Paper: TopoFisheye
Ch 13/14/15: Reduce, Embed, Case Studies

Tamara Munzner
Department of Computer Science
University of British Columbia

CPSC 547, Information Visualization
Week 9: 5 Nov 2019

http://www.cs.ubc.ca/~tmm/courses/547-19

News

- today
  - presentations: first 5
  - *break*
  - presentations: last 2
  - topo fisheye views paper
  - chapters: reduce, embed, case studies

Paper: TopoFisheye

Topological Fisheye Views

- derived data
  - input: laid-out network (spatial positions for nodes)
  - output: multilevel hierarchy from graph coarsening
- interaction
  - user changed selected focus point
- visual encoding
  - hybrid view made from cut through several hierarchy levels

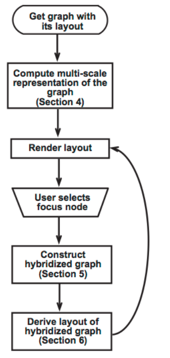

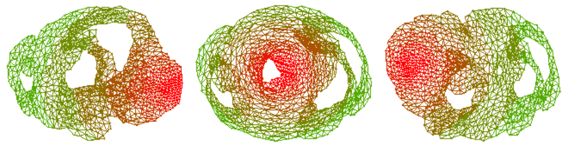

[Fig 4,8. Topological Fisheye Views for Visualizing Large Graphs. Gansner, Koren and North, IEEE TVCG 11(4), p 457-468, 2005]

Coarsening requirements

- uniform cluster/metanode size
- match coarse and fine layout geometries
- scalable

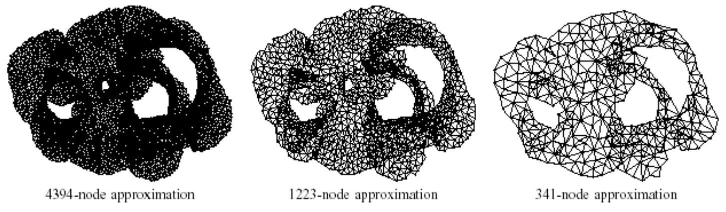

[Fig 3. Topological Fisheye Views for Visualizing Large Graphs. Gansner, Koren and North, IEEE TVCG 11(4), p 457-468, 2005]

Coarsening strategy

- must preserve graph-theoretic properties
- use both topology and geometry
  - topological distance (hops away)
  - geometric distance - but not just proximity alone!
    - just contracting nodes/edges could create new cycles
- derived data: proximity graph

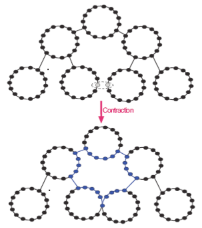

what **not** to do!

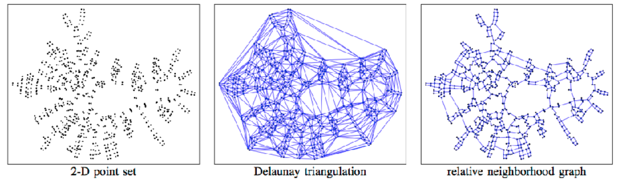

[Fig 10, 12. Topological Fisheye Views for Visualizing Large Graphs. Gansner, Koren and North, IEEE TVCG 11(4), p 457-468, 2005]

Candidate pairs: neighbors in original and proximity graph

- proximity graph: compromise between larger DT and smaller RNG
  - better than original graph neighbors alone
    - slow for cases like star graph
- maximize weighted sum of
  - geometric proximity
    - goal: preserve geometry
  - cluster size
    - goal: keep uniform cluster size
  - normalized connection strength
    - goal: preserve topology
  - neighborhood similarity
    - goal: preserve topology
  - degree
    - goal: penalize high-degree nodes to avoid salient artifacts and computational problems

Hybrid graph creation

- cut through coarsening hierarchy to get active nodes
  - animated transitions between states

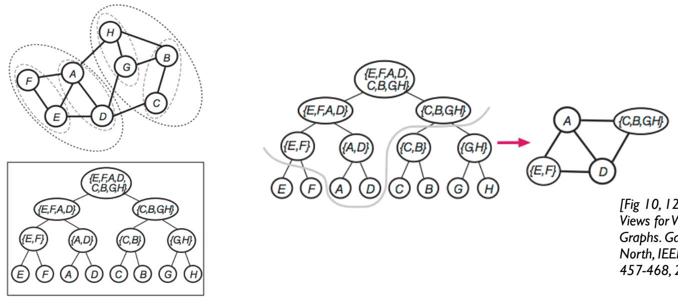

[Fig 10, 12. Topological Fisheye Views for Visualizing Large Graphs. Gansner, Koren and North, IEEE TVCG 11(4), p 457-468, 2005]

Final distortion

- geometric distortion for uniform density
- (colorcoded by hierarchy depth just to illustrate algorithm)
  - compare to original
  - compare to simple topologically unaware fisheye distortion

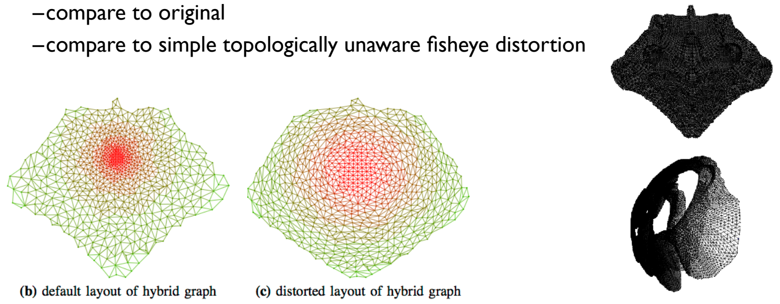

[Fig 2,15. Topological Fisheye Views for Visualizing Large Graphs. Gansner, Koren and North, IEEE TVCG 11(4), p 457-468, 2005]

Ch 13: Reduce

Reduce items and attributes

- reduce/increase: inverses
- filter
  - pro: straightforward and intuitive
    - to understand and compute
  - con: out of sight, out of mind
- aggregation
  - pro: inform about whole set
  - con: difficult to avoid losing signal
- not mutually exclusive
  - combine filter, aggregate
  - combine reduce, change, facet

