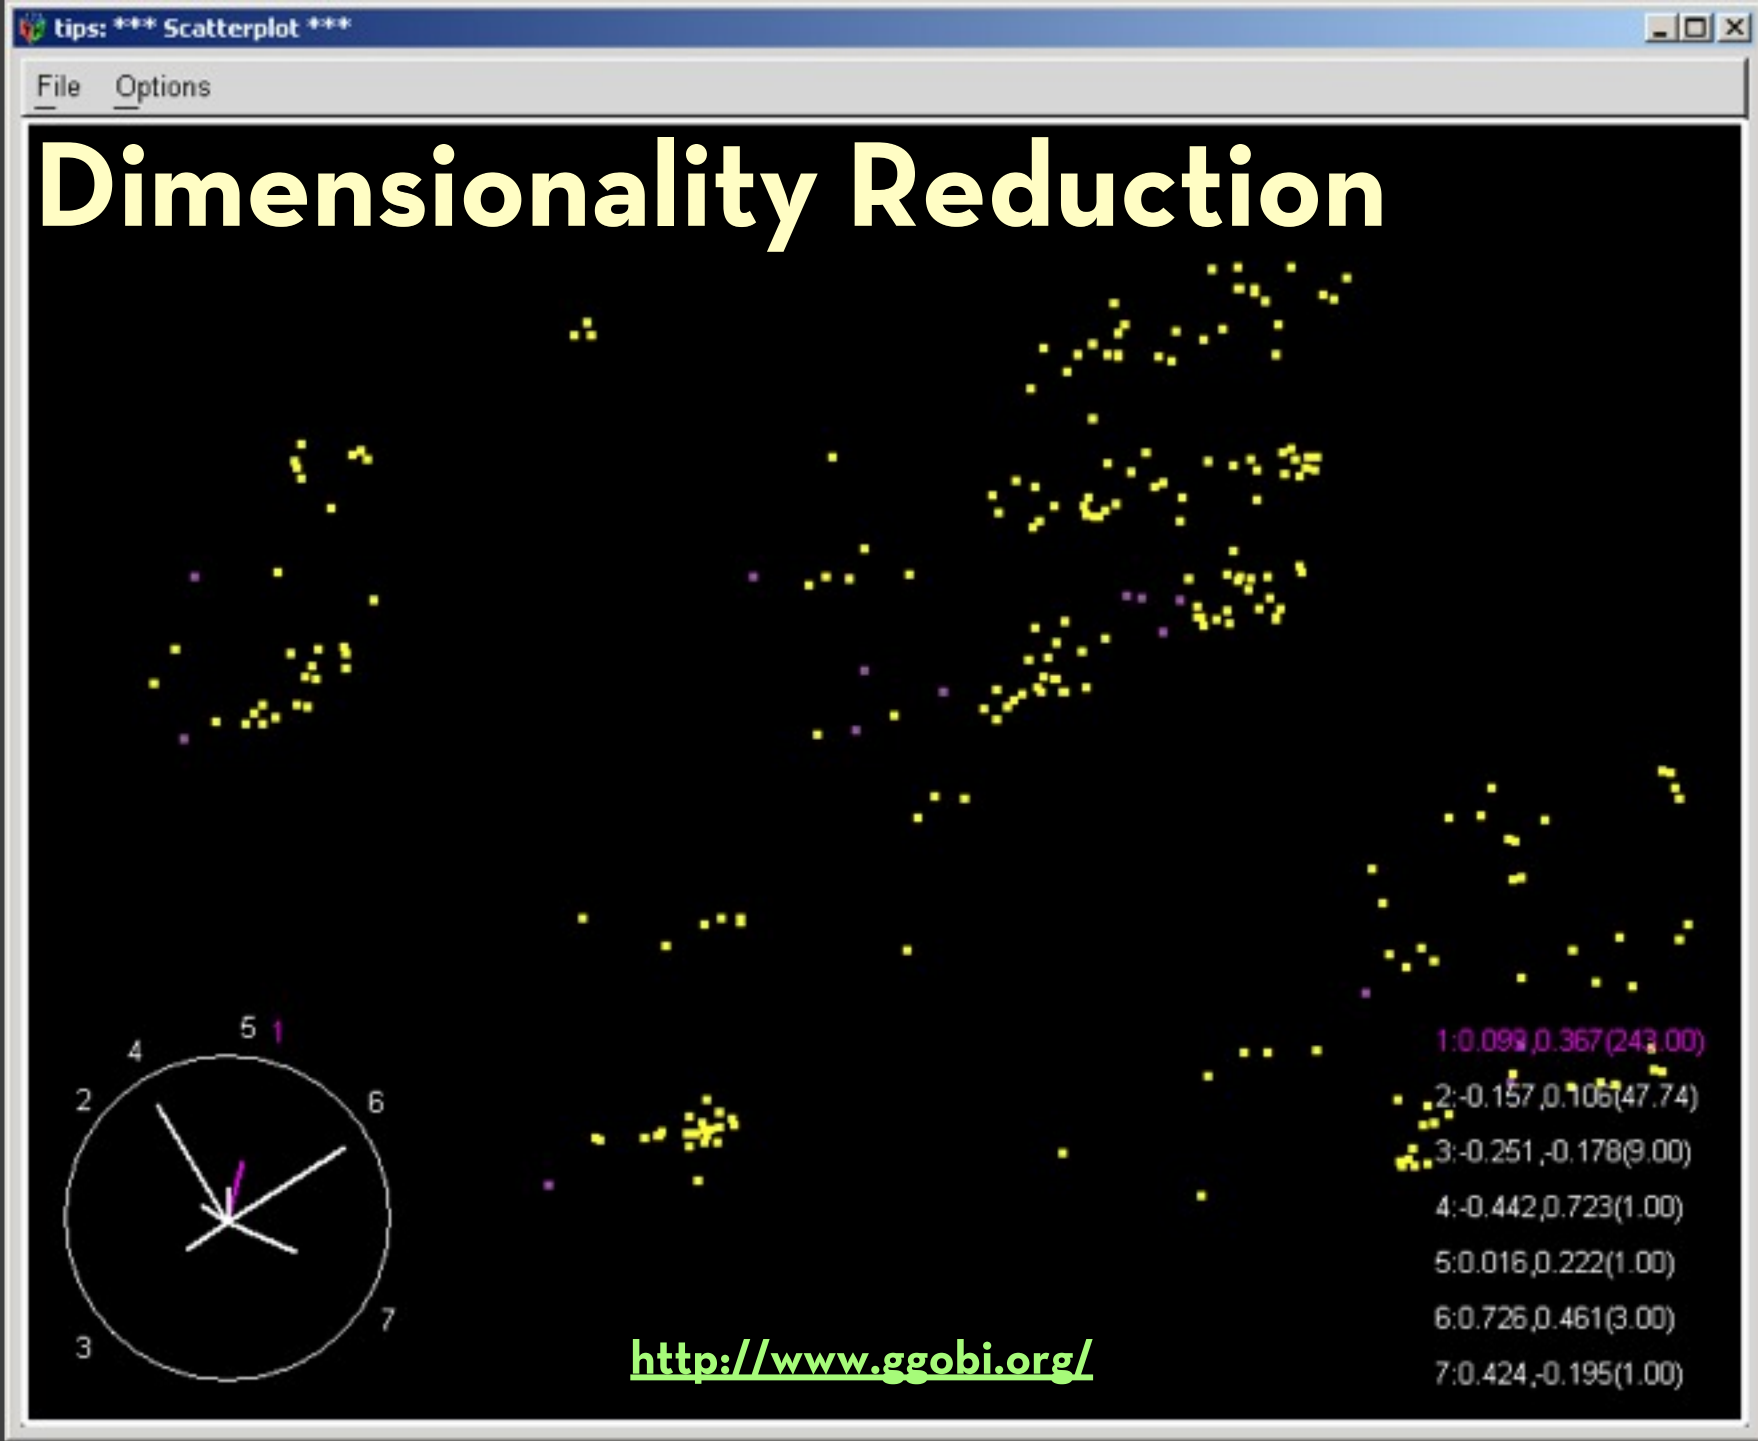

http://www.ggobi.org/

Principal Component Analysis

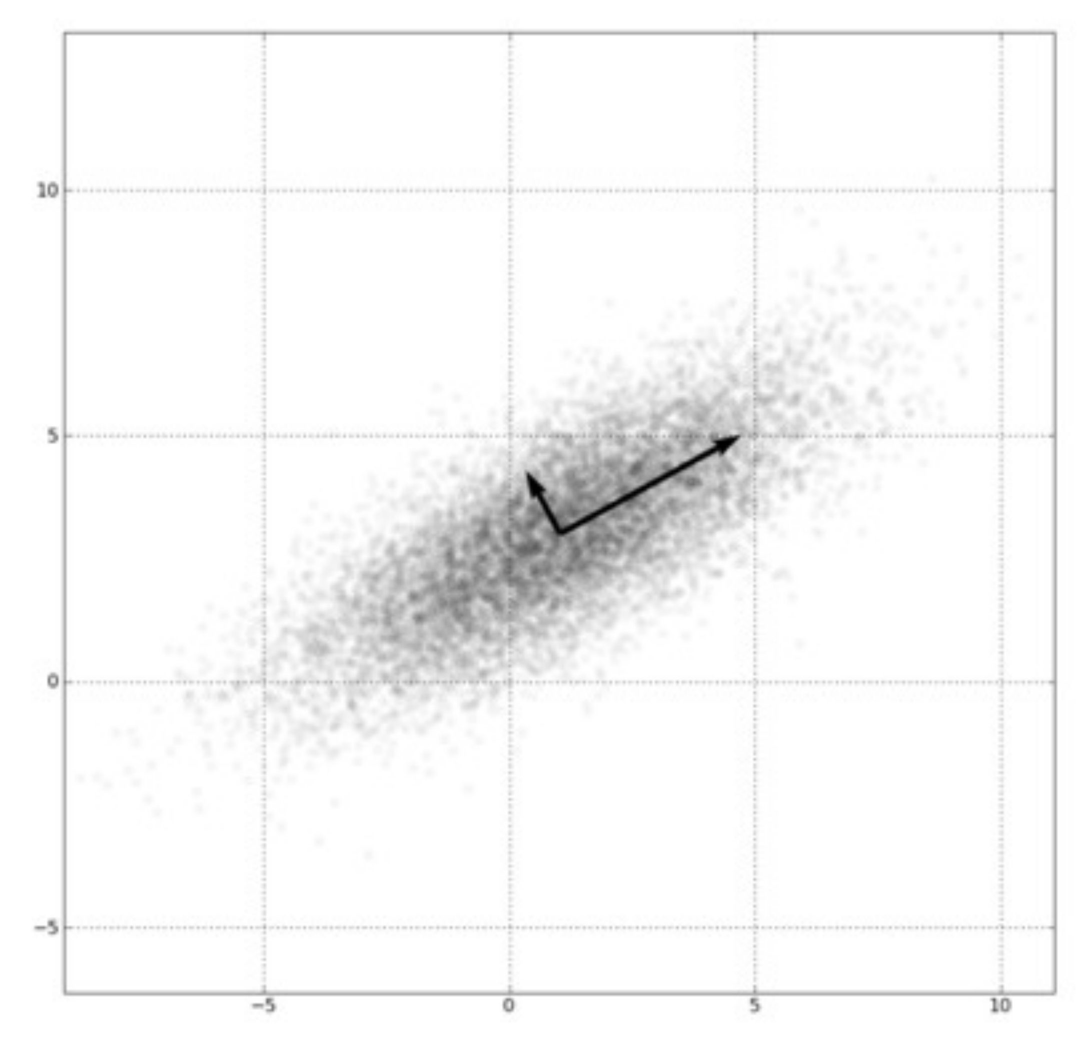

1. Mean-center the data.
2. Find ⊥ basis vectors that maximize the data variance.
3. Plot the data using the top vectors.

