

CSE 564
VISUALIZATION AND VISUAL ANALYTICS
TIME VARYING AND STREAMING DATA

KLAUS MUELLER

COMPUTER SCIENCE DEPARTMENT
STONY BROOK UNIVERSITY

Lecture	Topic	Projects
1	Intro and logistics	
2	Basic visualizations and tasks, data types, examples, ethical considerations	
3	Data preparation (cleaning, imputation, data set integration)	
4	AI-assisted coding for VIS applications (design, debugging, refactoring)	Project #1 out
5	Big data and data reduction (distance/sim metrics, intro to clustering)	
6	High-D data: concept, subspaces, dimension reduction, PCA	
7	Cluster analysis: hierarchical, density, model, embedding, temporal	
8	Perception and cognition (human visual system, color, contrast)	Project #2(a) out
9	Visual design and aesthetics	
10	Visualization of multivariate and high-D data: linear methods, projections	
11	Vis. of multivariate and high-D data: non-linear methods, embeddings	
12	Visualization and AI: mutual support and capabilities (VIS4AI, AI4VIS)	Project #2(b) out
13	Principles of interaction: drive what is visualized, analyzed & how (HCI4VIS)	
14	Visual analytics, human-centered AI, mixed-initiative & collaborative VA	
15	Midterm #1 (tentative date)	
16	VA system design and evaluation, the nested model	
17	Midterm #1 discussion (tentative date)	Final proj. proposal call out
18	Visualization of hierarchical data	
19	Visualization of maps and data with geo-reference	
20	Vis. of time-varying, time-series, streaming data, progressive visualization	Final project proposal due
21	Applications in Visualization Research	
22	Applications in Visualization Research	
23	Ed Tufte's principles and critiques, responsible visualization, uncertainty	
24	Design of effective infographics	Final proj. prelim report due
25	Foundations scientific and medical visualization	
26	Intro to volume rendering	Bonus project out (Vol Ren)
27	Scientific visualization	
28	Midterm #2 (tentative date)	
Final	Final project demo on zoom (public)	All final proj. materials due

Example (from Kuchar et al. IEEE CG&A May/June 2006)

Wednesday 28 April 1999; Posted: 11:33 p.m. EDT (03:33 GMT): Another robbery occurred in southwestern Ontario today, making this the fourth robbery in the past few months. Delaware Bank in Brantford was robbed by three masked individuals who stole \$150,000 in currency and several unknown items from the bank's vault. The bank robbery occurred at 2:30, lasting all of five minutes and injuring eight people. All injured parties were taken to the local hospital where one died on arrival. Two people were released and the remaining people are in intensive care. This robbery is similar to a crime spree that started on the Chinese New Year. The first robbery occurred in the morning at Allegiant Bank in Richmond Hill, with the robbers taking more than \$100,000 in currency. The second robbery occurred about two weeks later at Banner Bank in Ajax and was caught on tape. The robbers arrived just as the bank opened, riding in a white van and wearing black ski masks and black outfits, and carrying automatic machine guns. During the first three minutes, the robbers instructed all of the patrons to face the wall and place their hands on their heads. While two of the robbers watched the patrons, the other robber took the bank manager and instructed him to open the vault. Other video captured the movement in the vault. The vault was opened about five minutes after the robbers arrived. The safe was blown two minutes later, and the robber removed only two safety deposit boxes and placed them in a bag. He then continued to club the manager and was back upstairs, yelling instructions to his buddies as they left the bank, not more than 20 minutes after they arrived. What was unusual was that no alarms were sounded. The third robbery happened at night at Carter Bank in Brampton, with only nearby homeowners mentioning that they do not remember hearing the bank alarm, only dogs barking for a while....

Question

Why are temporal relationships difficult to discern?

Time Series Data Are Everywhere

Temporal relationships can be difficult to discern because

- temporal ordering can be hard to determine
- event may occur in spatially disjoint locations
- what came before what – cause and effect
- what time shifts are acceptable/plausible?

To understand temporal relationships, an analyst:

- might need to reread the paragraph many times
- needs to cognitively make inferences between pieces of information

Visualization is key to externalize these relationships

- put it all out on “paper” and reason with it

Time is Special

What actually is time?

- how can one work with the metaphor of time's flow?
- what is the proper formalism that captures the time's special role in reasoning

The time variable is different than the other variables

- people consider it as an independent quantity
- can't go back in time
- our perception is that we have no control over it
- time is an ever-present thread that can help tie events together

Time as a Reference Frame

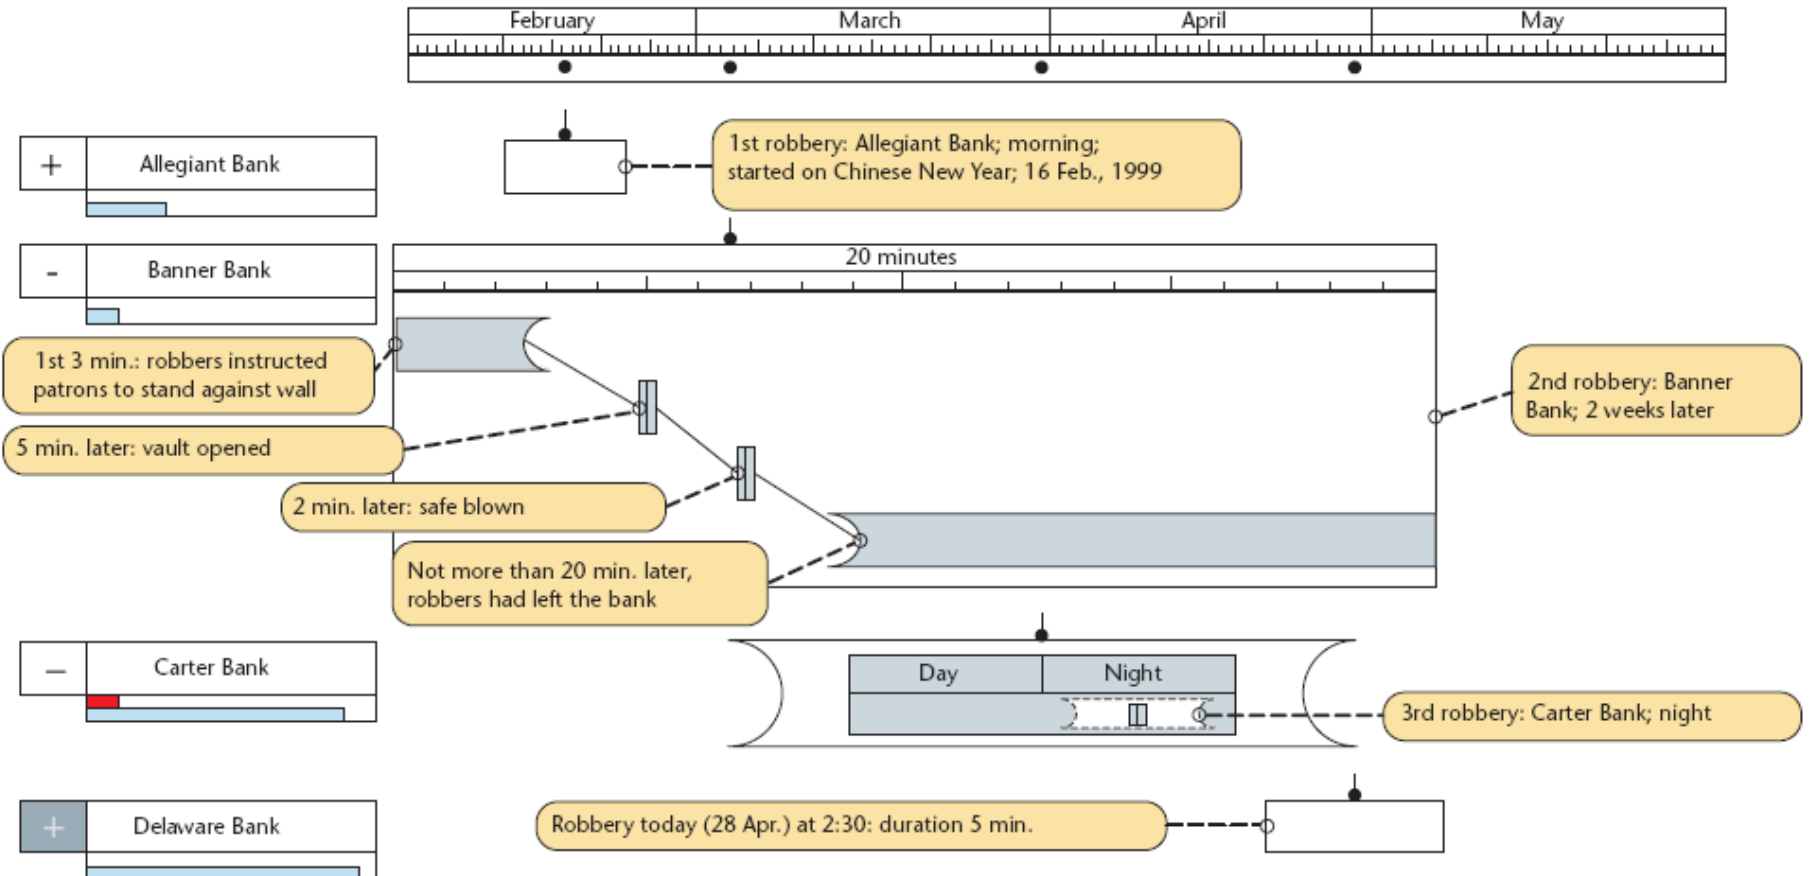
Calendars and time have reference frames