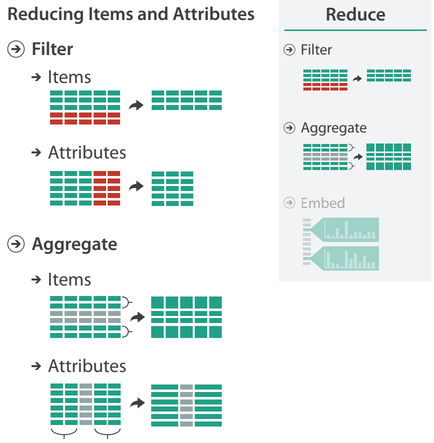

Idiom: **cross filtering** System: **Crossfilter**

- item filtering
- coordinated views/controls combined
  - all scented histogram bisliders update when any ranges change

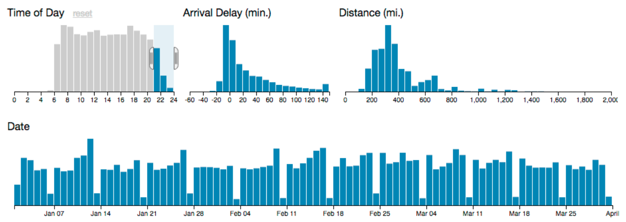

[http://square.github.io/crossfilter/]

Idiom: **cross filtering**

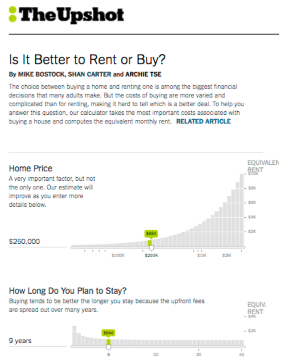

[https://www.nytimes.com/interactive/2014/upshot/buy-rent-calculator.html?_r=0]

Idiom: **histogram**

- static item aggregation
- task: find distribution
- data: table
- derived data
  - new table: keys are bins, values are counts
- bin size crucial
  - pattern can change dramatically depending on discretization
  - opportunity for interaction: control bin size on the fly

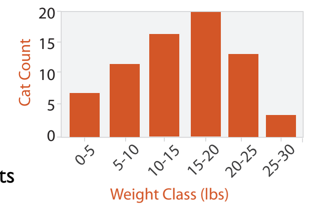

Idiom: **scented widgets**

- augmented widgets show information scent
  - cues to show whether value in drilling down further vs looking elsewhere
- concise use of space: histogram on slider

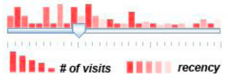

[Scented Widgets: Improving Navigation Cues with Embedded Visualizations. Willett, Heer, and Agrawala. IEEE TVCG (Proc. InfoVis 2007) 13:6 (2007), 1129–1136.]

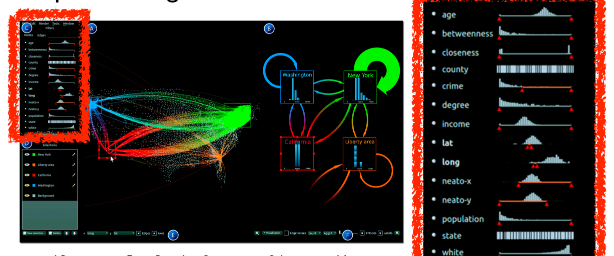

[Multivariate Network Exploration and Presentation: From Detail to Overview via Selections and Aggregations. van den Elzen, van Wijk, IEEE TVCG 20(12): 2014 (Proc. InfoVis 2014).]

Scented histogram bisliders: detailed

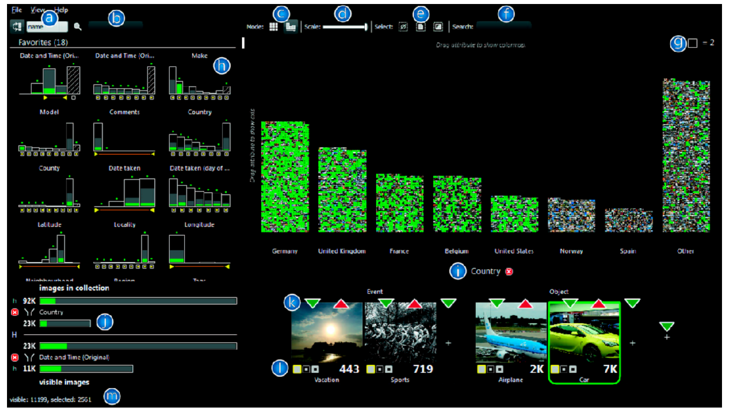

[ICLIC: Interactive categorization of large image collections. van der Corput and van Wijk. Proc. PacificVis 2016.]

Idiom: **Continuous scatterplot**

- static item aggregation
- data: table
- derived data: table
  - key attribs x,y for pixels
  - quant attrib: overplot density
- dense space-filling 2D matrix
- color: sequential categorical hue + ordered luminance colormap

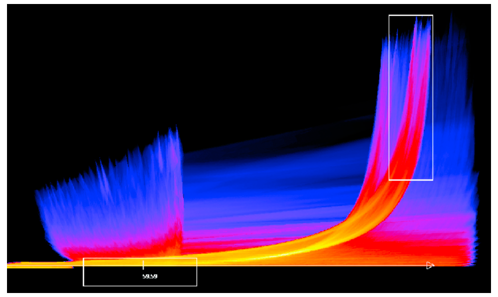

[Continuous Scatterplots. Bachthaler and Weiskopf. IEEE TVCG (Proc. Vis 08) 14:6 (2008), 1428–1435. 2008.]

Spatial aggregation

- MAUP: Modifiable Areal Unit Problem
  - gerrymandering (manipulating voting district boundaries) is only one example!
  - zone effects

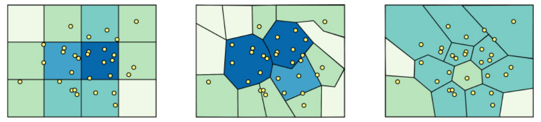

[http://www.e-education.psu.edu/geog486/l4_p7.html, Fig 4.cg.6]

  - scale effects

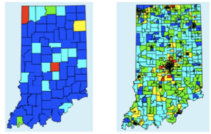

https://blog.cartographica.com/blog/2011/5/19/the-modifiable-areal-unit-problem-in-gis.html

Idiom: **boxplot**

- static item aggregation
- task: find distribution
- data: table
- derived data
  - 5 quant attribs
    - median: central line
    - lower and upper quartile: boxes
    - lower upper fences: whiskers
      - values beyond which items are outliers
  - outliers beyond fence cutoffs explicitly shown