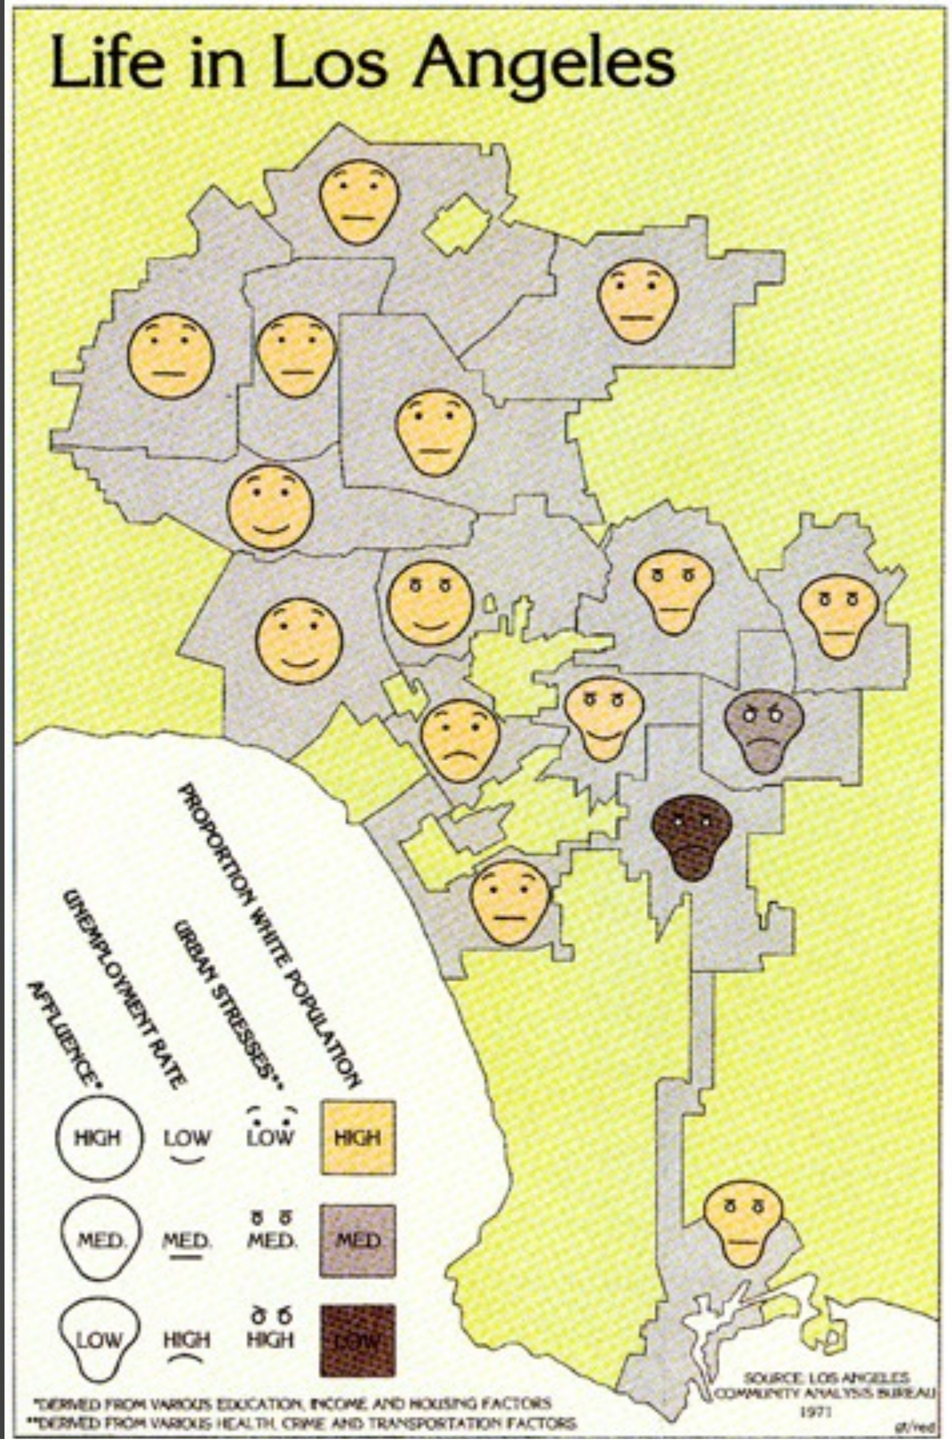

Life in Los Angeles
PROPORTION WHITE POPULATION
URBAN STRESSES**
UNEMPLOYMENT RATE
AFFLUENCE*
HIGH
LOW
LOW
HIGH
MED.
MED.
MED.
MED
LOW
HIGH
HIGH
LOW
*DERIVED FROM VARIOUS EDUCATION, INCOME AND HOUSING FACTORS
**DERIVED FROM VARIOUS HEALTH, CRIME AND TRANSPORTATION FACTORS
SOURCE: LOS ANGELES COMMUNITY ANALYSIS BUREAU 1971

Chernoff Faces (1973)

Observation: We have evolved a sophisticated ability to interpret faces.

Idea: Map data variables to facial features.

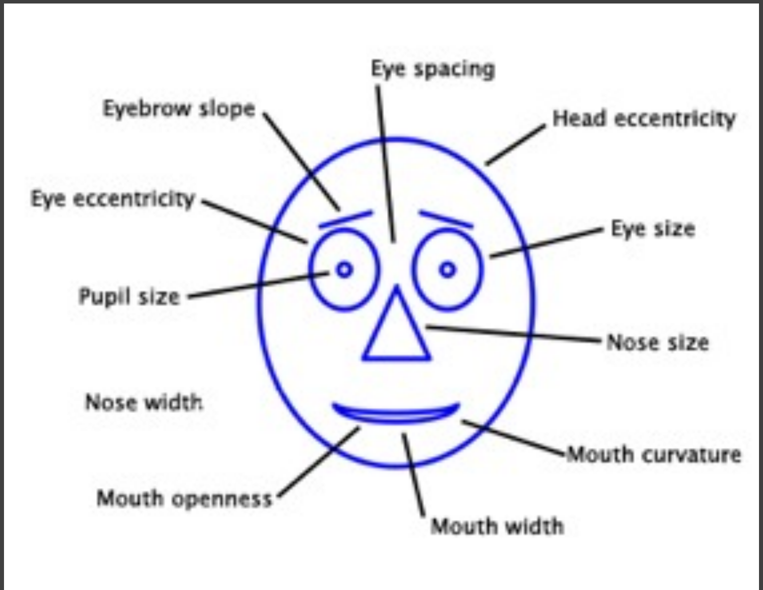

Question: Do we process facial features in an uncorrelated way? (i.e., are they *separable*?)

This is just one example of nD "glyphs"

Visualizing Multiple Dimensions

Strategies

- Avoid “over-encoding”
- Use space and small multiples intelligently
- Reduce the problem space
- Use interaction to generate *relevant* views

There is rarely a single visualization that answers all questions. Instead, the ability to generate appropriate visualizations quickly is key.

Parallel Coordinates

Parallel Coordinates [Inselberg]

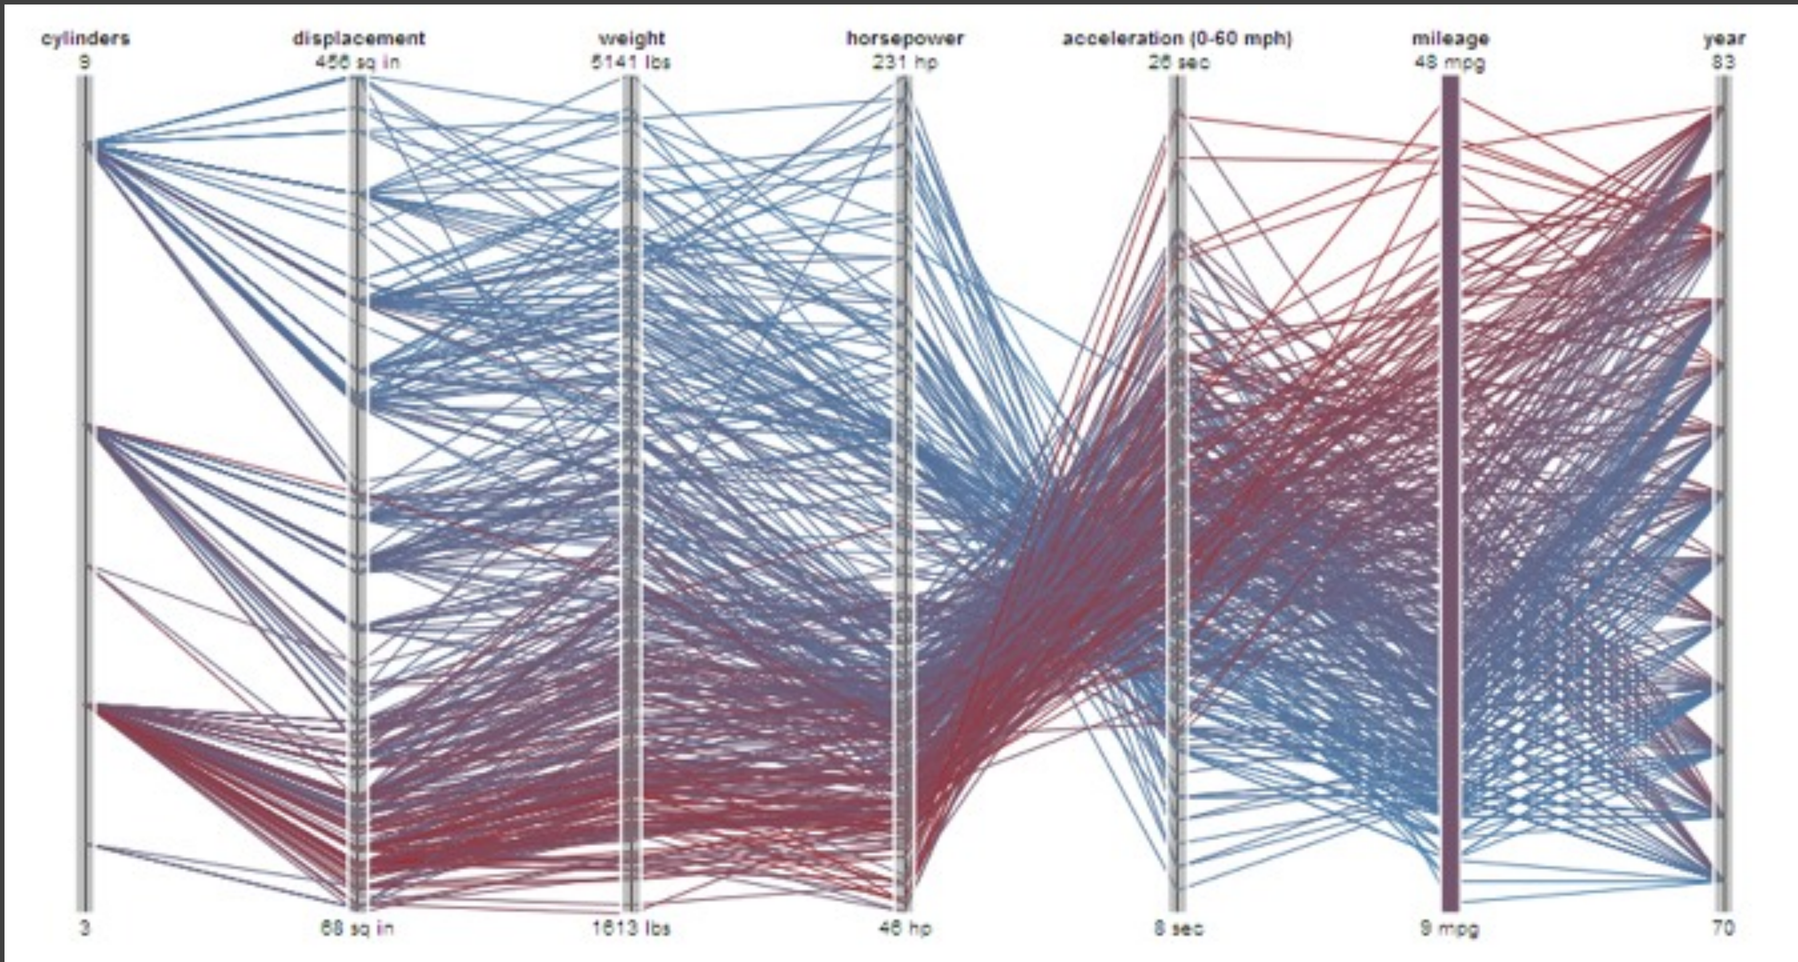

Parallel Coordinates [Inselberg]