- Gregorian, Greenwich, EDT

Time is also often used in relative terms:

- “today”, “yesterday”, “fortnight”, “before Tuesday”, ...
- must normalize different reference systems into a common framework
- but it might be unknown what reference system was used individually
- “This robbery is similar to a crime spree that started on the Chinese New Year ...” – when is Chinese New Year?
- causes ambiguities, uncertainties, biases, conflicts

Back to Example (from Kuchar et al. IEEE CG&A May/June 2006)



Time: Tasks and Taxonomies

Often asked questions:

- when was something greatest/least?
- is there a pattern?
- are two series similar?
- does a data element exist at time t , and when?
- how long does a data element exist and how often?
- how fast are data elements changing
- in what order do they appear?
- do data elements exist together?

Different types of time series data:

- discrete vs. interval
- linear vs. cyclic
- ordinal vs. continuous
- ordered vs. branching vs. time with multiple perspectives

Traditional Time Series Visualizations


NVIDIA stock vs. NASDAQ (from yahoo! finance)



Fun one... (found in J. Stasko lecture)



Water Consumption in Edmonton During Olympic Gold Medal Hockey Game



Next...

A few good visualization metaphors for time


- there are quite a few of them...

ThemeRiver (Havre et al., 2002)


Each vertical line is a stacked bar chart for data at time t

- assembled into a “river” by ordering the bar charts along t and centering them on y for at each t

ThemeRiver-SingleAxis(时间轴) ■ DQ ■ TY ■ SS ■ QG ■ SY ■ DD



ThemeRiver (Havre et al., 2002)



River widens or narrows to depict changes in the collective strength of selected themes in the underlying documents. Individual themes are represented as colored "currents" flowing within the river.