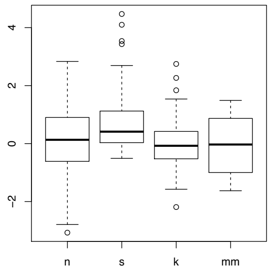

[40 years of boxplots. Wickham and Stryjewski. 2012. had.co.nz]

Idiom: **Hierarchical parallel coordinates**

- dynamic item aggregation
- derived data: ***hierarchical clustering***
- encoding:
  - cluster band with variable transparency, line at mean, width by min/max values
  - color by proximity in hierarchy

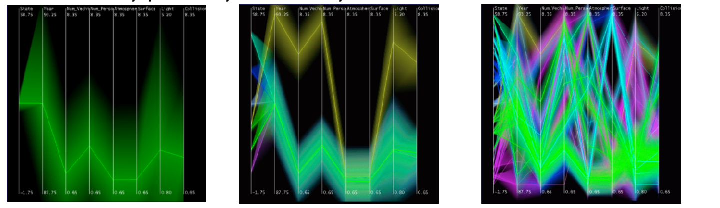

[Hierarchical Parallel Coordinates for Exploration of Large Datasets. Fua, Ward, and Rundensteiner. Proc. IEEE Visualization Conference (Vis '99), pp. 43– 50, 1999.]

Idiom: aggregation via hierarchical clustering (visible)

System: **Hierarchical Clustering Explorer**

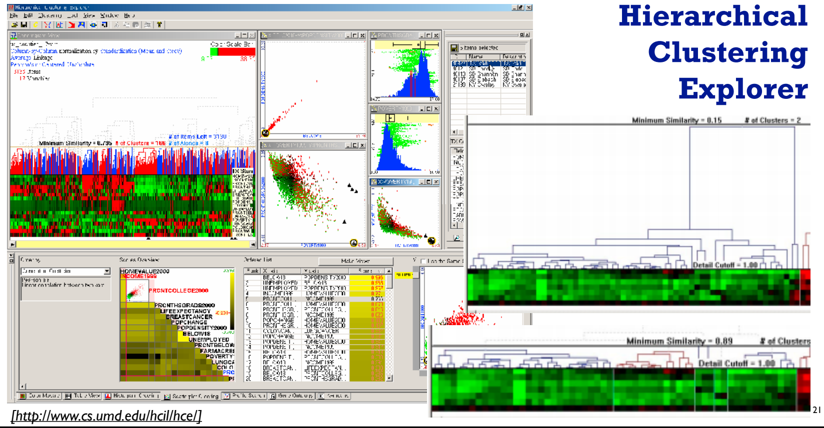

[http://www.cs.umd.edu/hcil/hce/]

Dimensionality reduction

- attribute aggregation
  - derive low-dimensional target space from high-dimensional measured space
    - capture most of variance with minimal error
  - use when you can't directly measure what you care about
    - true dimensionality of dataset conjectured to be smaller than dimensionality of measurements
    - latent factors, hidden variables

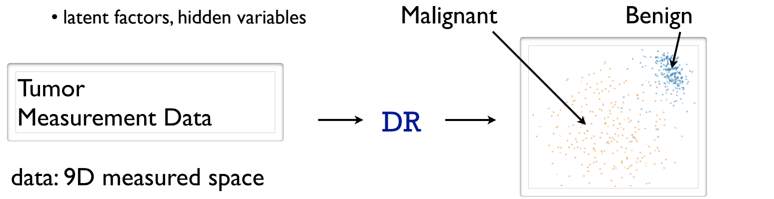

data: 9D measured space

derived data: 2D target space

Dimensionality vs attribute reduction

- vocab use in field not consistent
  - dimension/attribute
- attribute reduction: reduce set with filtering
  - includes orthographic projection
- dimensionality reduction: create smaller set of new dims/attribs
  - typically implies dimensional aggregation, not just filtering
  - vocab: projection/mapping

Dimensionality reduction & visualization

- why do people do DR?
  - improve performance of downstream algorithm
    - avoid curse of dimensionality
  - data analysis
    - if look at the output: visual data analysis
- abstract tasks when visualizing DR data
  - dimension-oriented tasks
    - naming synthesized dims, mapping synthesized dims to original dims
  - cluster-oriented tasks
    - verifying clusters, naming clusters, matching clusters and classes

[Visualizing Dimensionally-Reduced Data: Interviews with Analysts and a Characterization of Task Sequences. Brehmer, Sedlmair, Ingram, and Munzner. Proc. BELIV 2014.]

Dimension-oriented tasks

- naming synthesized dims: inspect data represented by lowD points

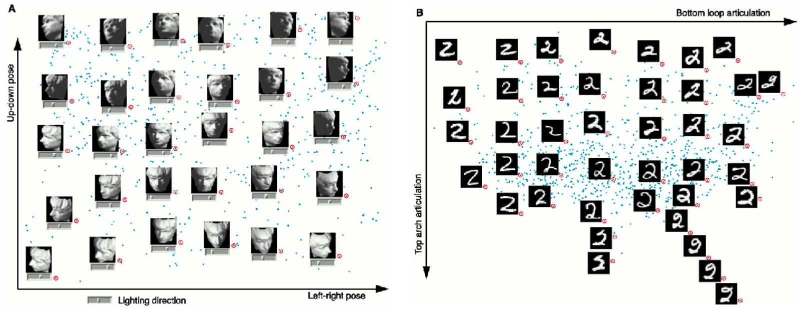

[A global geometric framework for nonlinear dimensionality reduction. Tenenbaum, de Silva, and Langford. Science, 290(5500):2319–2323, 2000.]

Cluster-oriented tasks

- verifying, naming, matching to classes

no discernable clusters	clearly discernable clusters	clear match cluster/class	partial match cluster/class	no match cluster/class

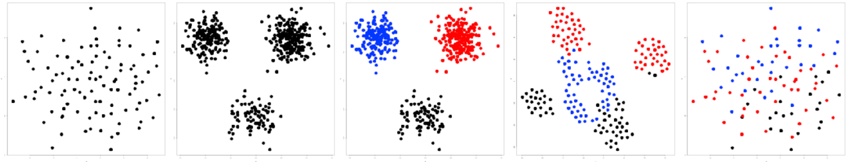

[Visualizing Dimensionally-Reduced Data: Interviews with Analysts and a Characterization of Task Sequences. Brehmer, Sedlmair, Ingram, and Munzner. Proc. BELIV 2014.]