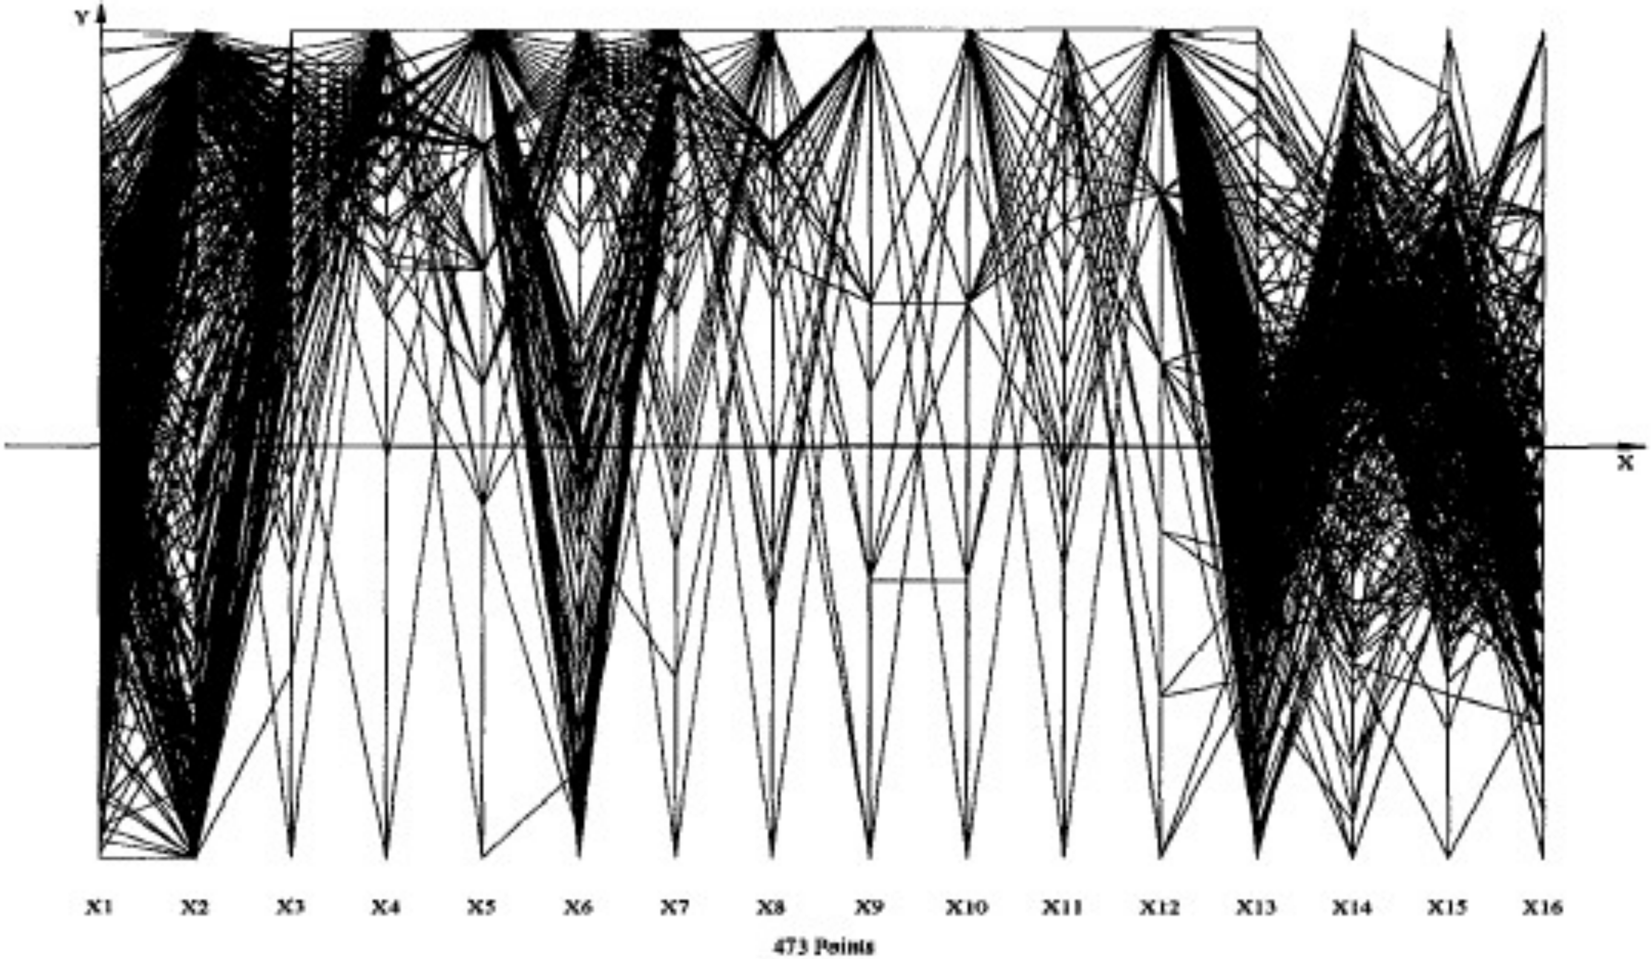

Figure 1: The full dataset consisting of 473 batches

The Multidimensional Detective

The Dataset:

- Production data for 473 batches of a VLSI chip
- 16 process parameters:

X_1: The yield: % of produced chips that are useful
X_2: The quality of the produced chips (speed)
X_3 ... X_{12}: 10 types of defects (zero defects shown at top)
X_{13} ... X_{16}: 4 physical parameters

The Objective:
Raise the yield (X_1) and maintain high quality (X_2)

A. Inselberg, Multidimensional Detective, Proceedings of IEEE Symposium on Information Visualization (InfoVis '97), 1997

Parallel Coordinates

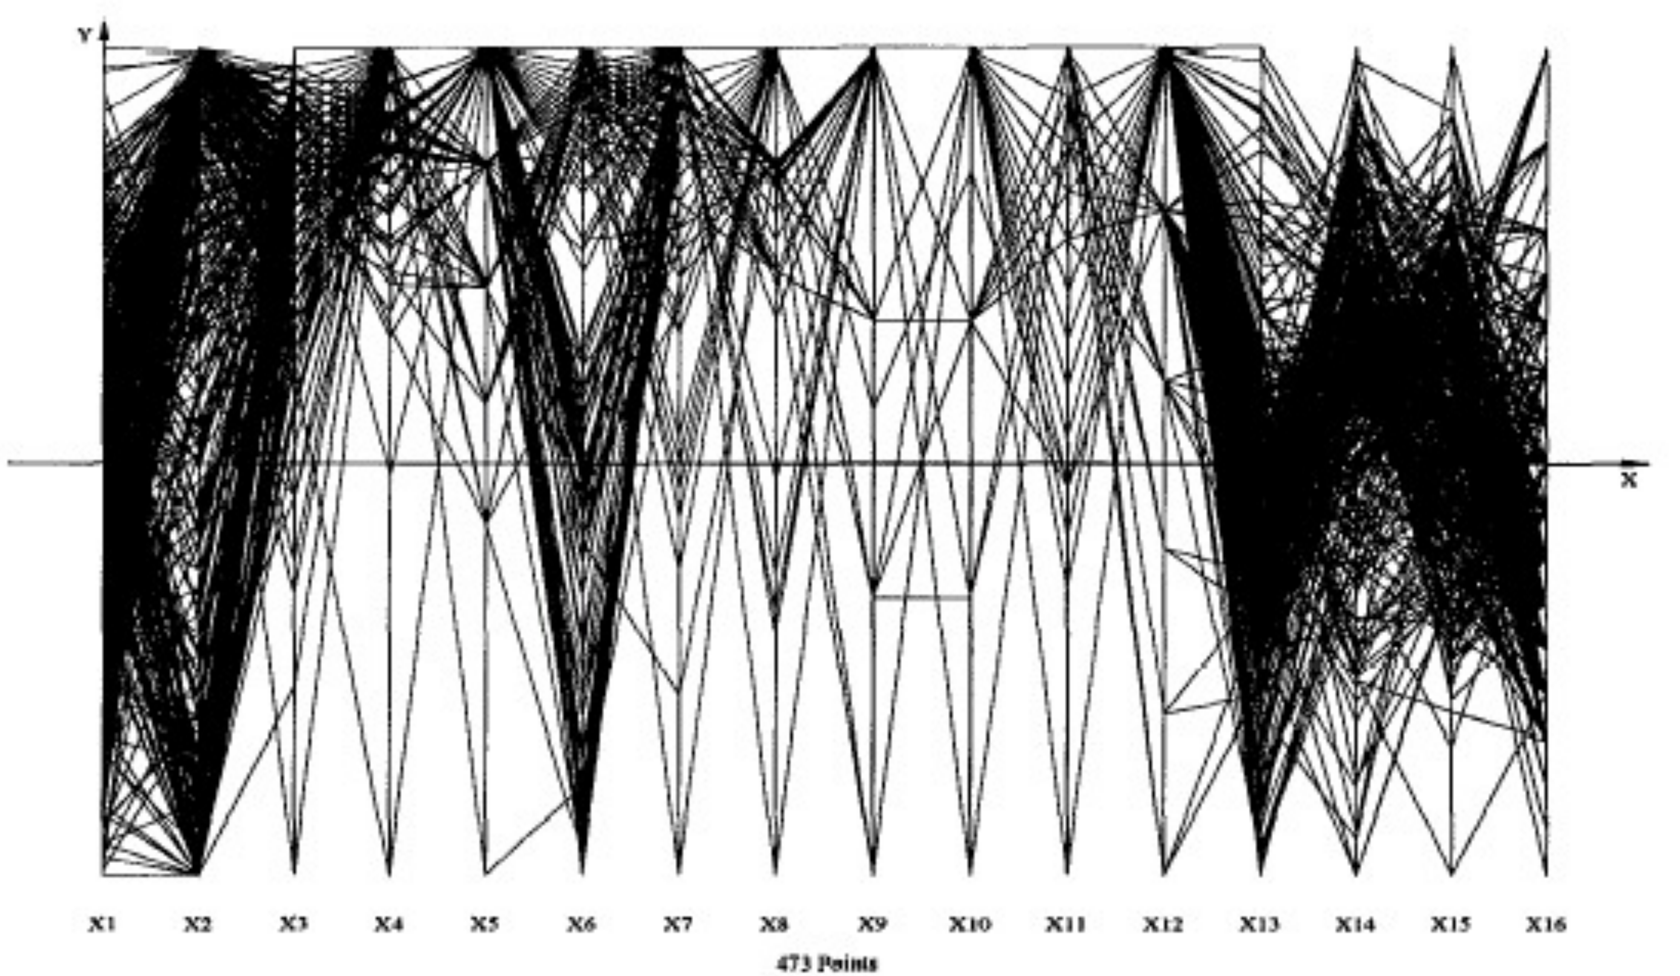

Figure 1: The full dataset consisting of 473 batches

Inselberg's Principles

1. Do not let the picture scare you
2. Understand your objectives
   - Use them to obtain visual cues
3. Carefully scrutinize the picture
4. Test your assumptions, especially the "I am really sure of's"
5. You can't be unlucky all the time!

Each line represents a tuple (e.g., VLSI batch)
Filtered below for high values of X1 and X2

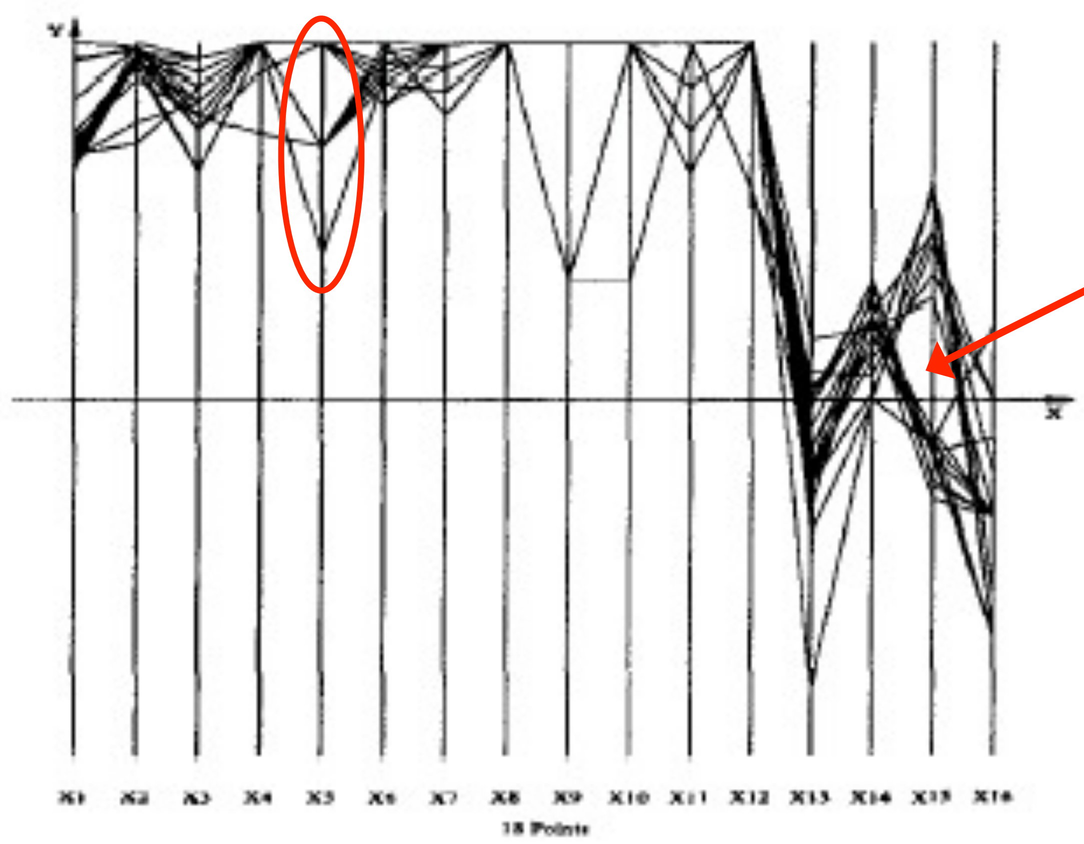

Figure 2: The batches high in Yield, $X1$, and Quality, $X2$.