Example shown here: newspaper themes around the Cuban Missile crisis

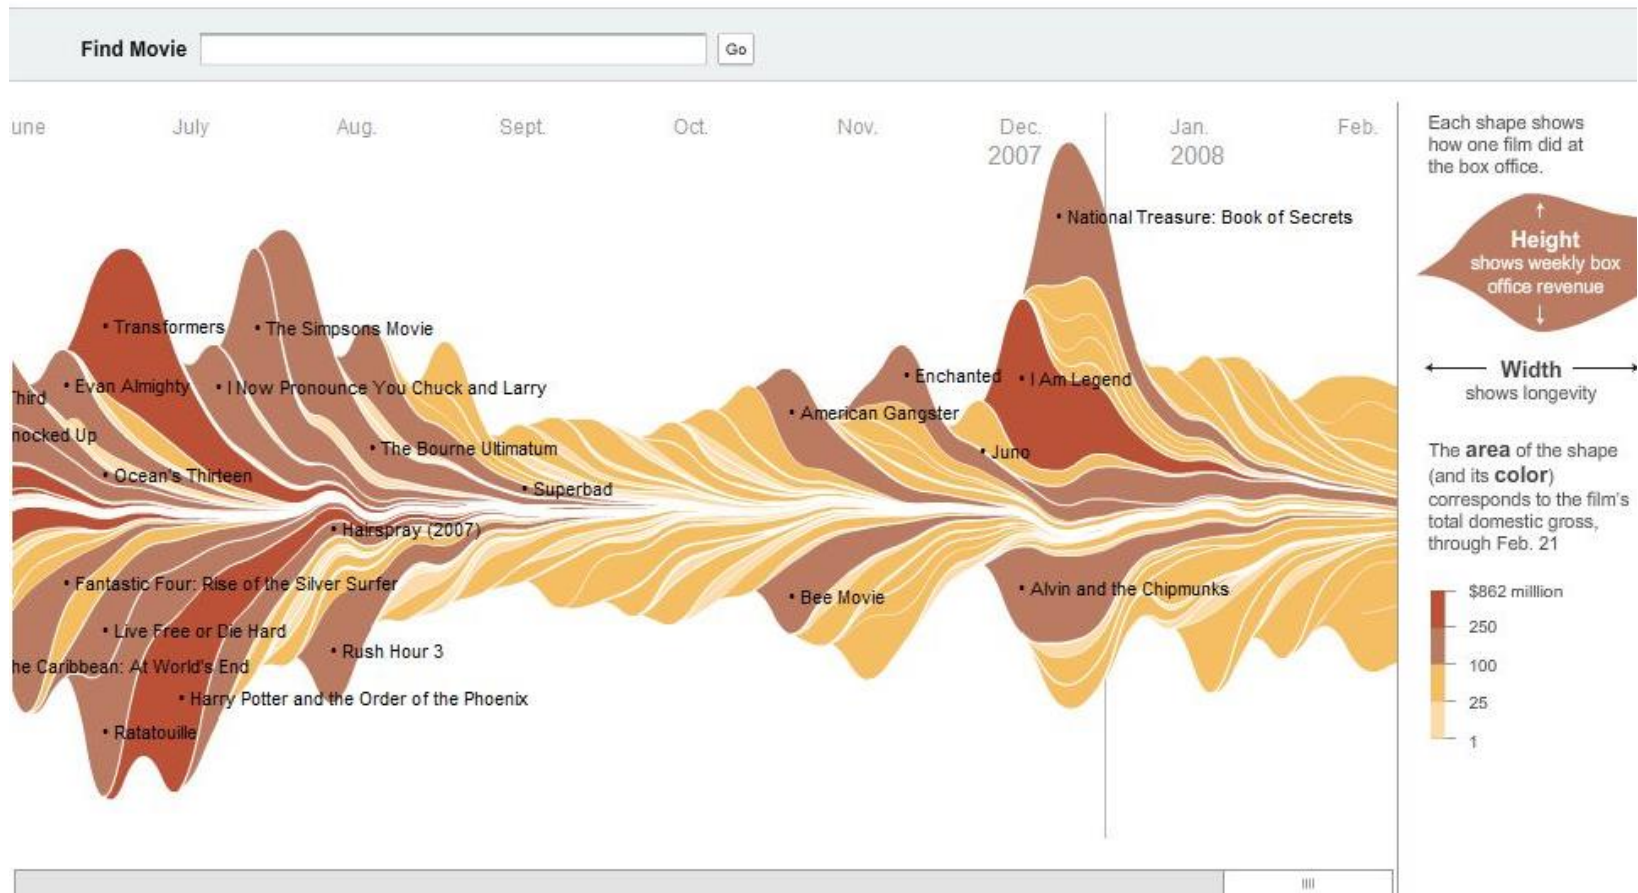
Stream Graphs

February 23, 2008

[SIGN IN TO E-MAIL OR SAVE THIS](#) | [FEEDBACK](#)

The Ebb and Flow of Movies: Box Office Receipts 1986 — 2008

Summer blockbusters and holiday hits make up the bulk of box office revenue each year, while contenders for the Oscars tend to attract smaller audiences that build over time. Here's a look at how movies have fared at the box office, after adjusting for inflation.



Stacked Area Charts


How Different Groups Spend Their Day

The American Time Use Survey asks thousands of American residents to recall every minute of a day. Here is how people over age 15 spent their time in 2008. [Related article](#)

Everyone

Sleeping, eating, working and watching television take up about two-thirds of the average day.

Everyone	Employed	White	Age 15-24	H.S. grads	No children
Men	Unemployed	Black	Age 25-64	Bachelor's	One child
Women	Not in lab...	Hispanic	Age 65+	Advanced	Two+ children



Name Voyager (<http://www.babynamewizard.com>)

New! Try the **NameMapper** to see where your favorite names are being used, and **Namipedia** for full info on every name

>

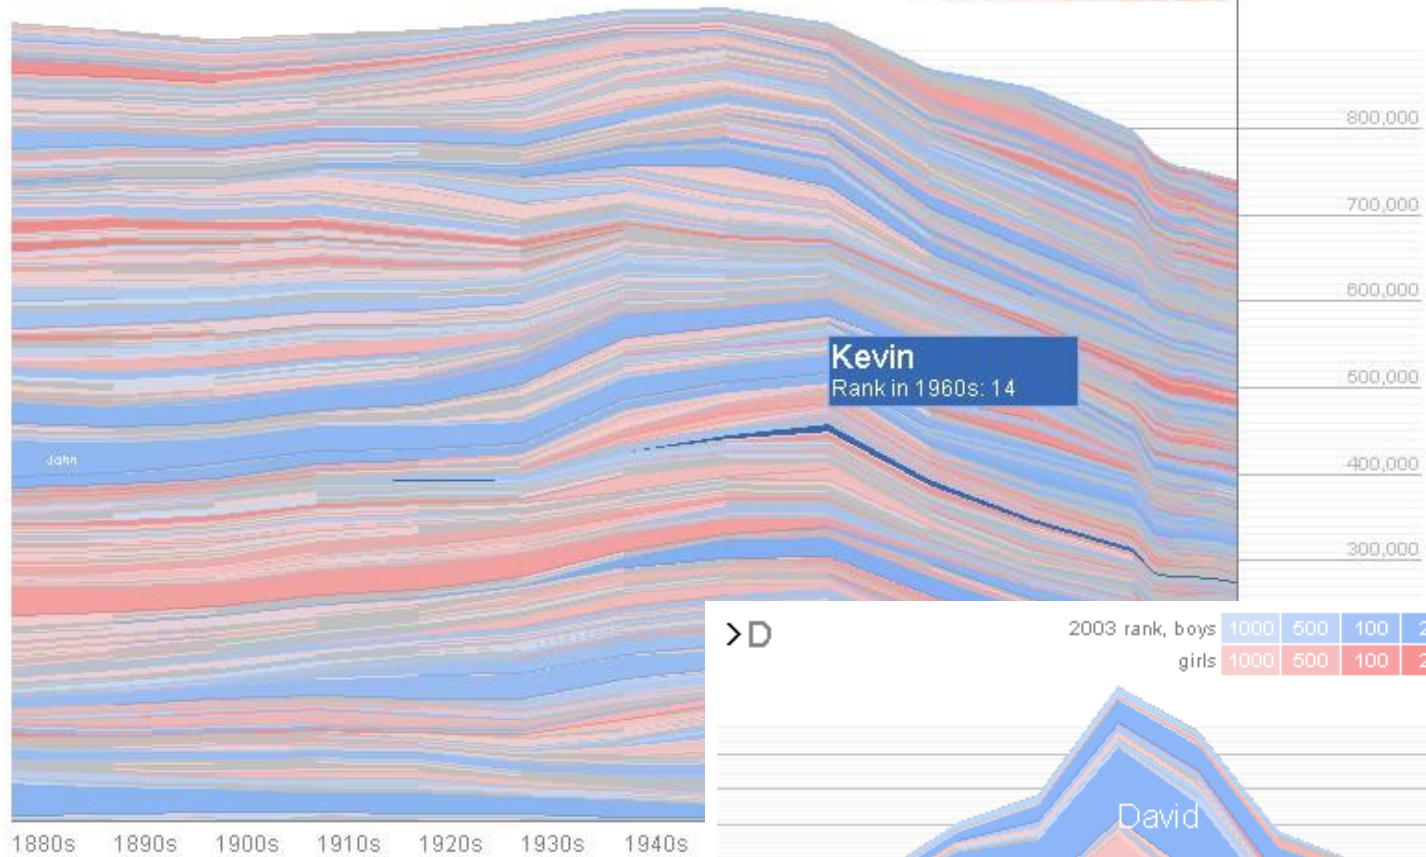
boys girls both

2007 rank, boys

1000	500	100	25	1
------	-----	-----	----	---

girls

1000	500	100	25	1
------	-----	-----	----	---



also [here](#)