Idiom: **Dimensionality reduction for documents**

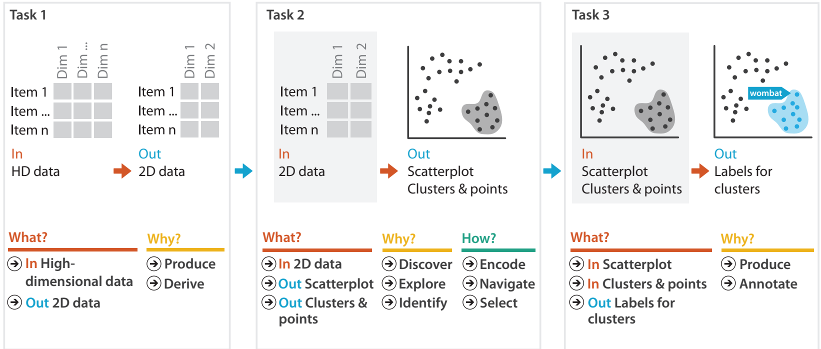

Interacting with dimensionally reduced data

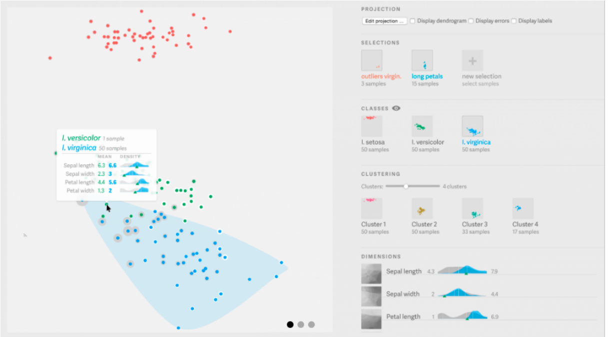

[https://uclab.fh-potsdam.de/projects/probing-projections/]

[Probing Projections: Interaction Techniques for Interpreting Arrangements and Errors of Dimensionality Reductions. Stahnke, Dörk, Müller, and Thom. IEEE TVCG (Proc. InfoVis 2015) 22(1):629-38 2016.]

Linear dimensionality reduction

- principal components analysis (PCA)
  - finding axes: first with most variance, second with next most, …
  - describe location of each point as linear combination of weights for each axis
    - mapping synthesized dims to original dims

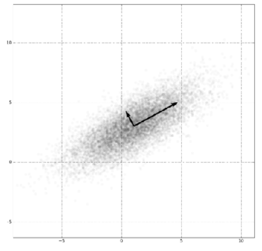

[http://en.wikipedia.org/wiki/File:GaussianScatterPCA.png]

Nonlinear dimensionality reduction

- pro: can handle curved rather than linear structure
- cons: lose all ties to original dims/attribs
  - new dimensions often cannot be easily related to originals
    - mapping synthesized dims to original dims task is difficult
- many techniques proposed
  - many literatures: visualization, machine learning, optimization, psychology, ...
  - techniques: t-SNE, MDS (multidimensional scaling), charting, isomap, LLE,…
    - t-SNE: excellent for clusters
      - but some trickiness remains: http://distill.pub/2016/misread-tsne/
    - MDS: confusingly, entire family of techniques, both linear and nonlinear
      - minimize stress or strain metrics
      - early formulations equivalent to PCA

VDA with DR example: nonlinear vs linear

- DR for computer graphics reflectance model
  - goal: simulate how light bounces off materials to make realistic pictures
    - computer graphics: BRDF (reflectance)
  - idea: measure what light does with real materials

[Fig 2. Matusik, Pfister, Brand, and McMillan. A Data-Driven Reflectance Model. SIGGRAPH 2003]

Capturing & using material reflectance

- reflectance measurement: interaction of light with real materials (spheres)
- result: 104 high-res images of material
  - each image 4M pixels
- goal: image synthesis
  - simulate completely new materials
- need for more concise model
  - 104 materials * 4M pixels = 400M dims
  - want concise model with meaningful knobs
    - how shiny/greasy/metallic
    - DR to the rescue!

[Figs 5/6. Matusik et al. A Data-Driven Reflectance Model. SIGGRAPH 2003]

Linear DR

- first try: PCA (linear)
- result: error falls off sharply after ~45 dimensions
  - scree plots: error vs number of dimensions in lowD projection
- problem: physically impossible intermediate points when simulating new materials
  - specular highlights cannot have holes!

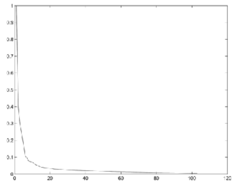

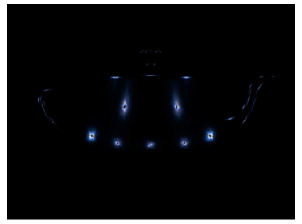

[Figs 6/7. Matusik et al. A Data-Driven Reflectance Model. SIGGRAPH 2003]

Nonlinear DR

- second try: charting (nonlinear DR technique)
  - scree plot suggests 10-15 dims
  - note: dim estimate depends on technique used!

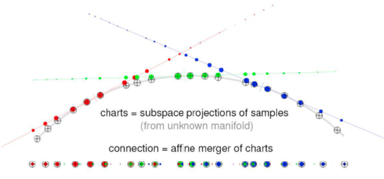

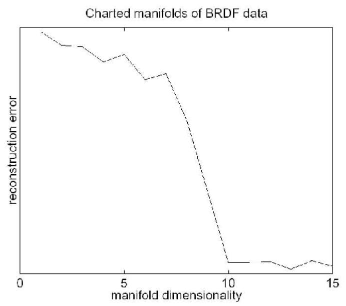

[Fig 10/11. Matusik et al. A Data-Driven Reflectance Model. SIGGRAPH 2003]

Finding semantics for synthetic dimensions

- look for meaning in scatterplots
  - synthetic dims created by algorithm but named by human analysts
  - points represent real-world images (spheres)
  - people inspect images corresponding to points to decide if axis could have meaningful name
- cross-check meaning
  - arrows show simulated images (teapots) made from model
  - check if those match dimension semantics