Look for batches with *nearly* zero defects (9/10)
Most of these have low yields -> defects OK.

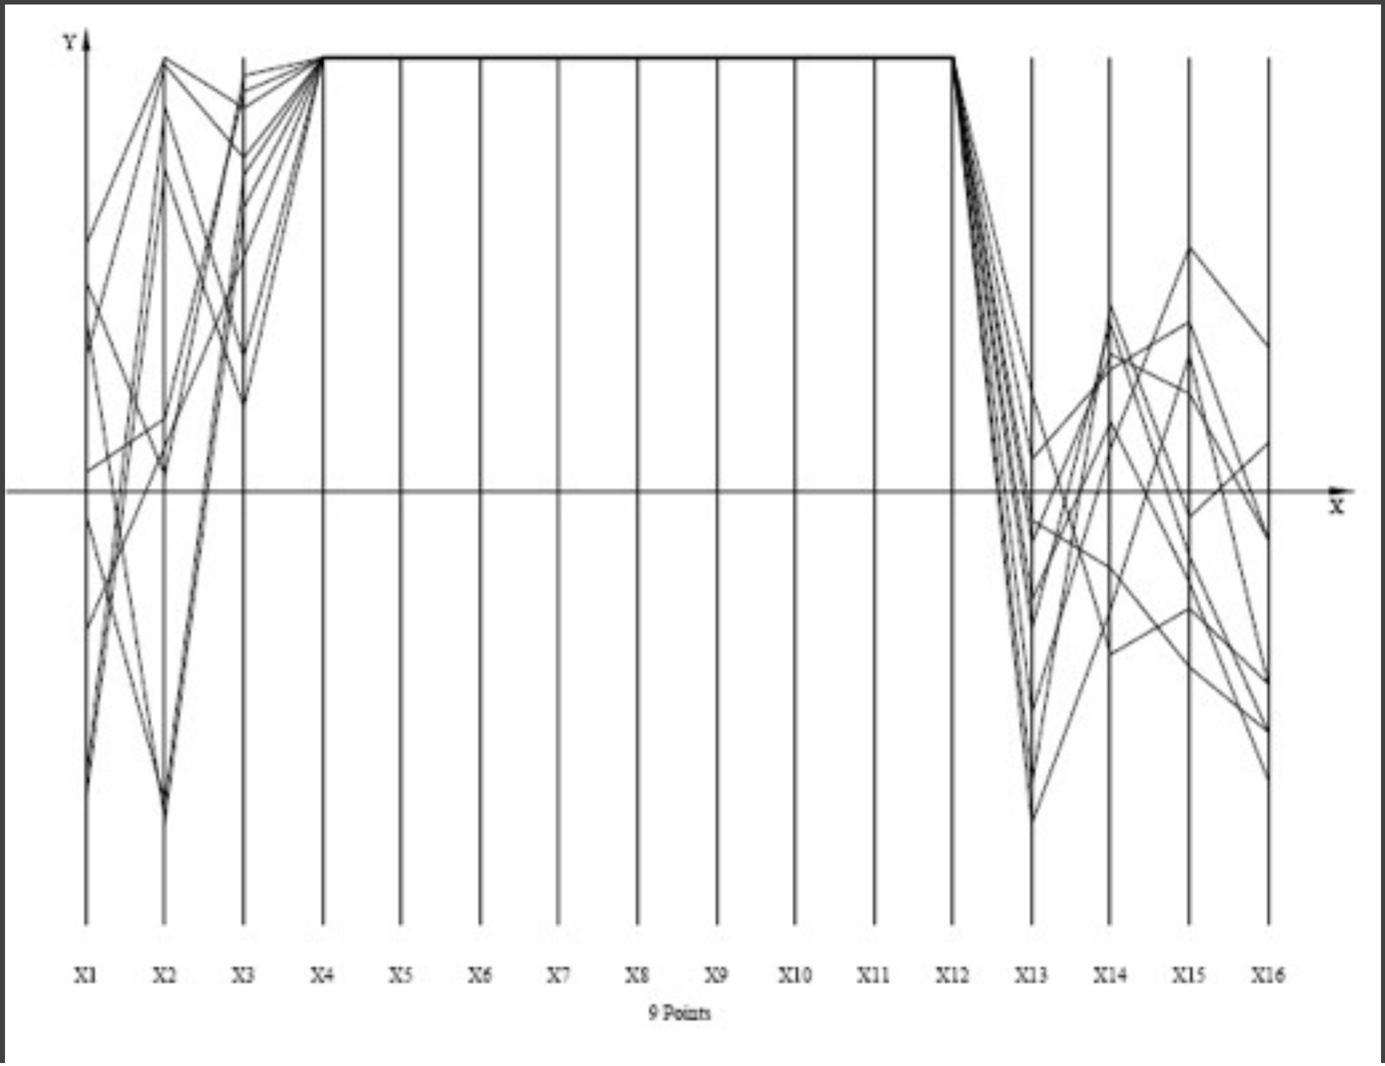

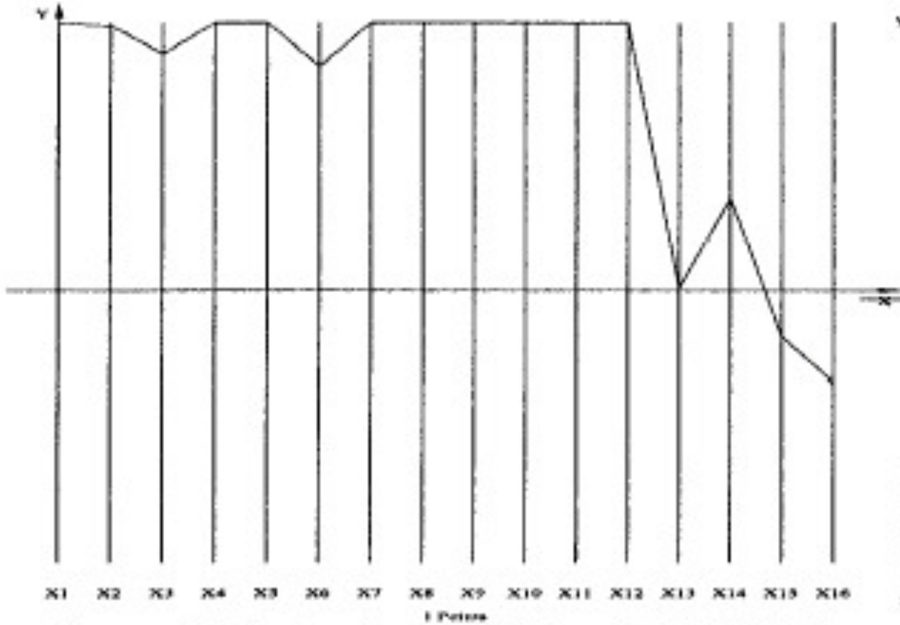

Figure 5: The best batch. Highest in Yield, $X1$, and very high in Quality, $X2$.

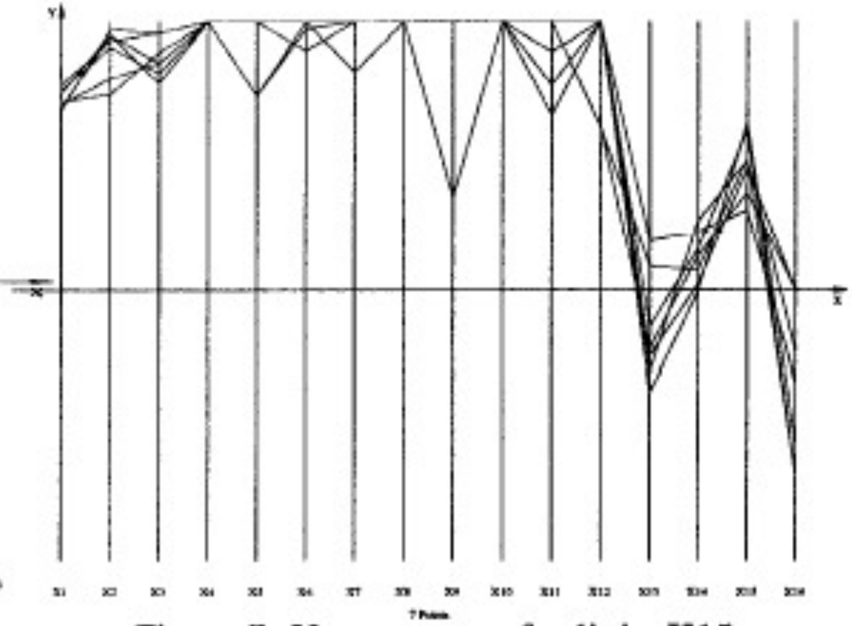

Figure 7: Upper range of split in $X15$

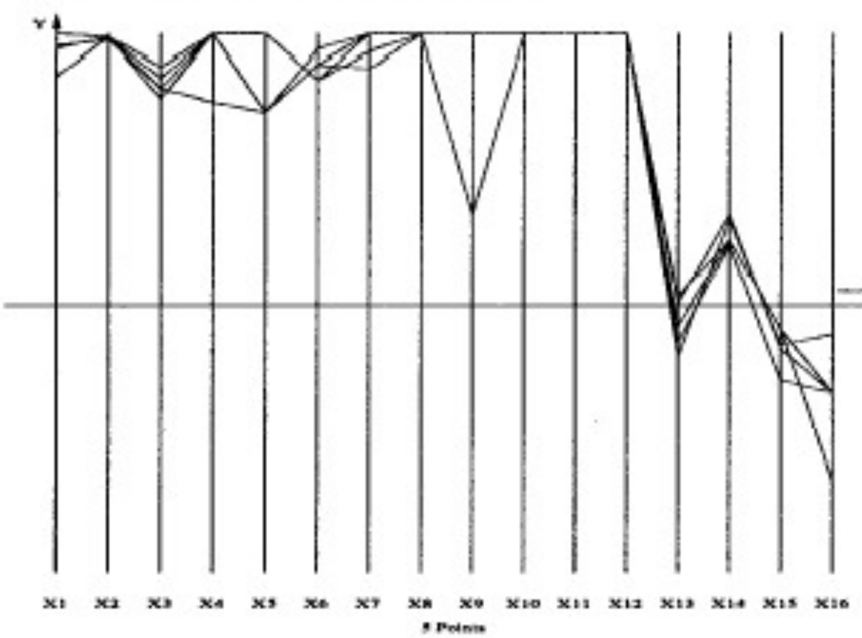

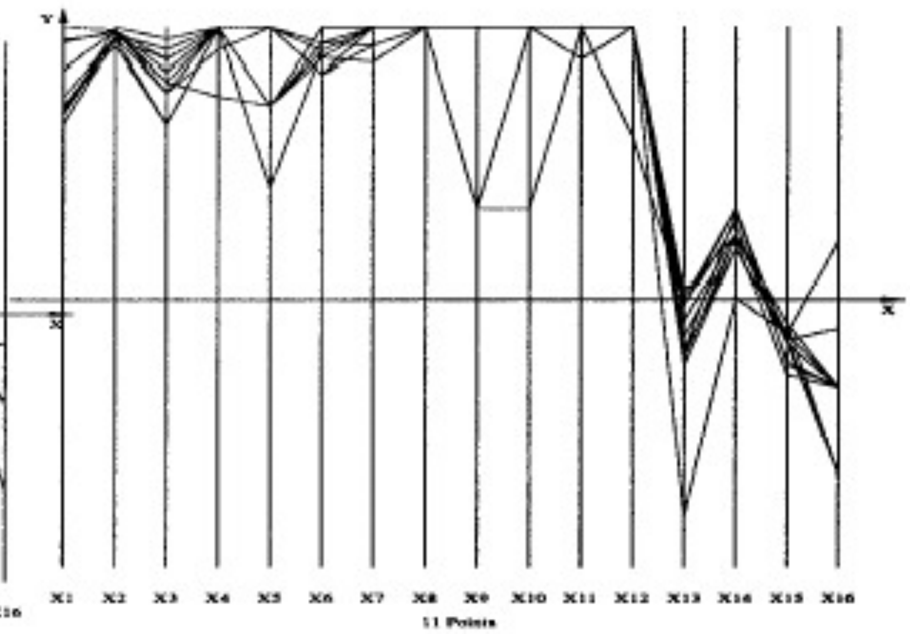

Notice that X6 behaves differently.