>D

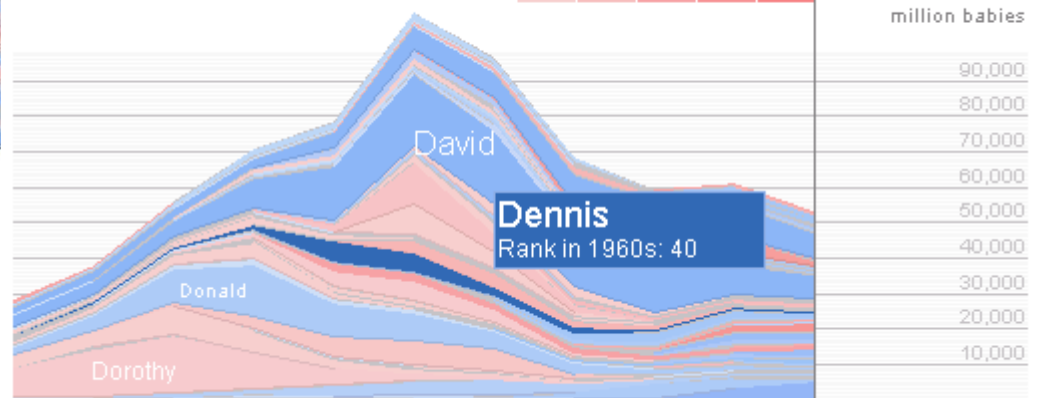
2003 rank, boys

1000	500	100	25	1
------	-----	-----	----	---


girls

1000	500	100	25	1
------	-----	-----	----	---

Names starting with "D", per million babies



1900s 1910s 1920s 1930s 1940s 1950s 1960s 1970s 1980s 1990s 2003



Question

Can you tell me who is who?

- I tell you all the names there are and the age of each person
- can you assign them? (hint: use the Name Voyager)



Medical Data

Medical data are often displayed along time

- natural to humans
- progression of disease
- appearance of symptoms
- time course of treatment and outcome
- but also time signals (ECG, blood pressure, etc.)

A popular example is Lifelines and Lifelines2


- Shneiderman and Plaisant et al.
- <http://www.cs.umd.edu/hcil/lifelines/>

LifeLines: Patient-Centric

In progress version of LifeLine Frame [not for demo purposes - Microsoft Internet Explorer - [Working Offline]

File Edit View Go Favorites Help

Back Links Best of the Web Channel Guide Customize Links Internet Explorer News Internet Start

 **Linda Simpson**
Female 40

Line from input file: %:3-10-1997,3-12-1997_black,p10,Sonogram,images/babysonogra

LifeLine

92 93 94 95 96 97

Notes: Tobacco, Depression, Lyme, Arthritis, Obesity, AtrialFlutter, Flu, Pneumonia, KneePain, Fatigue>Diabetes, Diabe, Pregnancy

Hosp.: Appendectomy, Pneumonia, KneeSurgery

Tests: BloodEKG, EKG, Xray, Blood, Blood, Blood, Sonogr

Meds.: Prozac, Heartdrug, Ventolyn, Antib., Advil, Advil, Insulin, Insulir

Others: PhysicalTherapy, LowSaltFatDiet

Immun.: TBtest, Tetanos, Flu

DATE: 5-18-88 PATIENT ID: [redacted] 17aa/3.8

LifeLine Control Panel

Layout Label

Default

Quick Compact


Slow Compact

Chronologically Ordered

Event Ordered

Apply OK Cancel

Warning: Applet Window



LifeLines2: Pattern-Centric

Goals:

- bring out temporal categorical patterns across multiple records
- categorical event data such as complaints, diagnoses, treatments
- play important roles in health providers decision making

Features

- allows users to manipulate multiple records simultaneously
- understand relative temporal relationships across records
- 3 operators: align, rank, filter
- temporal summaries allow multiple groups of records to be compared

LifeLines2: Screenshot

LifeLines 2 - contrast+creatine-large.bt.out(deID).bt (deidentified data)

File About

Record View Comparison View

In 2 weeks Records 377/377

Zoom to: Days Annotate!

Rank by [# of CREAT-H]

Align by [1st occurrence of Radiology Contrast]

Rank	Record ID	CREAT-H	CREAT-	Radiology Contrast
49	0000333	49	4	1
29	0002830	29	1	1
25	0001471	25	4	1
25	0000125	25	1	1
24	0003295	24	1	1
24	0002107	24	1	1
24	0002215	24	1	1
24	0001003	24	1	1

Events

Distribution

Records: 377

Day	CREAT-H	CREAT-	CREAT-L
-5 D	15	12	19
-4 D	30	41	116
-3 D	1	116	116
-2 D	434	305	263
-1 D	305	263	206
0 D	263	206	191
1 D	206	191	161
2 D	191	161	143
3 D	161	143	113
4 D	143	113	100
5 D	113	100	96
6 D	100	96	81
7 D	96	81	66
8 D	81	66	56
9 D	66	56	50
10 D	56	50	50
11 D	50	50	50
12 D	50	50	50
13 D	50	50	50

ARF Comparison

Align by... Radiology Contrast(3598) 1st occurrence

Rank by... # of CREAT-H(1190)

Filter by... No Filter

Keep Selected Remove Selected Reset Filters

Collapse all Save Records as New Group

Show Distribution of... CREAT-H(1190) CREAT-(2346) CREAT-L(207)

