DATA ANALYSIS AT SCALE

PETRA ISENBERG (slides by WESLEY WILLETT)

VISUAL ANALYTICS 04 NOV 2014

DATA ANALYSIS AT SCALE

CHALLENGES

ANALYSIS AND CLUSTER COMPUTING

INTERACTING WITH BIG DATA

PARALLELIZING HUMAN INTELLIGENCE

CHALLENGES FOR ANALYZING LARGE DATA SETS

SIZE

SPEED

ATTENTION

SIZE

KILOBYTES OF DATA

MEGABYTES OF DATA

GIGABYTES OF DATA

TERABYTES OF DATA

PETABYTES OF DATA

...

SIZE

KILOBYTES OF DATA

MEGABYTES OF DATA

GIGABYTES OF DATA

TERABYTES OF DATA

PETABYTES OF DATA

...

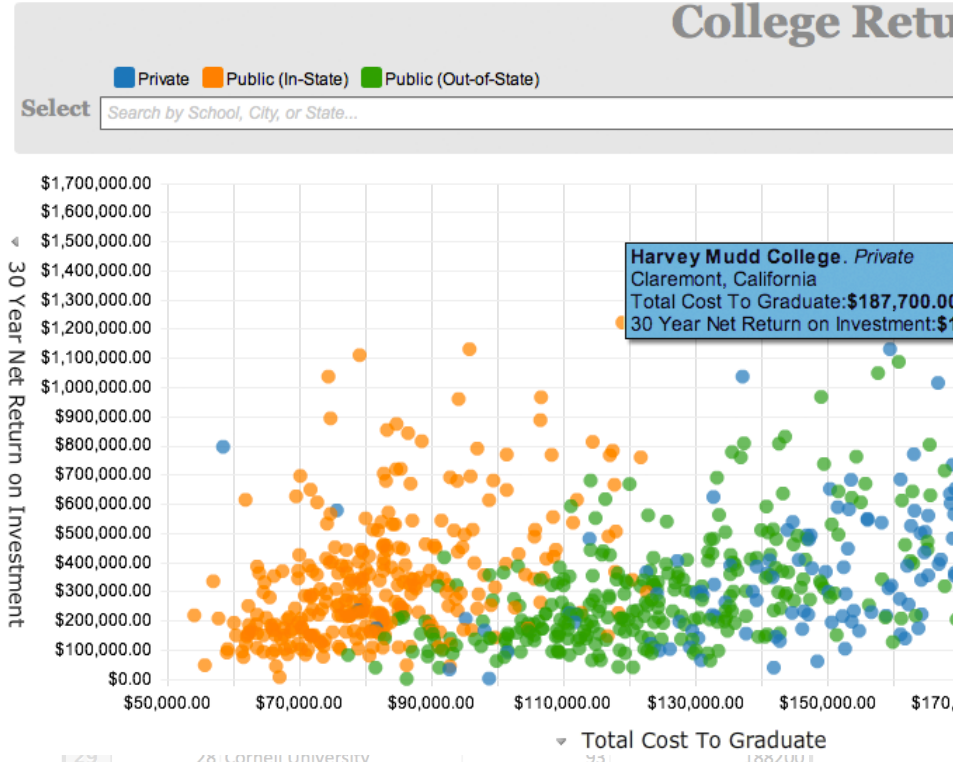

30	29	Lafayette College	89	179500
31	30	Stevens Institute of Technolo	79	222600
32	31	Georgia Institute of Technolo	77	79140
33	32	Colorado School of Mines	68	160700
34	33	Santa Clara University	85	187500
35	34	College of the Holy Cross	98	182600
36	35	University of Chicago	92	185800

SIZE

KILOBYTES OF DATA

MEGABYTES OF DATA

GIGABYTES OF DATA

TERABYTES OF DATA

PETABYTES OF DATA

...

EVEN A MEGABYTE IS MORE BITS OF DATA THAN THERE ARE PIXELS ON A SCREEN!

SIZE

KILOBYTES OF DATA

MEGABYTES OF DATA

GIGABYTES OF DATA

TERABYTES OF DATA

PETABYTES OF DATA

...

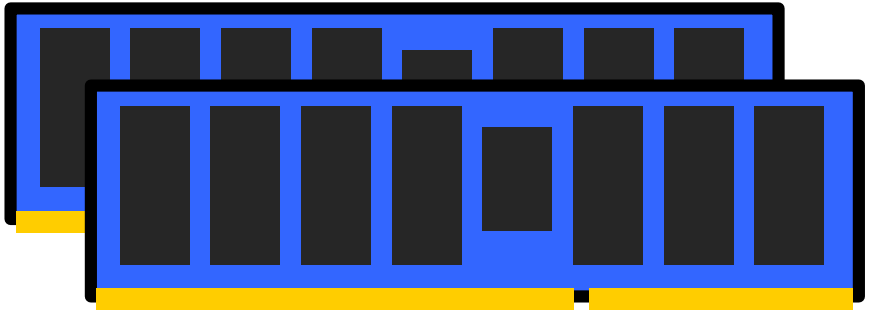

MORE DATA THAN CAN FIT IN MEMORY

SIZE

KILOBYTES OF DATA

MEGABYTES OF DATA

GIGABYTES OF DATA

TERABYTES OF DATA

PETABYTES OF DATA

...