Allow 2 defects, including X6 -> best batches

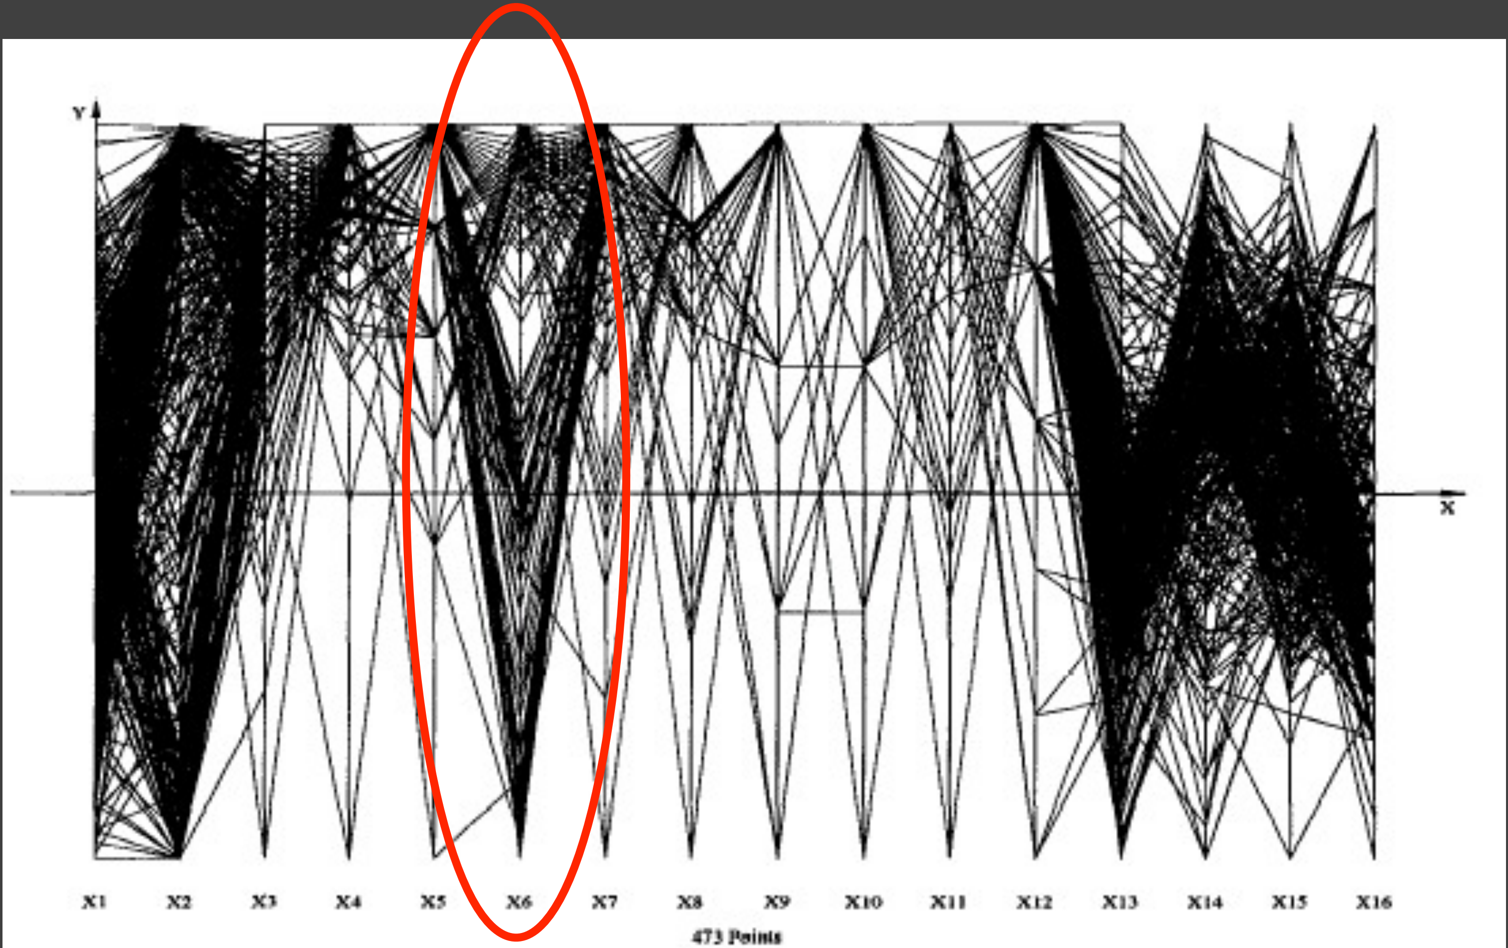

Figure 1: The full dataset consisting of 473 batches

Radar Plot / Star Graph

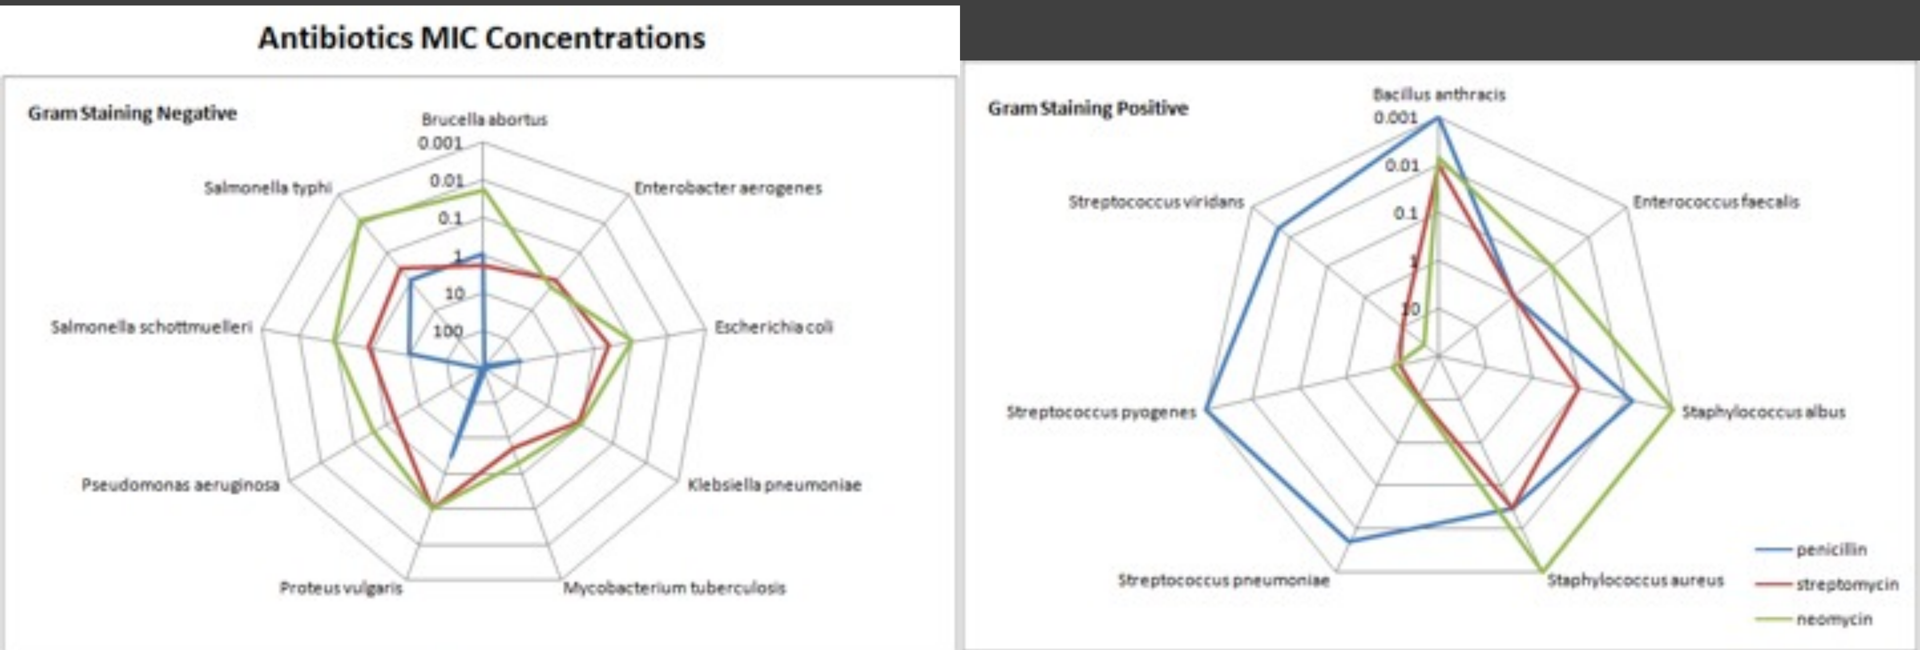

"Parallel" dimensions in polar coordinate space

Best if same units apply to each axis

Tableau / Polaris

Polaris

Research at Stanford by Stolte, Tang, and Hanrahan.