View Record Group... In 2 weeks Less Controls

Events Icons Labels Alignment

Legend

- CREAT-H
- CREAT-
- CREAT-L
- 24H UCREAT-H
- 24H UCREAT-

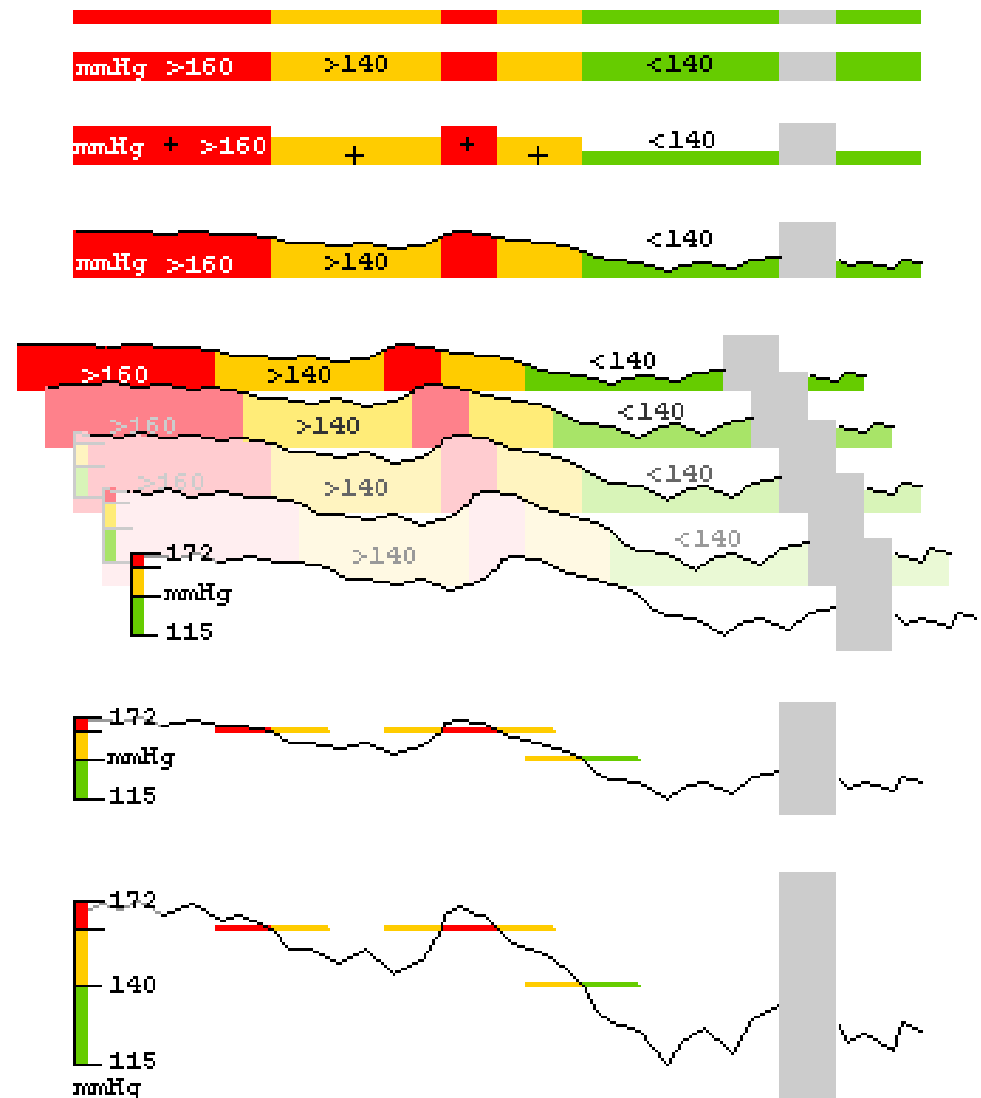
Edit Events

HCIL UNIVERSITY OF MARYLAND Copyright 2007 - 2009 HCIL UMD

Multi-Scale and Abstractions (Aigner et al., IEEE TVCG, 2008)

Deal with different levels of detail


- illustrative abstraction
- overview + detail
- used here for medical data




Cyclic Patterns

Time data are often cyclic

- spiral displays are good to bring out cyclic patterns
- one period per loop (for example, a year)



linear layout



radial layout

sunshine pattern





Figure 2. An indented spiral, with spokes, showing monthly consumption percentages for *Baphia Capparidifolia* during the period 1980 - 1988.