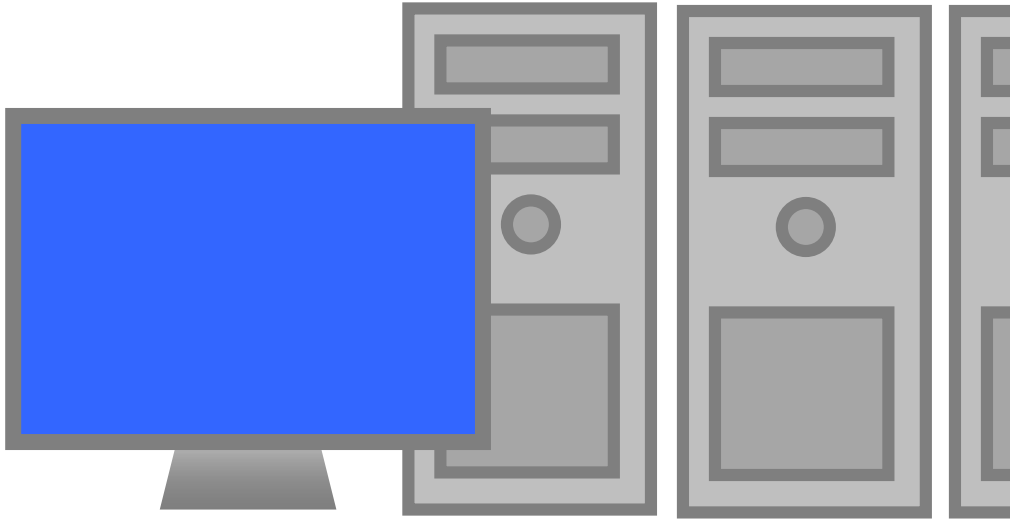

MORE DATA THAN CAN FIT ON ONE MACHINE!

SIZE

KILOBYTES OF DATA

MEGABYTES OF DATA

GIGABYTES OF DATA

TERABYTES OF DATA

PETABYTES OF DATA

SPEED

~0.1 SECOND	DIRECT MANIPULATION
~1 SECOND	INTERACTIVE
~10 SECONDS	QUERY / RESPONSE
MINUTES	…
HOURS	BATCH PROCESSING (VERY SLOW)

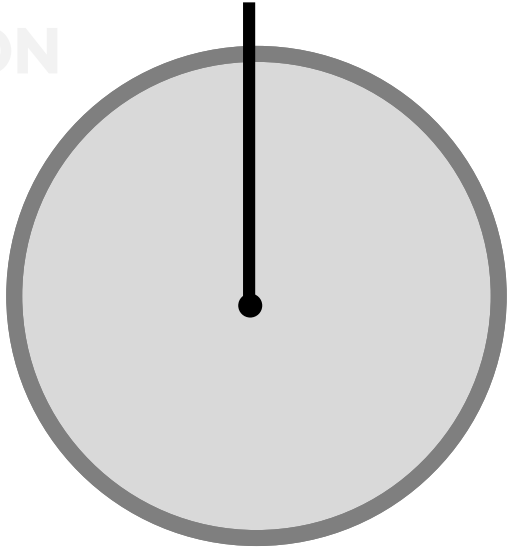

SPEED

~0.1 SECOND DIRECT MANIPULATION

~1 SECOND INTERACTIVE

~10 SECONDS QUERY / RESPONSE

MINUTES ...

HOURS BATCH PROCESSING (VERY SLOW)

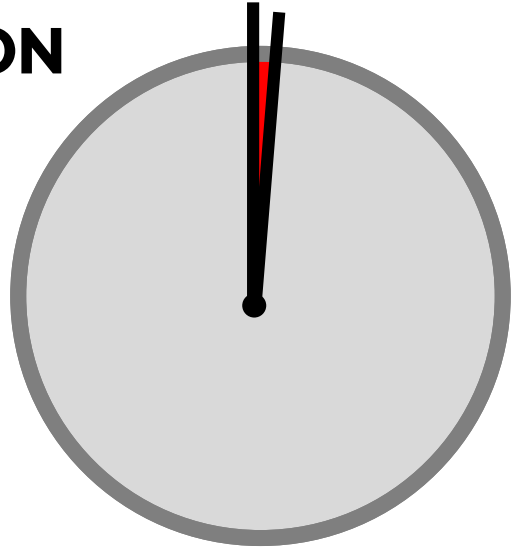

SPEED

~0.1 SECOND	DIRECT MANIPULATION
~1 SECOND	INTERACTIVE
~10 SECONDS	QUERY / RESPONSE
MINUTES	...
HOURS	BATCH PROCESSING (VERY SLOW)