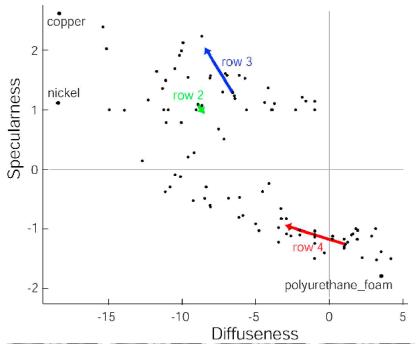

[Fig 12/16. Matusik et al. A Data-Driven Reflectance Model. SIGGRAPH 2003]

Understanding synthetic dimensions

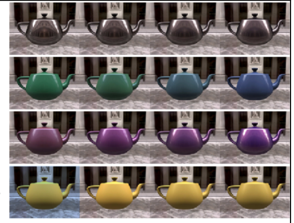

Specular-Metallic

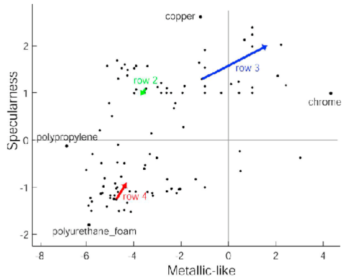

Diffuseness-Glossiness

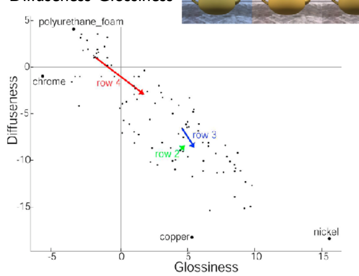

[Fig 13/14/16. Matusik et al. A Data-Driven Reflectance Model. SIGGRAPH 2003]

Ch 14: Embed

Embed: Focus+Context

- combine information within single view
- elide
  - selectively filter and aggregate
- superimpose layer
  - local lens
- distortion design choices
  - region shape: radial, rectilinear, complex
  - how many regions: one, many
  - region extent: local, global
  - interaction metaphor

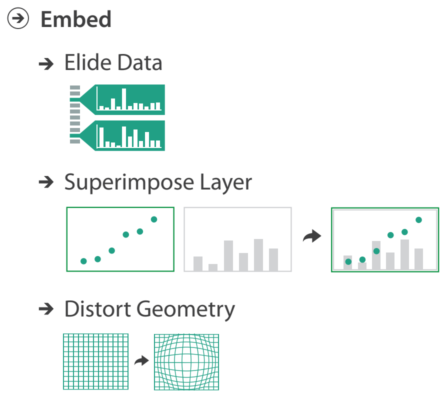

Idiom: **DOITrees Revisited**

- elide
  - some items dynamically filtered out
  - some items dynamically aggregated together
  - some items shown in detail

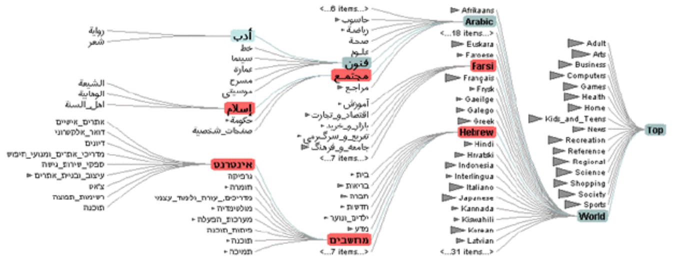

[DOITrees Revisited: Scalable, Space-Constrained Visualization of Hierarchical Data. Heer and Card. Proc. Advanced Visual Interfaces (AVI), pp. 421–424, 2004.]

Idiom: **Fisheye Lens**

- distort geometry
  - shape: radial
  - focus: single extent
  - extent: local
  - metaphor: draggable lens

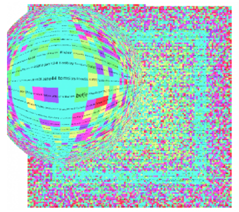

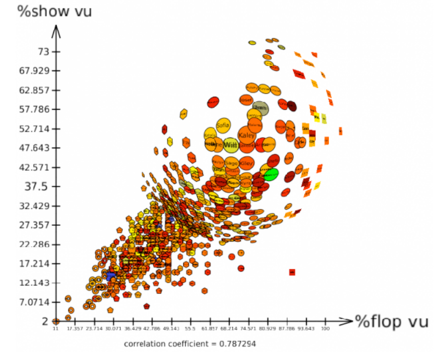

http://tulip.labri.fr/TulipDrupal/?q=node/351
http://tulip.labri.fr/TulipDrupal/?q=node/371

Idiom: **Fisheye Lens**

System: **D3**

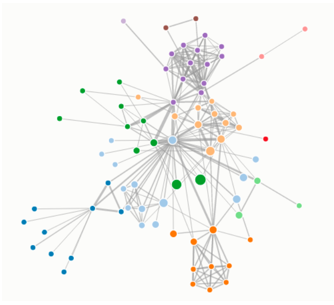

[D3 Fisheye Lens](https://bost.ocks.org/mike/fisheye/)

Idiom: **Stretch and Squish Navigation**

- distort geometry
  - shape: rectilinear
  - foci: multiple
  - impact: global
  - metaphor: stretch and squish, borders fixed

System: **TreeJuxtaposer**

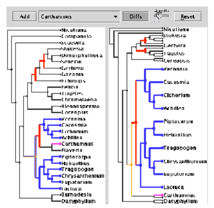

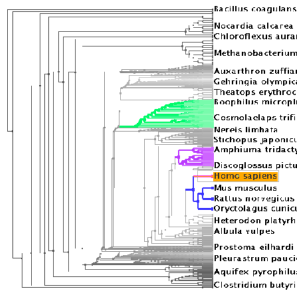

[TreeJuxtaposer: Scalable Tree Comparison Using Focus+Context With Guaranteed Visibility. Munzner, Guimbretiere, Tasiran, Zhang, and Zhou. ACM Transactions on Graphics (Proc. SIGGRAPH) 22:3 (2003), 453– 462.]