Cyclic Patterns

May have to play around to discover the cycles



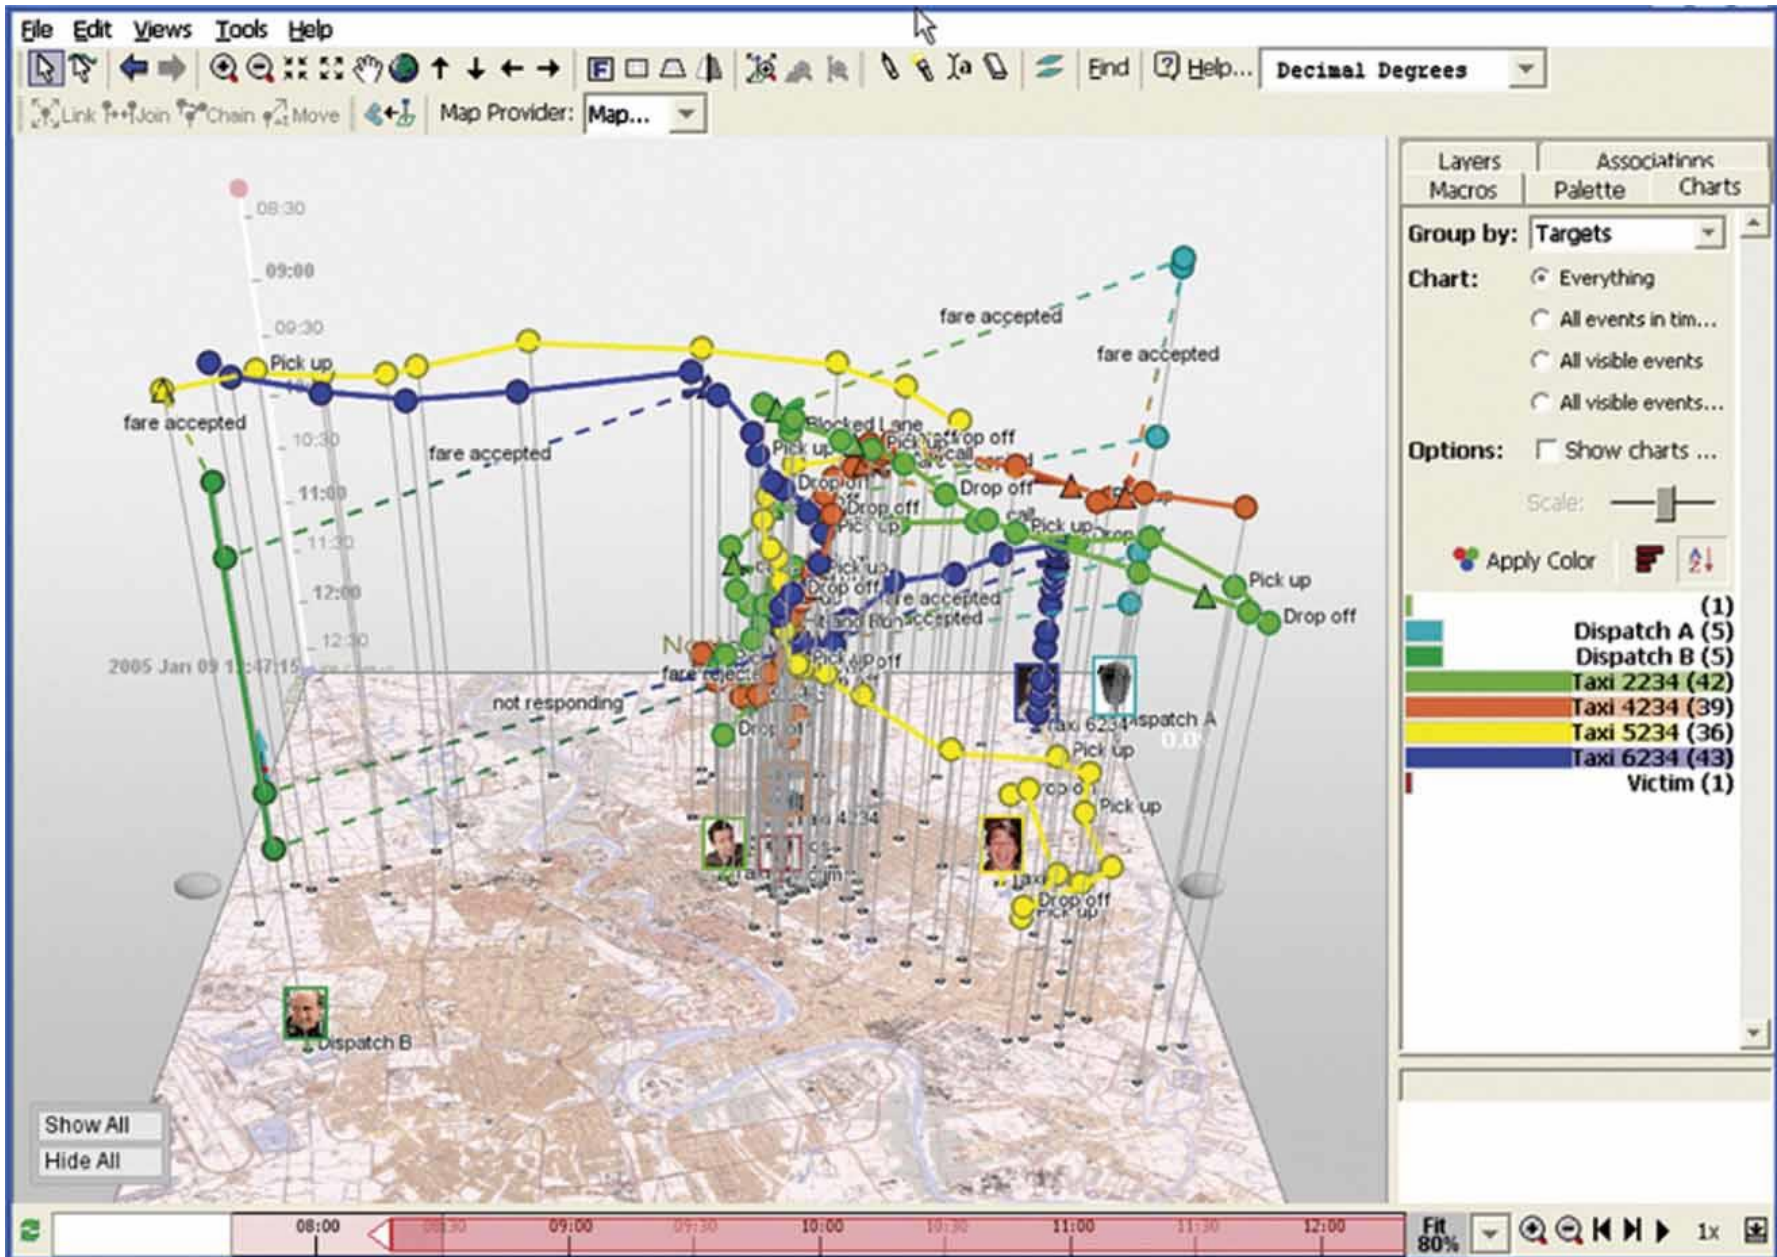
Combine Space and Time

OculusInfo Geotime application

- events are represented in an X,Y,T coordinate space
- the X,Y plane shows geography
- the vertical T axis represents time
- events animate in time vertically through the 3-D space as the time slider bar is moved.



Geotime



Interaction

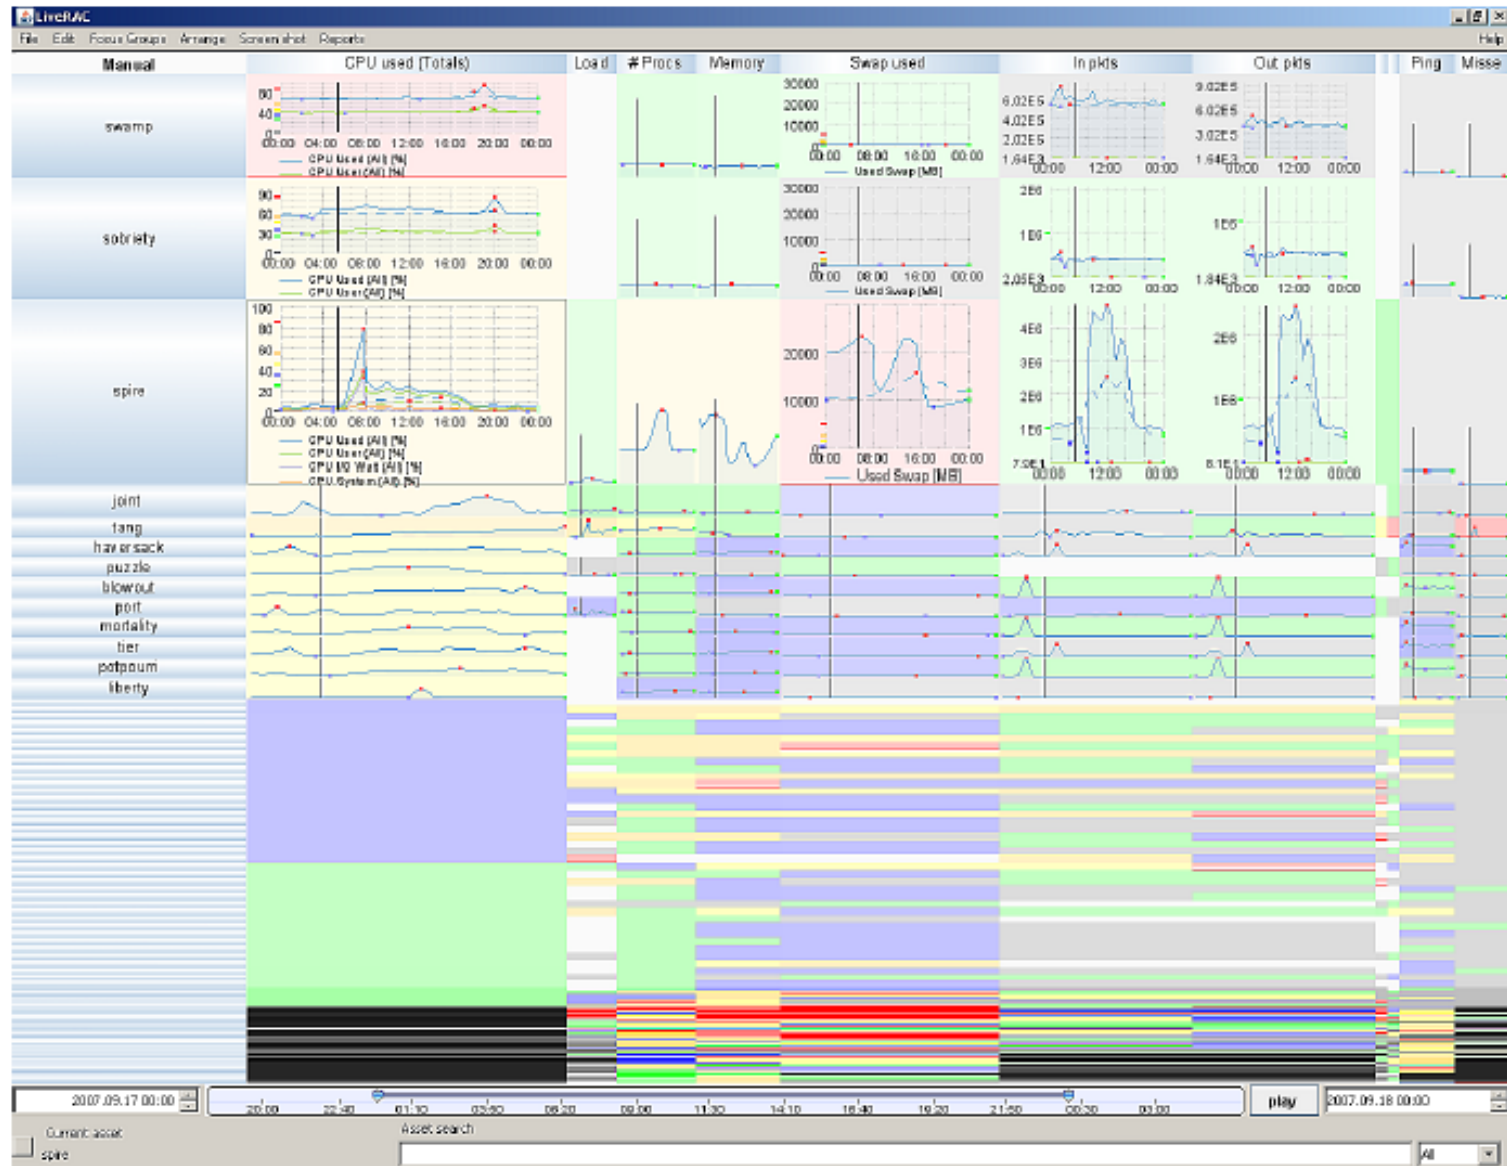
As complexity increases, interaction capabilities are key