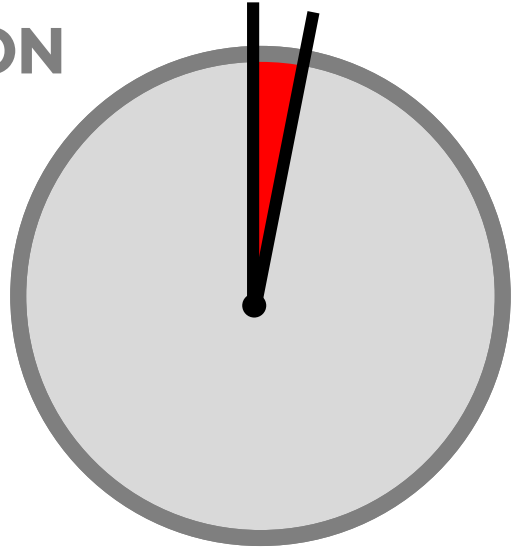

SPEED

~0.1 SECOND	DIRECT MANIPULATION
~1 SECOND	INTERACTIVE
~10 SECONDS	QUERY / RESPONSE
MINUTES	…
HOURS	BATCH PROCESSING (VERY SLOW)

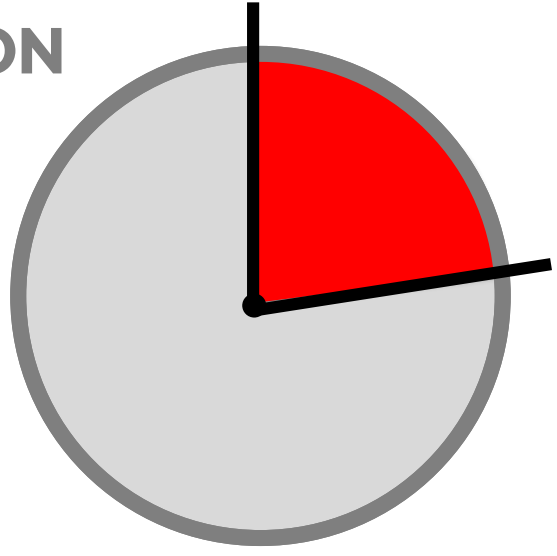

SPEED

~0.1 SECOND DIRECT MANIPULATION

~1 SECOND INTERACTIVE

~10 SECONDS QUERY / RESPONSE

MINUTES

HOURS

...

BATCH PROCESSING (VERY SLOW)

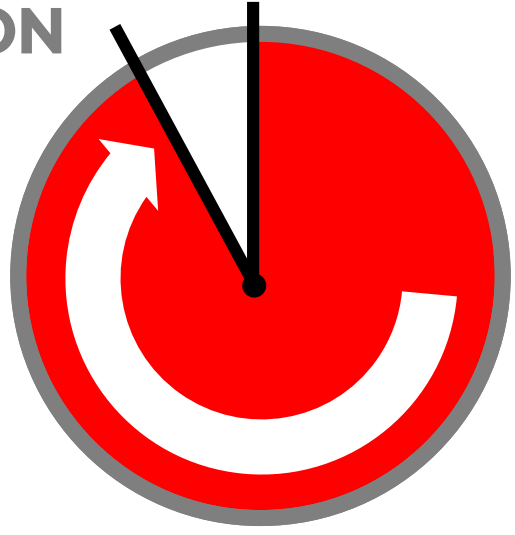

ATTENTION

EVERY PERSON ONLY HAS A FINITE NUMBER OF WORKING HOURS

5-8 PERSON-HOURS PER DAY

1,489 PERSON-HOURS PER YEAR (FRANCE)

(1,388 GERMANY 2,163 IN S. KOREA 1,788 IN USA) [OECD STATS]

HOW LONG CAN YOU AFFORD TO SPEND FINDING EXAMPLES, PROCESSING A DATASET, OR ANSWERING A QUESTION?