Distortion costs and benefits

- benefits
  - combine focus and context information in single view
- costs
  - length comparisons impaired
    - network/tree topology comparisons unaffected: connection, containment
  - effects of distortion unclear if original structure unfamiliar
  - object constancy/tracking maybe impaired

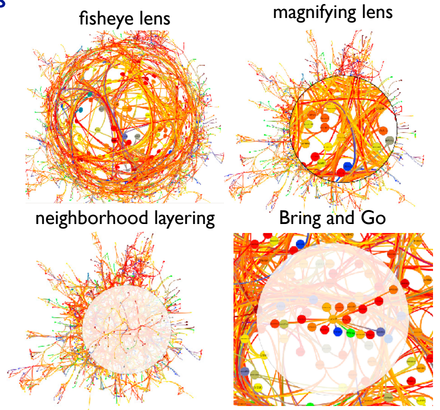

[Living Flows: Enhanced Exploration of Edge-Bundled Graphs Based on GPU-Intensive Edge Rendering. Lambert, Auber, and Melançon. Proc. Intl. Conf. Information Visualisation (IV), pp. 523–530, 2010.]

Ch 15: Case Studies

Analysis Case Studies

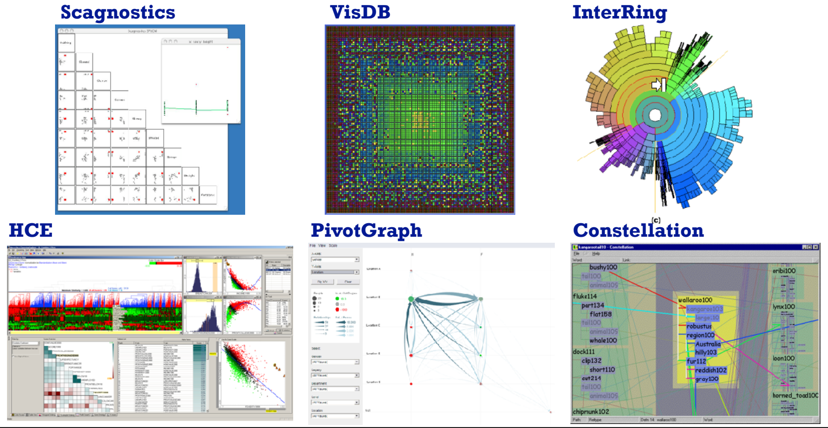

Graph-Theoretic Scagnostics

- scatterplot diagnostics
  - scagnostics SPLOM: each point is one original scatterplot

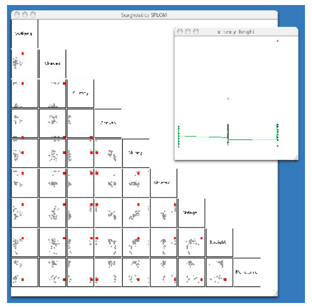

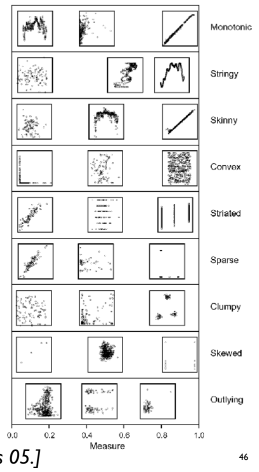

[Graph-Theoretic Scagnostics Wilkinson, Anand, and Grossman. Proc InfoVis 05.]

Scagnostics analysis

System	Scagnostics
What: Data	Table.
What: Derived	Nine quantitative attributes per scatterplot (pairwise combination of original attributes).
Why: Tasks	Identify, compare, and summarize; distributions and correlation.
How: Encode	Scatterplot, scatterplot matrix.
How: Manipulate	Select.
How: Facet	Juxtaposed small-multiple views coordinated with linked highlighting, popup detail view.
Scale	Original attributes: dozens.

VisDB

- table: draw pixels sorted, colored by relevance
- group by attribute or partition by attribute into multiple views

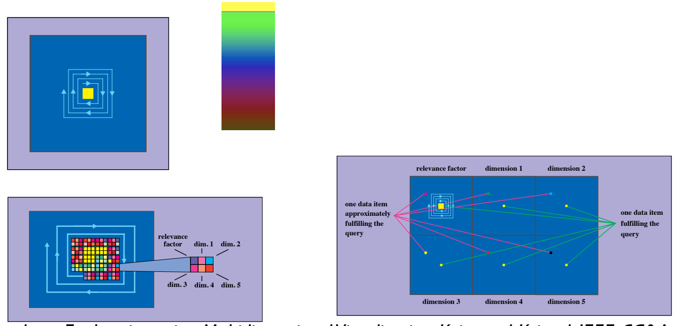

[VisDB: Database Exploration using Multidimensional Visualization, Keim and Kriegel, IEEE CG&A, 1994]

VisDB Results

- partition into many small regions: dimensions grouped together

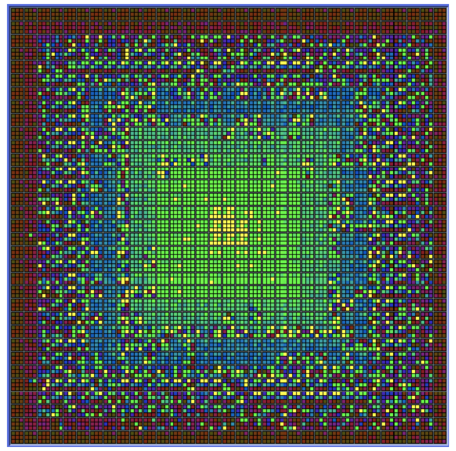

[VisDB: Database Exploration using Multidimensional Visualization, Keim and Kriegel, IEEE CG&A, 1994]

VisDB Results

- partition into small number of views
  - inspect each attribute

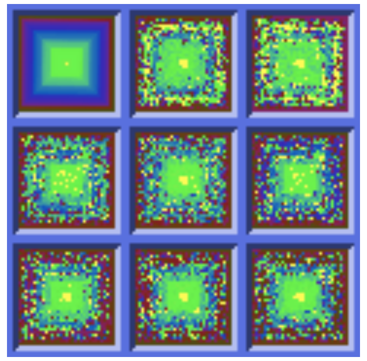

[VisDB: Database Exploration using Multidimensional Visualization, Keim and Kriegel, IEEE CG&A, 1994]

VisDB Analysis

System	VisDB
What: Data	Table (database) with k attributes; query returning table subset (database query).
What: Derived	$k+1$ quantitative attributes per original item: query relevance for the k original attributes plus overall relevance.
Why: Tasks	Characterize distribution within attribute, find groups of similar values within attribute, find outliers within attribute, find correlation between attributes, find similar items.
How: Encode	Dense, space-filling; area marks in spiral layout; colormap: categorical hues and ordered luminance.
How: Facet	Layout 1: partition by attribute into per-attribute views, small multiples. Layout 2: partition by items into per-item glyphs.
How: Reduce	Filtering
Scale	Attributes: one dozen. Total items: several million. Visible items (using multiple views, in total): one million. Visible items (using glyphs): 100,000