- show more context of what else was going on at that time
- likely have to abstract some of the information
- allow several different levels of detail at once
- allow drill-down for details
- use dashboard design with many linked information displays


Example: Computer system management

- LiveRAC system (McLachlan et al.)
- next two slides

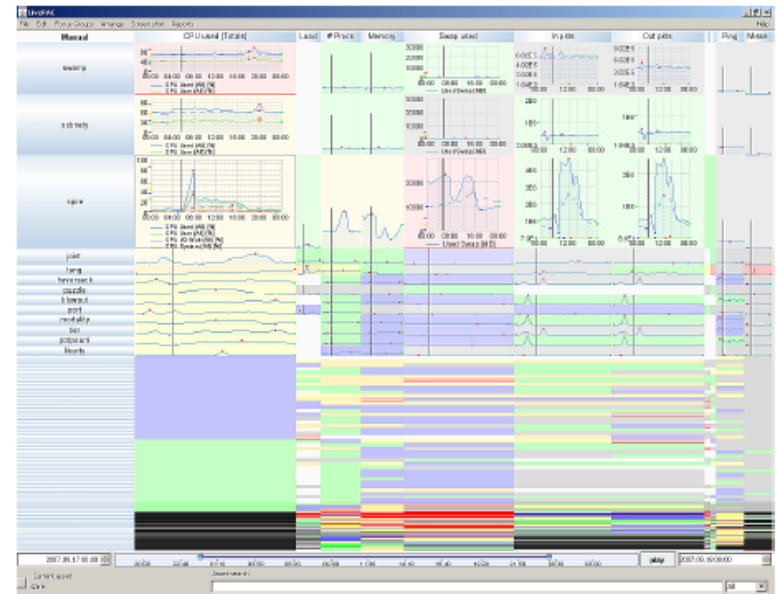
LiveRAC



LiveRAC




(a)



(b)

Figure 3. LiveRAC shows a full day of system management time-series data using a reordable matrix of area-aware charts. Over 4000 devices are shown in rows, with 11 columns representing groups of monitored parameters. (a): The user has sorted by the maximum value in the *CPU* column. The first several dozen rows have been stretched to show sparklines for the devices, with the top 13 enlarged enough to display text labels. The time period of business hours has been selected, showing the increase in the *In pkts* parameter for many devices. (b): The top three rows have been further enlarged to show fully detailed charts in the *CPU* column and partially detailed ones in *Swap* and two other columns. The time marker (vertical black line on each chart) indicates the start of anomalous activity in several of *spire's* parameters. Below the labeled rows, we see many blocks at the lowest semantic zoom level, and further below we see a compressed region of highly saturated blocks that aggregate information from many charts.

LiveRAC Video



video

Next – Streaming Data

Time series data with no end...

Types of Streaming data

Transaction streams

- credit card, point-of-sale transaction
- at a supermarket, or online purchase of an item

Web click-streams

Social streams

- online social networks such as Twitter
- speed and volume of the stream typically scale super-linearly with the number of actors

Network streams

- communication networks contain large volumes of traffic streams
- often mined for intrusions, outliers, or other unusual activity

Challenges (1)


One-pass constraint

- data is generated continuously and rapidly
- it is assumed that the data can be processed only once
- archival for future processing is not possible
- prevents use of iterative mining or model building algorithms that require multiple passes over the data

Concept drift, concept evolution, feature evolution

- data may evolve over time
- various statistical properties, such as correlations between attributes, correlations between attributes and class labels, and cluster distributions may change over time

Concept Drift

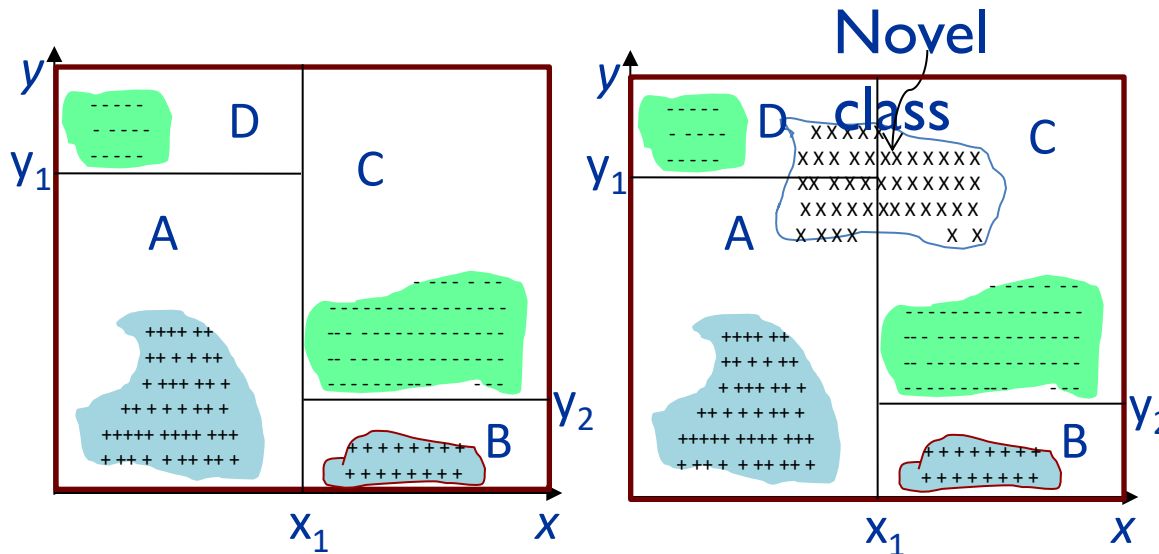


Negative instance •

Positive instance ○

Instances victim of concept-drift ●

Concept Evolution



Classification rules:

R1. if $(x > x_1 \text{ and } y < y_2)$ or $(x < x_1 \text{ and } y < y_1)$ then class = +

R2. if $(x > x_1 \text{ and } y > y_2)$ or $(x < x_1 \text{ and } y > y_1)$ then class = -

Existing classification models misclassify novel class instances

CluStream Clustering

The concept drift in an evolving data stream changes the clusters significantly over time

- need a clustering algorithm that can deal with this
- CluStream is such an algorithm

CluStream's online microclustering clustering stage

- processes the stream in real time to continuously maintain summarized but detailed (micro-)cluster statistics of the stream

CluStream's offline macroclustering stage

- further summarizes these detailed clusters
- provides the user with a more concise understanding of the clusters over different time horizons and levels of temporal granularity.