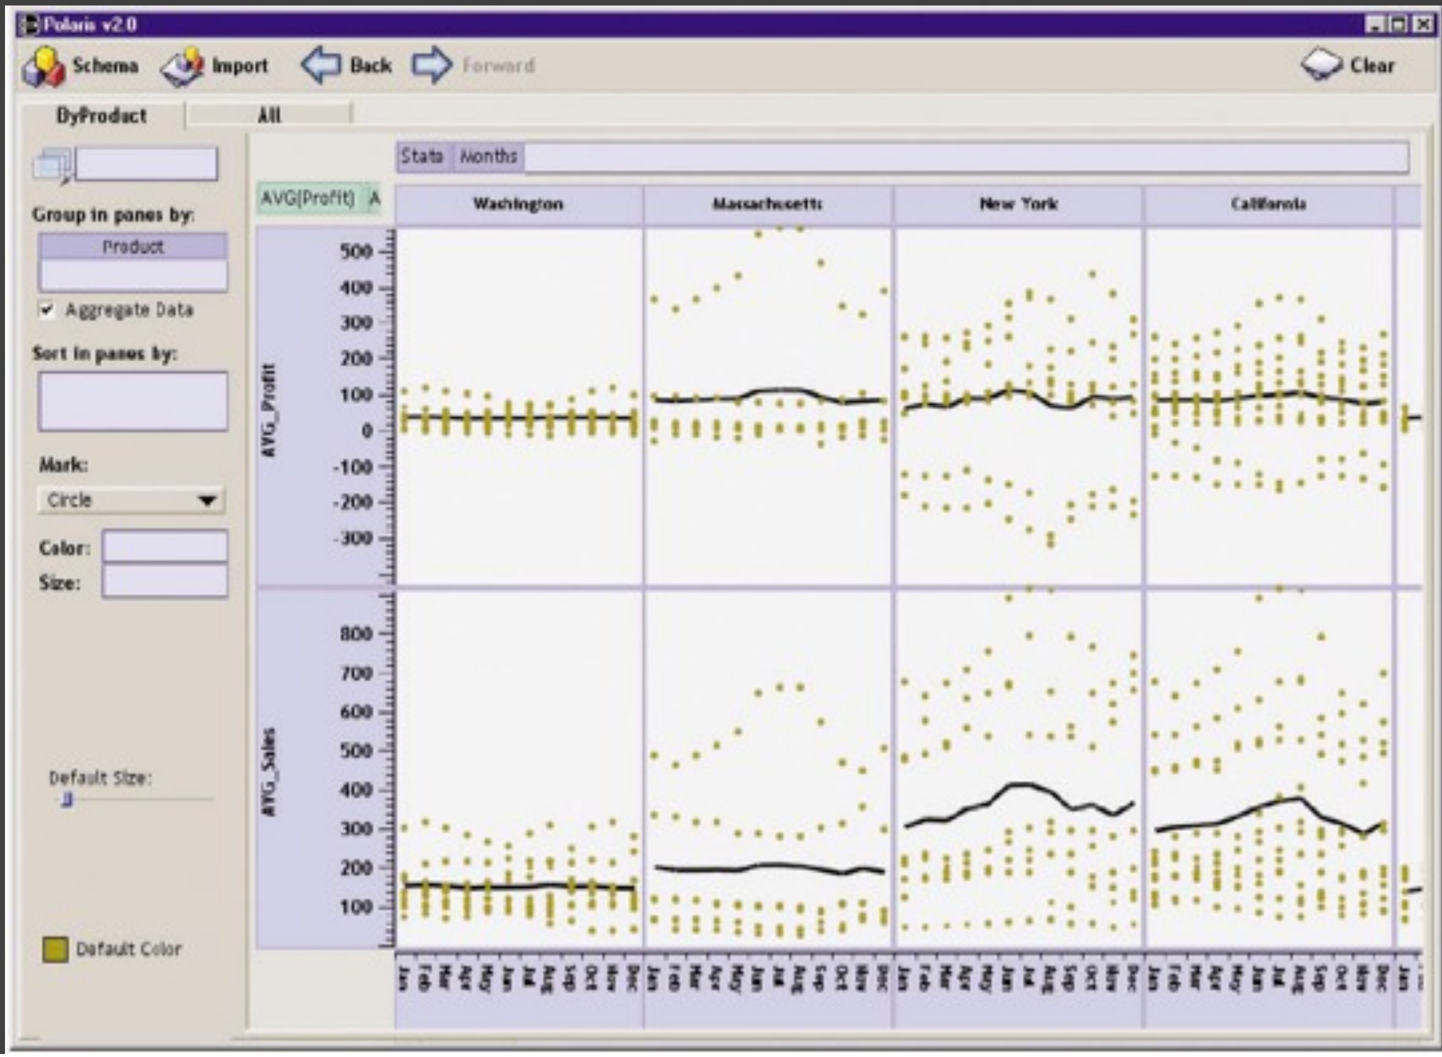

Tableau

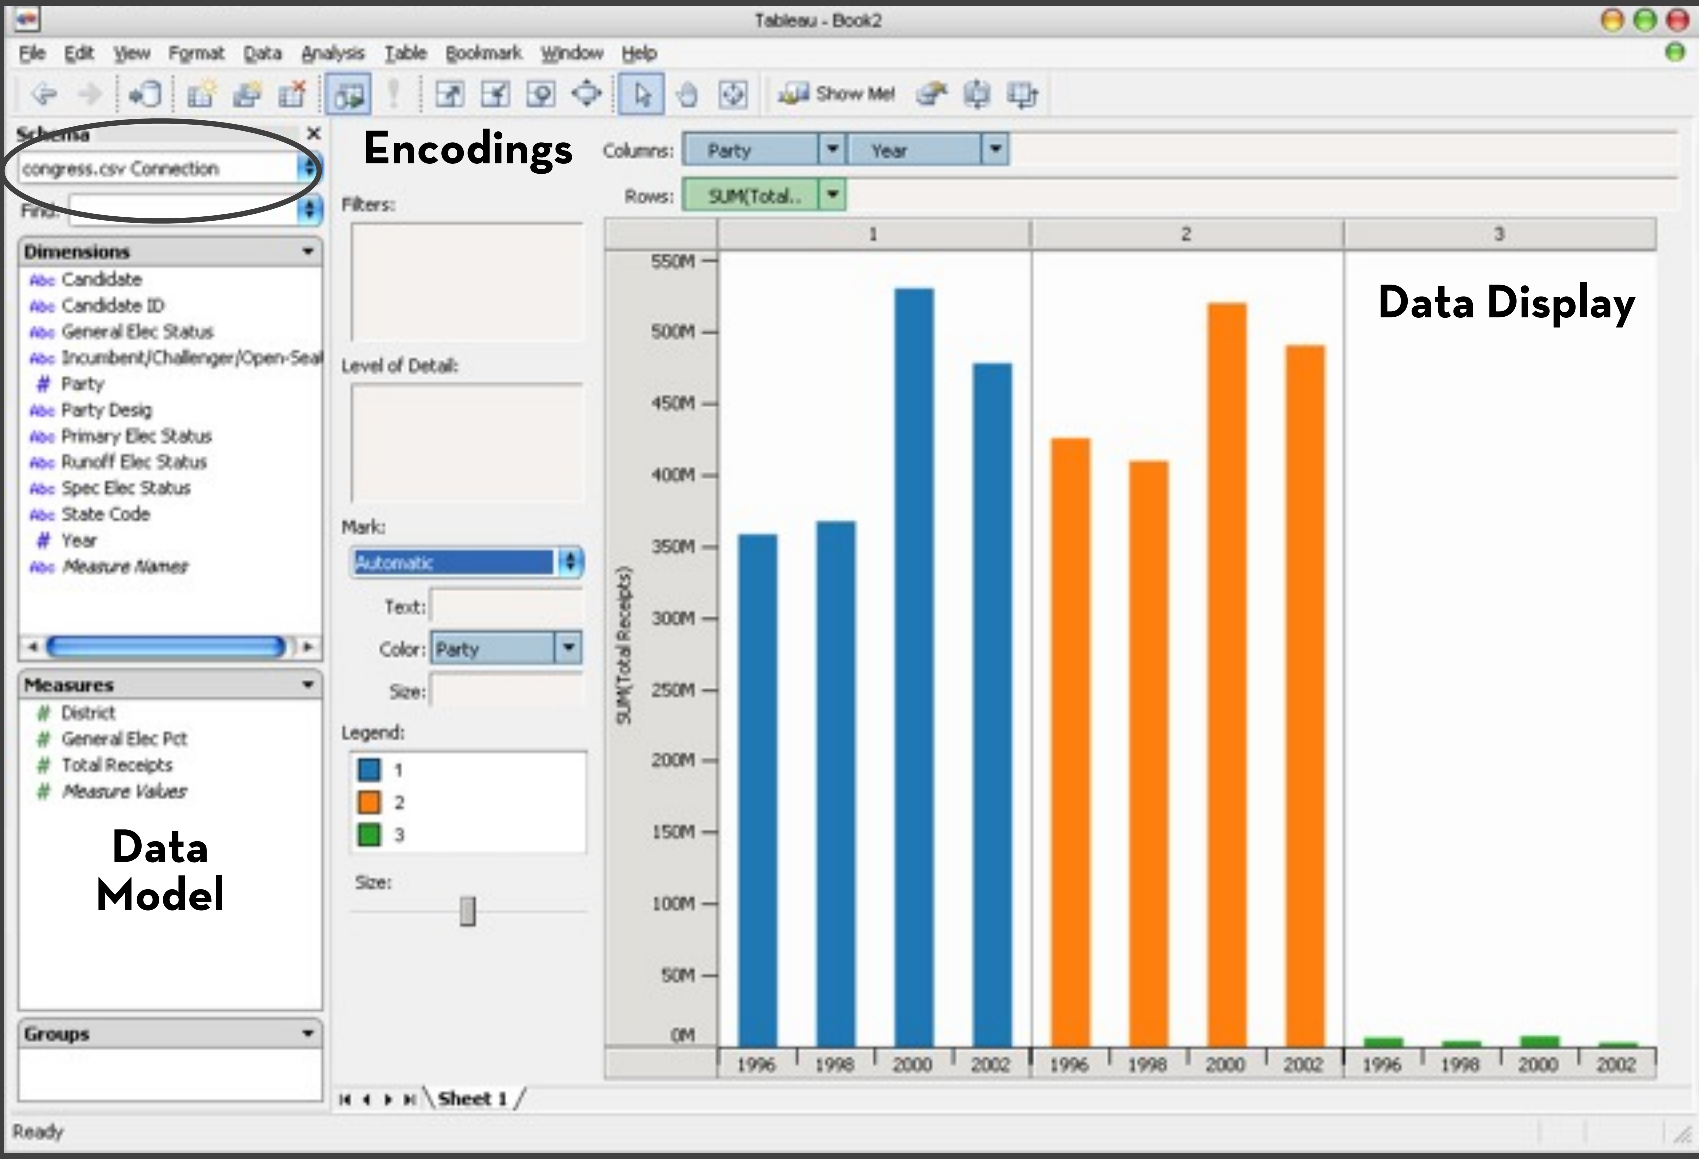
Tableau - Book2
File Edit View Format Data Analysis Table Bookmark Window Help
Show Me!
Schema
congress.csv Connection
Find:
Dimensions
Candidate
Candidate ID
General Elec Status
Incumbent/Challenger/Open-Seat
Party
Party Desig
Primary Elec Status
Runoff Elec Status
Spec Elec Status
State Code
Year
Measure Names
Measures
District
General Elec Pct
Total Receipts
Measure Values
Groups
Data Model
Encodings
Filters:
Level of Detail:
Mark:
Automatic
Text:
Color: Party
Size:
Legend:
1
2
3
Size:
Columns: Party Year
Rows: SUM(Total..
Data Display
SUM(Total Receipts)
550M
500M
450M
400M
350M
300M
250M
200M
150M
100M
50M
0M
1
2
3
1996 1998 2000 2002
Sheet 1
Ready

Tableau Demo

The dataset:

Federal Elections Commission Receipts

Every Congressional Candidate from 1996 to 2002

4 Election Cycles

9216 Candidacies

Data Set Schema

Year (Qi)

Candidate Code (N)

Candidate Name (N)

Incumbent / Challenger / Open-Seat (N)

Party Code (N) [1=Dem,2=Rep,3=Other]

Party Name (N)

Total Receipts (Qr)

State (N)

District (N)

This is a subset of the larger data set available from the FEC

Hypotheses?

What might we learn from this data?

- ??

Hypotheses?

What might we learn from this data?

Correlation between receipts and winners?

Do receipts increase over time?

Which states spend the most?

Which party spends the most?

Margin of victory vs. amount spent?

Amount spent between competitors?

Tableau Demo

Assignment 2: Exploratory Data Analysis

Use visualization software to form & answer questions.

First steps:

- Step 1: Pick domain & data
- Step 2: Pose questions
- Step 3: Profile the data
- Iterate as needed

Create visualizations

- Interact with data
- Refine your questions

Make wiki notebook

- Keep record of your analysis
- Prepare a final graphic and caption

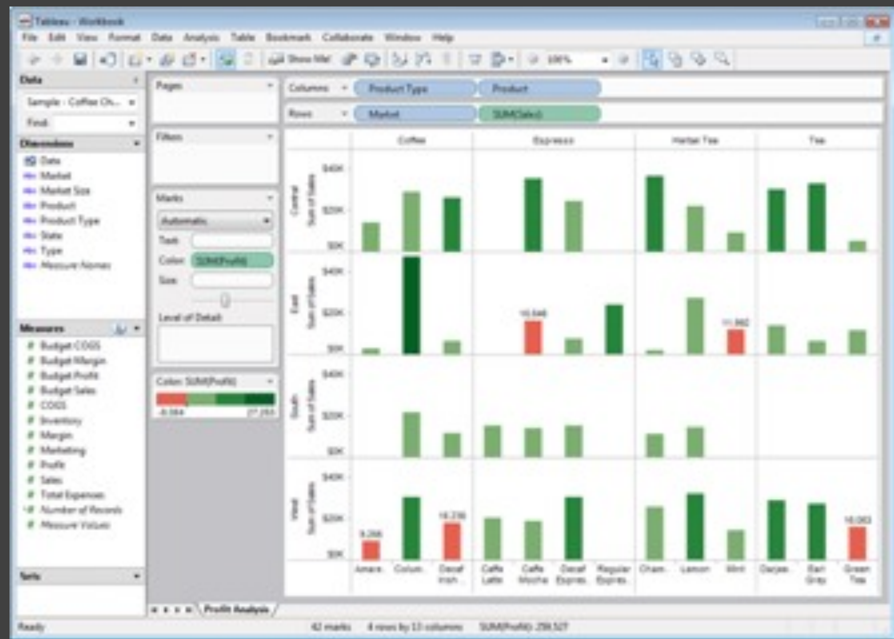

Due by 5:00pm
Monday, Jan 27

Polaris/Tableau Approach

Insight: can simultaneously specify both database queries and visualization

Choose data, then visualization, not vice versa

Use smart defaults for visual encodings

More recently: automate visualization design

Specifying Table Configurations

Operands are the database fields

- Each operand interpreted as a set {...}
- Quantitative and Ordinal fields treated differently

Three operators:

- **concatenation** (+)
- **cross product** (x)
- **nest** (/)

Table Algebra: Operands

Ordinal fields: interpret domain as a set that partitions table into rows and columns.