Hierarchical Clustering Explorer

- heatmap, dendrogram
- multiple views

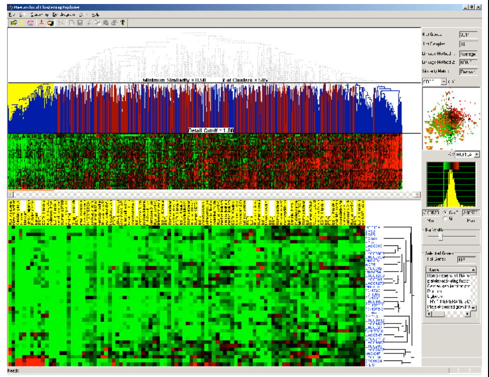

[Interactively Exploring Hierarchical Clustering Results. Seo and Shneiderman, IEEE Computer 35(7): 80-86 (2002)]

HCE

- rank by feature idiom
  - 1D list
  - 2D matrix

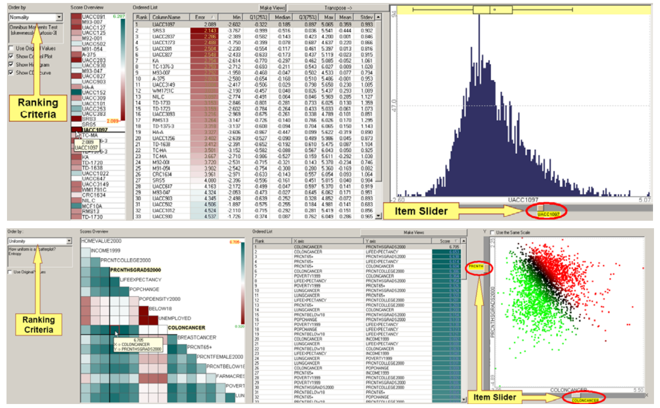

A rank-by-feature framework for interactive exploration of multidimensional data. Seo and Shneiderman. Information Visualization 4(2): 96-113 (2005)

HCE

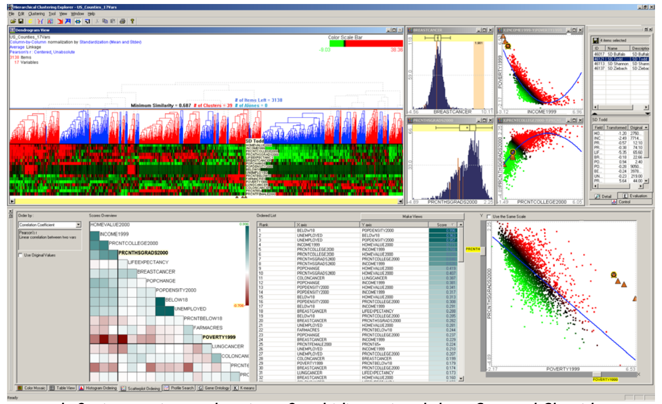

A rank-by-feature framework for interactive exploration of multidimensional data. Seo and Shneiderman. Information Visualization 4(2): 96-113 (2005)

HCE Analysis

System	Hierarchical Clustering Explorer (HCE)
What: Data	Multidimensional table: two categorical key attributes (genes, conditions); one quantitative value attribute (gene activity level in condition).
What: Derived	Hierarchical clustering of table rows and columns (for cluster heatmap); quantitative derived attributes for each attribute and pairwise attribute combination; quantitative derived attribute for each ranking criterion and original attribute combination.
Why: Tasks	Find correlation between attributes; find clusters, gaps, outliers, trends within items.
How: Encode	Cluster heatmap, scatterplots, histograms, boxplots. Rank-by-feature overviews: continuous diverging colormaps on area marks in reorderable 2D matrix or 1D list alignment.
How: Reduce	Dynamic filtering; dynamic aggregation.
How: Manipulate	Navigate with pan/scroll.
How: Facet	Multiform with linked highlighting and shared spatial position; overview–detail with selection in overview populating detail view.
Scale	Genes (key attribute): 20,000. Conditions (key attribute): 80. Gene activity in condition (quantitative value attribute): 20,000 × 80 = 1,600,000.

InterRing

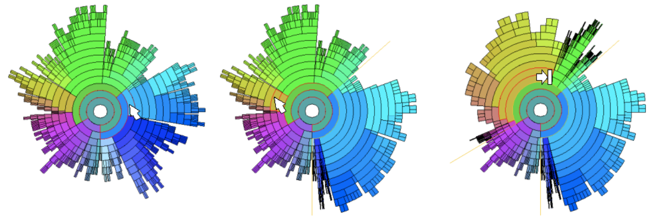

original hierarchy — blue subtree expanded — tan subtree expanded

[InterRing: An Interactive Tool for Visually Navigating and Manipulating Hierarchical Structures. Yang, Ward, Rundensteiner. Proc. InfoVis 2002, p 77-84.]

InterRing Analysis

System	InterRing
What: Data	Tree.
Why: Tasks	Selection, rollup/drilldown, hierarchy editing.
How: Encode	Radial, space-filling layout. Color by tree structure.
How: Facet	Linked coloring and highlighting.
How: Reduce	Embed: distort; multiple foci.
Scale	Nodes: hundreds if labeled, thousands if dense. Levels in tree: dozens.

PivotGraph

- derived rollup network

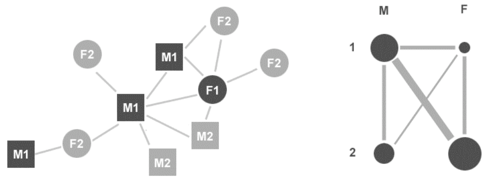

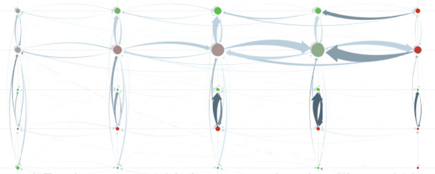

[Visual Exploration of Multivariate Graphs, Martin Wattenberg, CHI 2006.]