Microclustering Algorithm

There are k microclusters

- a new data point either needs to be absorbed by a microcluster, or it needs to be put in a cluster of its own

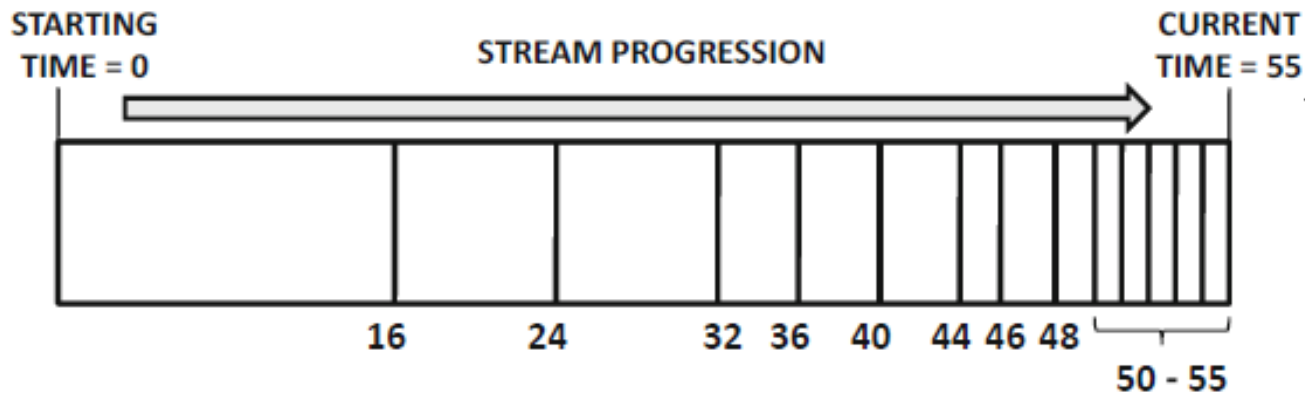
Algorithm

- determine distance of the new data point to all current microcluster centroids
- assign the point to the closest cluster and update the statistics
- if the point does not fall within the maximum boundary of any microcluster create a new microcluster
- to create this new microcluster, the number of other microclusters must be reduced by 1 to free memory availability
- achieve this by either deleting an old microcluster or merging two of the older clusters
- decide by examining the staleness (using the time stamp statistics) of the different clusters, and the number of points in them
- determine whether one of them is “sufficiently” stale to merit removal
- if no microcluster is stale, then a merging of the two microclusters is initiated

Pyramidal Time Frame

Store microclusters statistics periodically to enable time horizon-specific analysis of the clusters

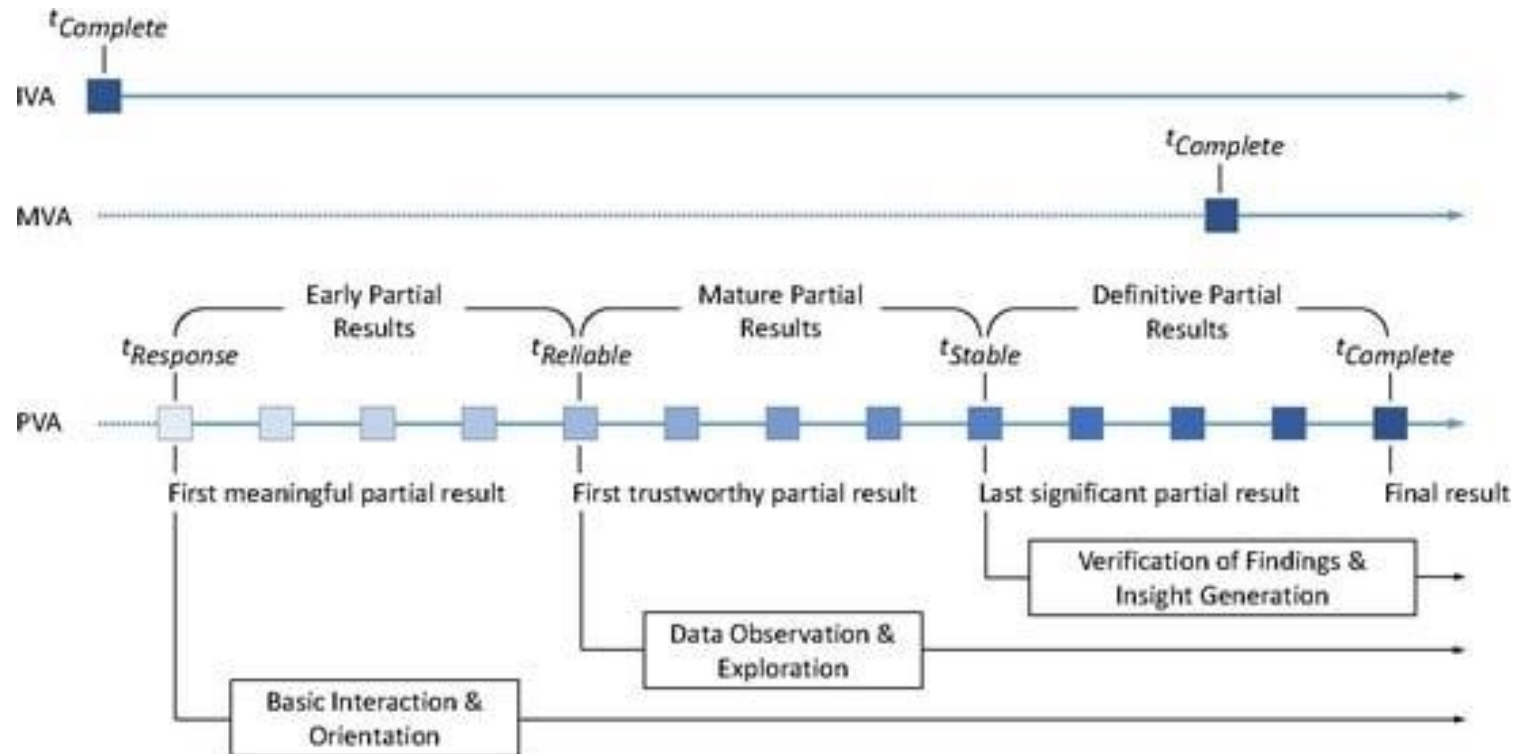
- the microcluster snapshots are stored at varying levels of granularity depending on the recency of the snapshot



PROGRESSIVE VISUAL ANALYTICS: REFINE VISUALIZATIONS OVER TIME

CONCEPT

Enables users to interact with and analyze partial, intermediate results of long-running computations rather than waiting for complete results.



KEY ASPECTS

Partial Results:

- Algorithms are adapted to output refined results iteratively, allowing analysis to begin immediately.

User Steering:

- Analysts can interact with the system to guide computations towards relevant data subspaces, prioritizing interesting data points before the full analysis finishes.

Improved Efficiency:

- By providing early results, PVA prevents long waiting periods and allows users to terminate, adjust, or refine analyses in real-time.

BENEFITS

User-driven exploration:

- Analysts can pivot quickly based on partial insights, improving efficiency.

Handling Big Data:

- Allows visualizing trends in massive datasets where traditional full-computation is too slow.

Immediate feedback:

- Reduces the "latency gap" between asking a question and getting an answer, leading to more responsive and engaging

EXAMPLE

Steering by Example:

- Side-by-side view from a progression that uses the 'steering-by-example' approach to steer the computation on the left, compared to a progression using random uniform sampling on the right.