Quarter = {(Qtr1),(Qtr2),(Qtr3),(Qtr4)} →

Qtr1	Qtr2	Qtr3	Qtr4
95892	101760	105282	98225

Quantitative fields: treat domain as single element set and encode spatially as axes:

Profit = {(Profit[-410,650])} →

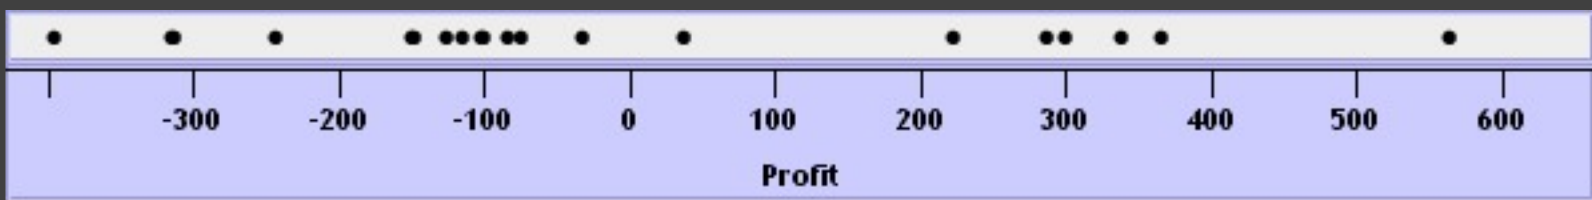

Concatenation (+) Operator

Ordered union of set interpretations

Quarter + Product Type

= {(Qtr1),(Qtr2),(Qtr3),(Qtr4)} + {(Coffee), (Espresso)}

= {(Qtr1),(Qtr2),(Qtr3),(Qtr4),(Coffee),(Espresso)}

Qtr1	Qtr2	Qtr3	Qtr4	Coffee	Espresso
48	59	57	53	151	21

Profit + Sales = {(Profit[-310,620]),(Sales[0,1000])}

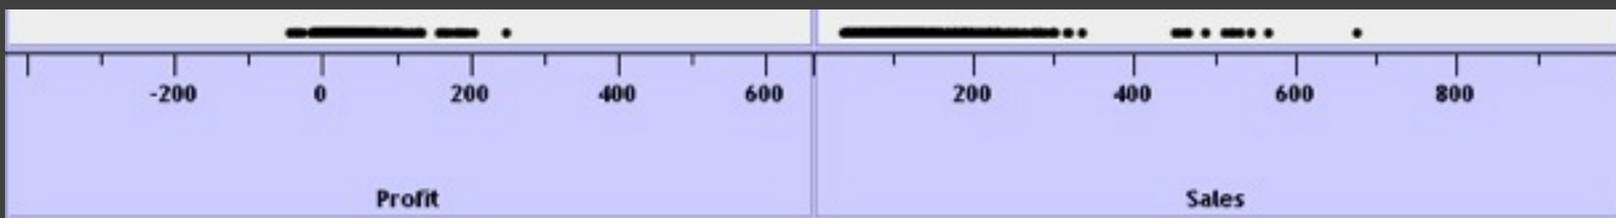

Cross (x) Operator

Cross-product of set interpretations

Quarter x Product Type

= {(Qtr1,Coffee), (Qtr1, Tea), (Qtr2, Coffee), (Qtr2, Tea), (Qtr3, Coffee), (Qtr3, Tea), (Qtr4, Coffee), (Qtr4,Tea)}

Qtr1		Qtr2		Qtr3		Qtr4	
Coffee	Espresso	Coffee	Espresso	Coffee	Espresso	Coffee	Espresso
131	19	160	20	178	12	134	33

Product Type x Profit =

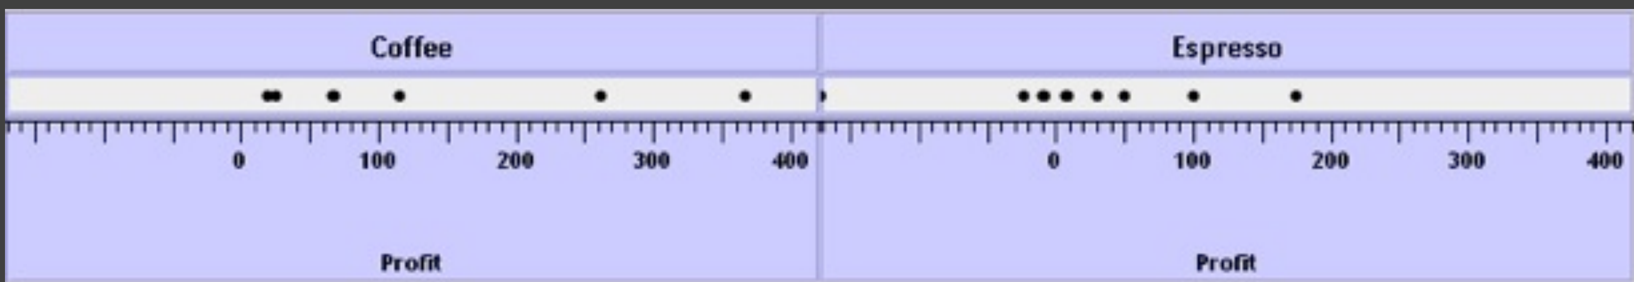

Nest (/) Operator

Cross-product filtered by existing records

Quarter x Month

- creates twelve entries for each quarter. i.e., (Qtr1, December)

Quarter / Month

- creates three entries per quarter based on tuples in database (not semantics)

Table Algebra

The operators (+, x, /) and operands (O, Q) provide an *algebra* for tabular visualization.

Algebraic statements are then mapped to:

- **Visualizations** - trellis plot partitions, visual encodings
- **Queries** - selection, projection, group-by aggregation

In Tableau, users make statements via drag-and-drop

- Note that this specifies operands NOT operators!
- Operators are inferred by data type (O, Q)