PivotGraph

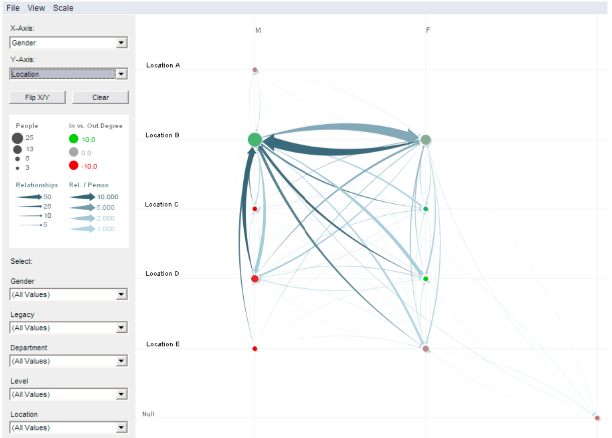

[Visual Exploration of Multivariate Graphs, Martin Wattenberg, CHI 2006.]

PivotGraph Analysis

Idiom	PivotGraph
What: Data	Network.
What: Derived	Derived network of aggregate nodes and links by roll-up into two chosen attributes.
Why: Tasks	Cross-attribute comparison of node groups.
How: Encode	Nodes linked with connection marks, size.
How: Manipulate	Change: animated transitions.
How: Reduce	Aggregation, filtering.
Scale	Nodes/links in original network: unlimited. Roll-up attributes: 2. Levels per roll-up attribute: several, up to one dozen.

Analysis example: Constellation

- data
  - multi-level network
    - node: word
    - link: words used in same dictionary definition
    - subgraph for each definition
      - not just hierarchical clustering
  - paths through network
    - query for high-weight paths between 2 nodes
      - quant attrib: plausibility

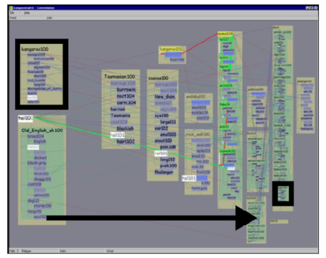

[Interactive Visualization of Large Graphs and Networks. Munzner. Ph.D. Dissertation, Stanford University, June 2000.]

[Constellation: A Visualization Tool For Linguistic Queries from MindNet. Munzner, Guimbretière and Robertson. Proc. IEEE Symp. InfoVis 1999, p.132-135.]

Using space: Constellation

- visual encoding
  - link connection marks between words
  - link containment marks to indicate subgraphs
  - encode plausibility with horiz spatial position
  - encode source/sink for query with vert spatial position
- spatial layout
  - curvilinear grid: more room for longer low-plausibility paths

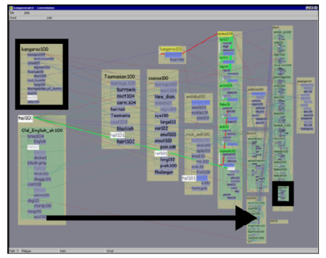

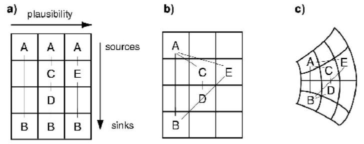

[Interactive Visualization of Large Graphs and Networks. Munzner. Ph.D. Dissertation, Stanford University, June 2000.]

Using space: Constellation

- edge crossings
  - cannot easily minimize instances, since position constrained by spatial encoding
  - instead: minimize perceptual impact
- views: superimposed layers
  - dynamic foreground/background layers on mouseover, using color
  - four kinds of constellations
    - definition, path, link type, word
      - not just 1-hop neighbors

https://youtu.be/7sJC3QVpSkQ

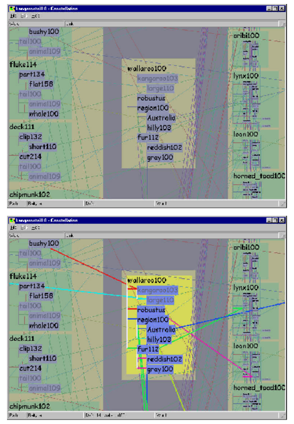

[Interactive Visualization of Large Graphs and Networks. Munzner. Ph.D. Dissertation, Stanford University, June 2000.]

Constellation Analysis

System	Constellation
What: Data	Three-level network of paths, subgraphs (definitions), and nodes (word senses).
Why: Tasks	Discover/verify: browse and locate types of paths, identify and compare.
How: Encode	Containment and connection link marks, horizontal spatial position for plausibility attribute, vertical spatial position for order within path, color links by type.
How: Manipulate	Navigate: semantic zooming. Change: Animated transitions.
How: Reduce	Superimpose dynamic layers.
Scale	Paths: 10–50. Subgraphs: 1–30 per path. Nodes: several thousand.

What-Why-How Analysis

- this approach is not the only way to analyze visualizations!
  - one specific framework intended to help you think
  - other frameworks support different ways of thinking
    - following: one interesting example

Algebraic Process for Visualization Design

- which mathematical structures in data are preserved and reflected in vis
  - negation, permutation, symmetry, invariance

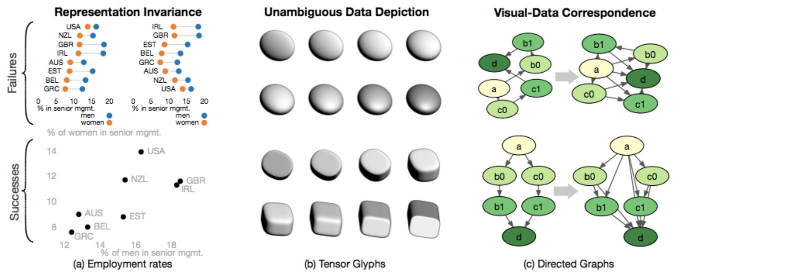

[Fig 1. An Algebraic Process for Visualization Design. Carlos Scheidegger and Gordon Kindlmann. IEEE TVCG (Proc. InfoVis 2014), 20(12):2181-2190.]

Algebraic process: Vocabulary

- **invariance** violation: single dataset, many visualizations
  - **hallucinator**
- **unambiguity** violation: many datasets, same vis
  - data change invisible to viewer
    - **confuser**
- **correspondence** violation:
  - can't see change of data in vis
    - **jumbler**
  - salient change in vis not due to significant change in data
    - **misleader**
  - match mathematical structure in data with visual perception
- we can X the data; can we Y the image?
  - are important data changes well-matched with obvious visual changes?