Ordinal - Ordinal

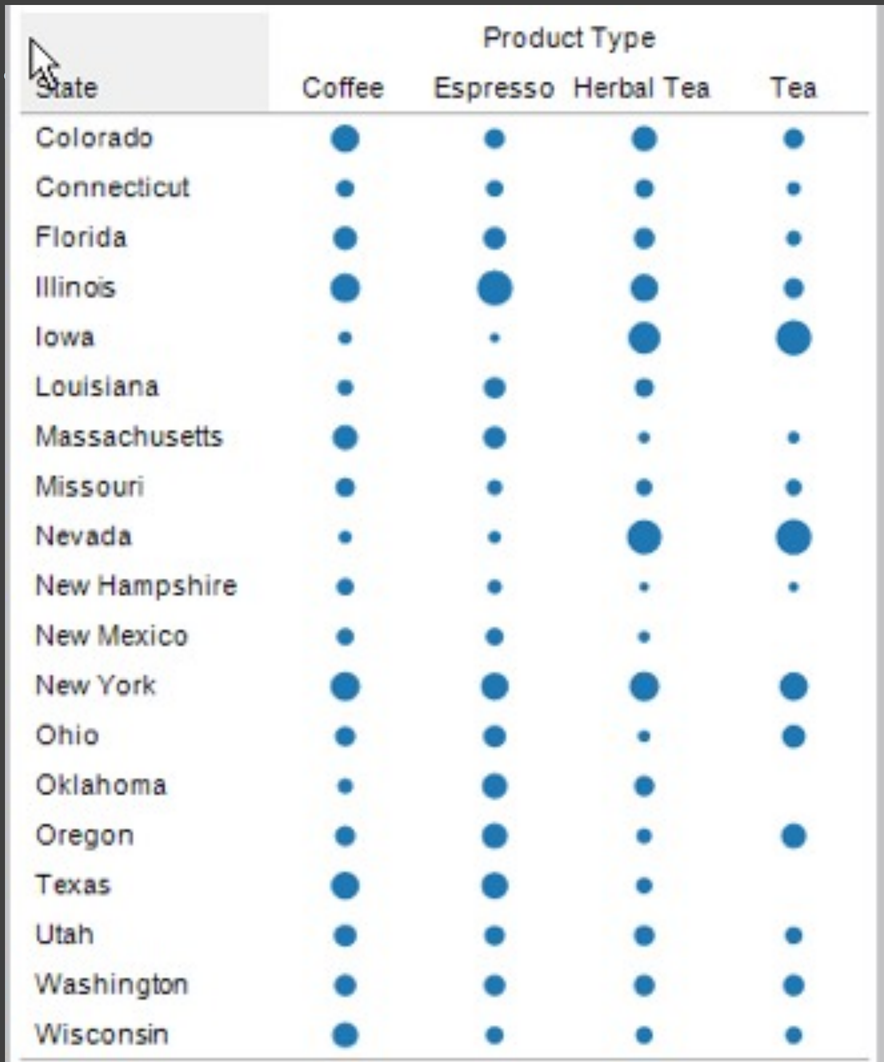

Quantitative - Quantitative

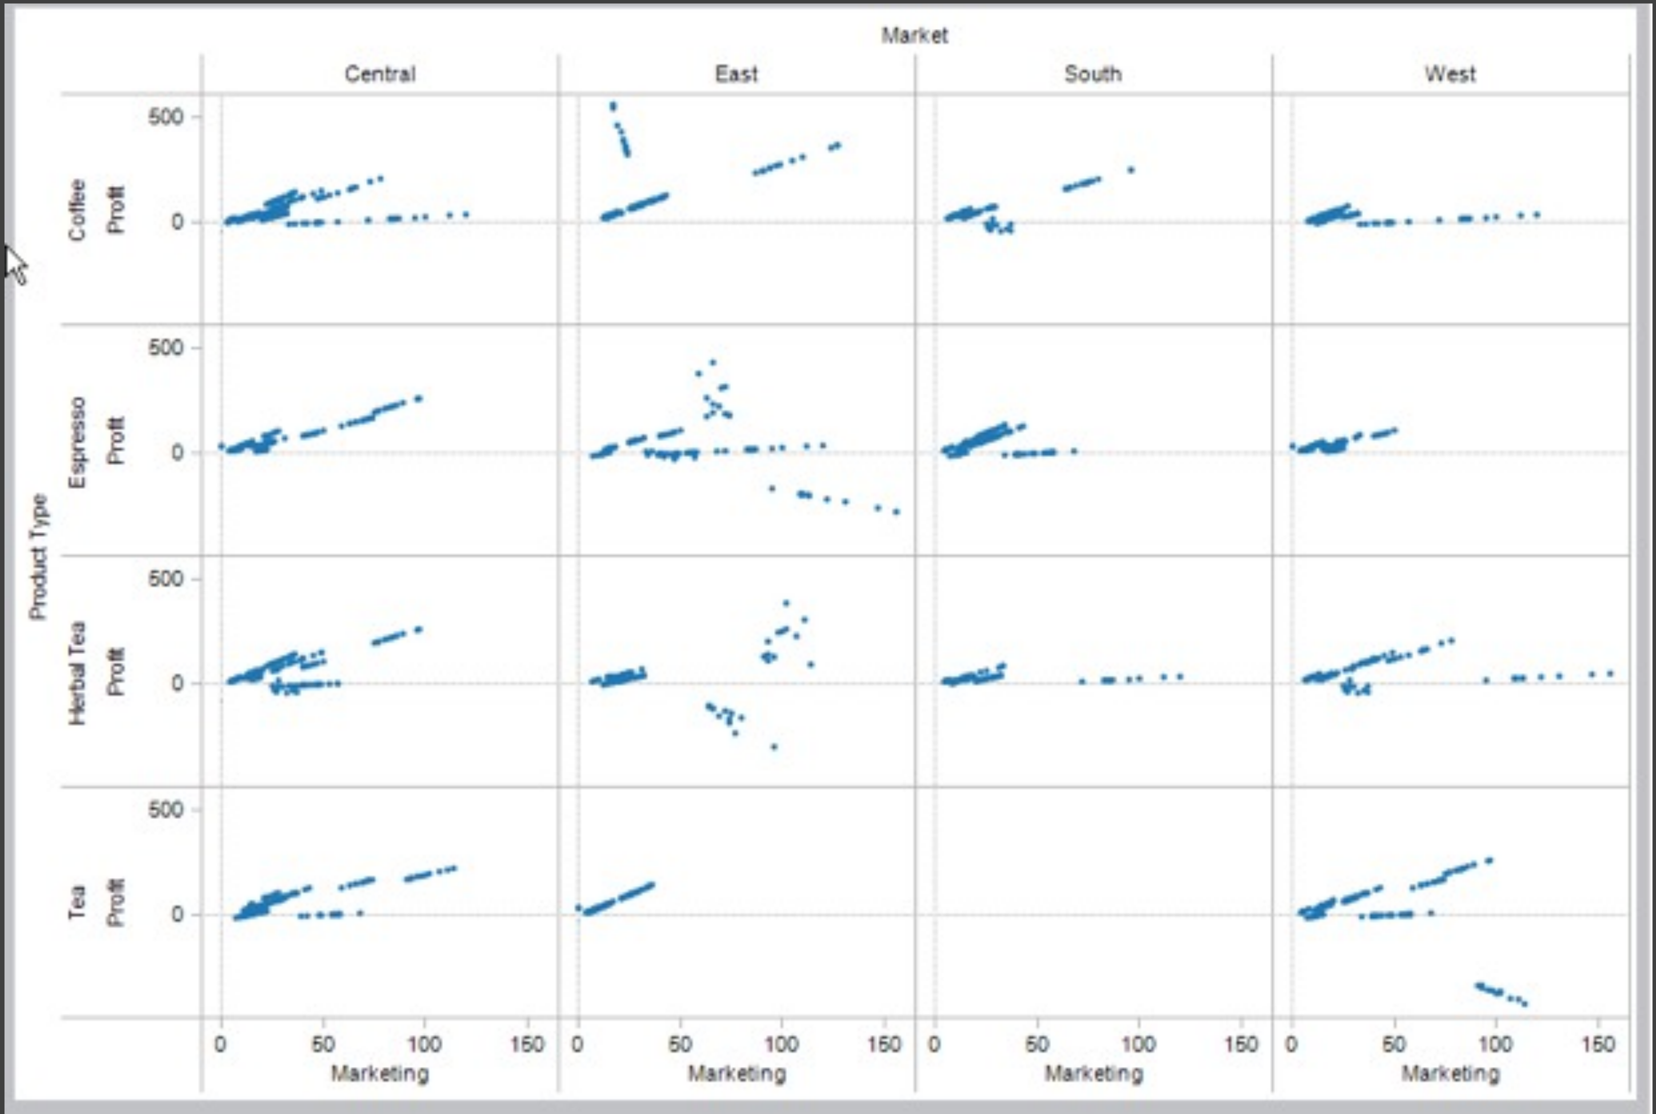

Ordinal - Quantitative

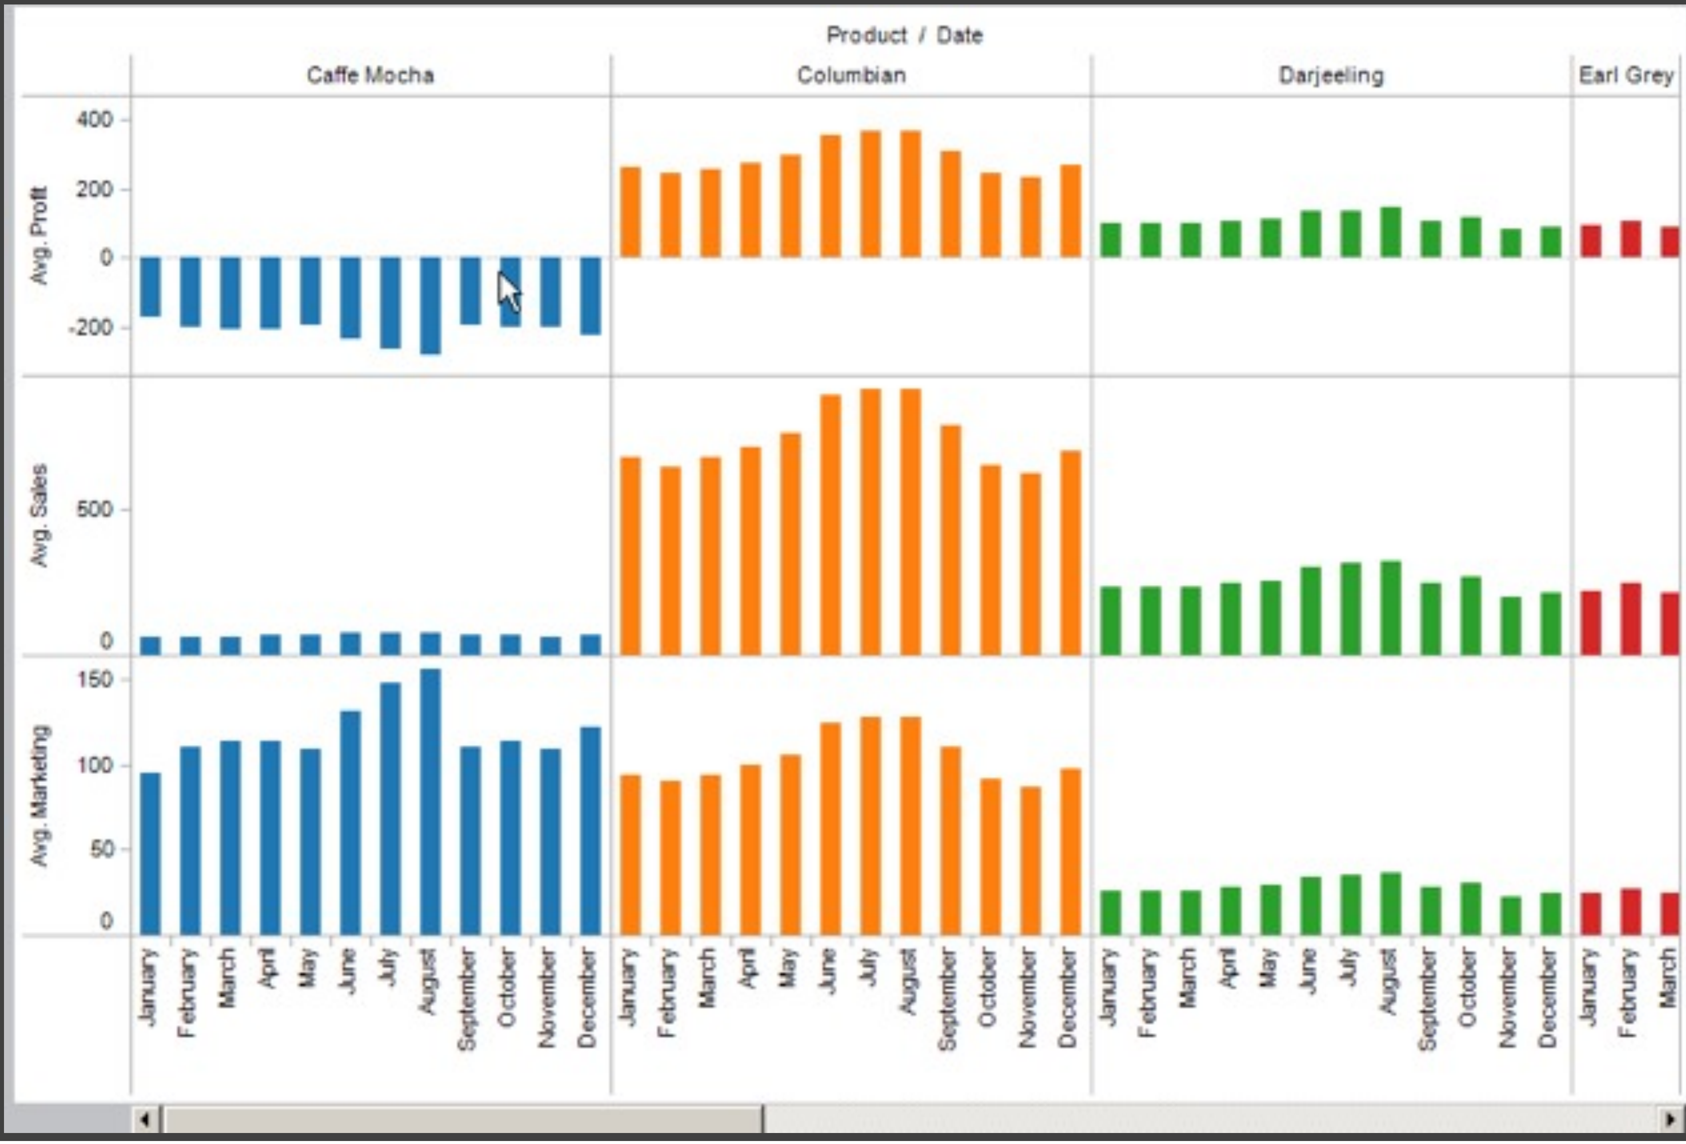

Querying the Database

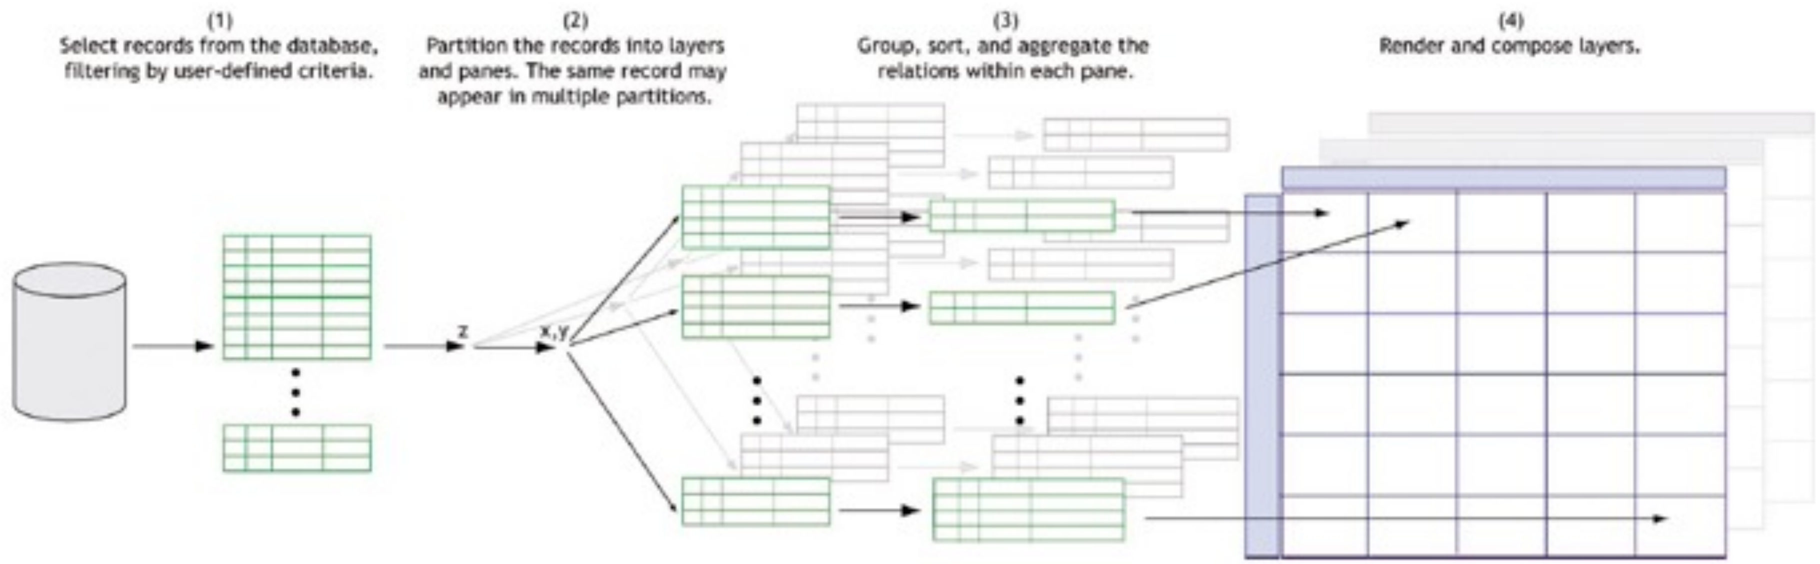

Visualizing Multiple Dimensions

Strategies

- Start by visualizing individual dimensions
- Avoid “over-encoding”
- Use space and small multiples intelligently
- Use interaction to generate *relevant* views

There is rarely a single visualization that answers all questions. Instead, the ability to generate appropriate visualizations quickly is key.