ATTENTION

AN INDIVIDUAL ANALYST IS UNLIKELY TO BE ABLE TO SEE DATA FROM MANY PERSPECTIVES

"MANY EYES FIND MORE BUGS"

DATA ANALYSIS AT SCALE

CHALLENGES

ANALYSIS AND CLUSTER COMPUTING

INTERACTING WITH BIG DATA

PARALLELIZING HUMAN INTELLIGENCE

ANALYSIS & CLUSTER COMPUTING

BIG DATASETS ARE LIKELY TO BE SPREAD OUT ACROSS A CLUSTER (OR CLUSTERS)

ANALYSIS REQUIRES DISTRIBUTED DATA PROCESSING

HOW CAN WE PERFORM ANALYSIS ACROSS A CLUSTER?

How can we split work across machines?

MAP-REDUCE

Jeffrey Dean and Sanjay Ghemawat (Google) 2004

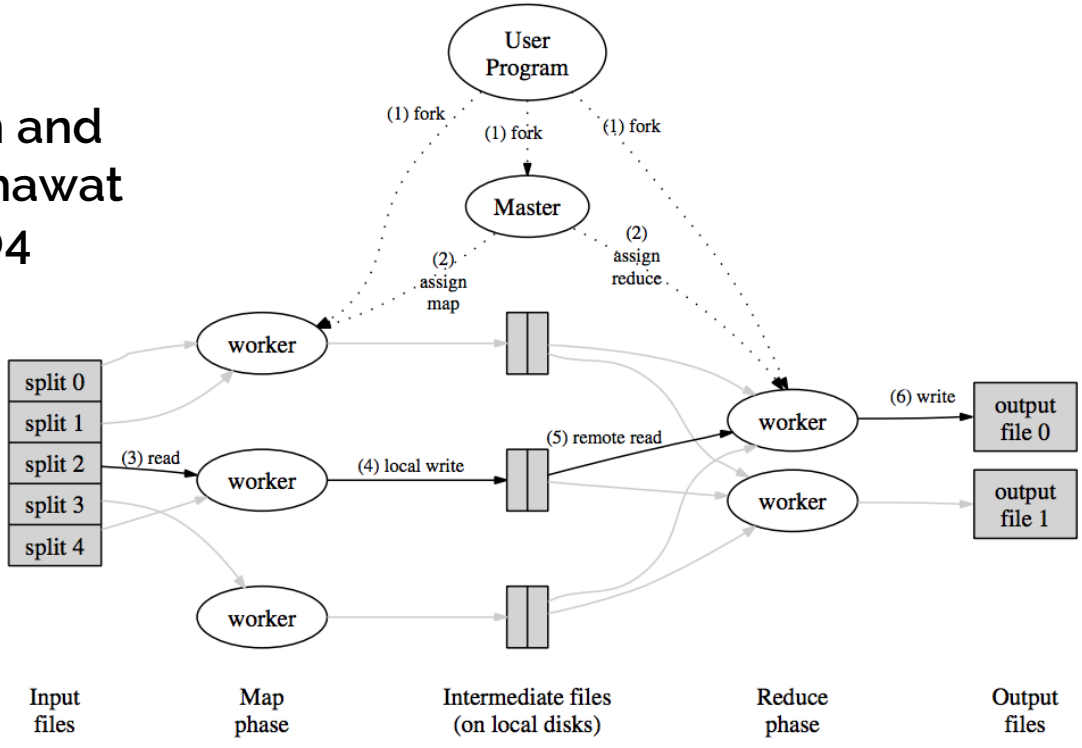

A SIMPLE EXAMPLE

HOW TO COUNT NUMBER OF TIMES WORDS OCCUR IN A DOCUMENT?

(IF THAT DOCUMENT IS SPREAD ACROSS MANY MACHINES)

"I am Sam
I am Sam
Sam I am
Do you like
Green eggs and ham?"

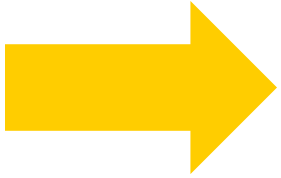

I: 3
am: 3
Sam: 3
do: 1
you: 1
like: 1
…

[K. Ousterhout - UCB 194-16]

JUST A HASH TABLE

“I am Sam
I am Sam
Sam I am
Do you like
Green eggs and ham?”

{}

[K. Ousterhout - UCB 194-16]

JUST A HASH TABLE

"I am Sam
I am Sam
Sam I am
Do you like
Green eggs and ham?"

{I:1}

JUST A HASH TABLE

“I am Sam
I am Sam
Sam I am
Do you like
Green eggs and ham?”

{I:1,
am:1}

JUST A HASH TABLE

“I am Sam
I am Sam
Sam I am
Do you like
Green eggs and ham?”

{I:1,
am:1,
Sam:1}

JUST A HASH TABLE

“I am Sam
I am Sam
Sam I am
Do you like
Green eggs and ham?”

{I:2,
am:1,
Sam:1}

BUT YOU SAID THE DOCUMENT IS REALLY BIG?

[K. Ousterhout - UCB 194-16]

COMPUTE IN PARALLEL

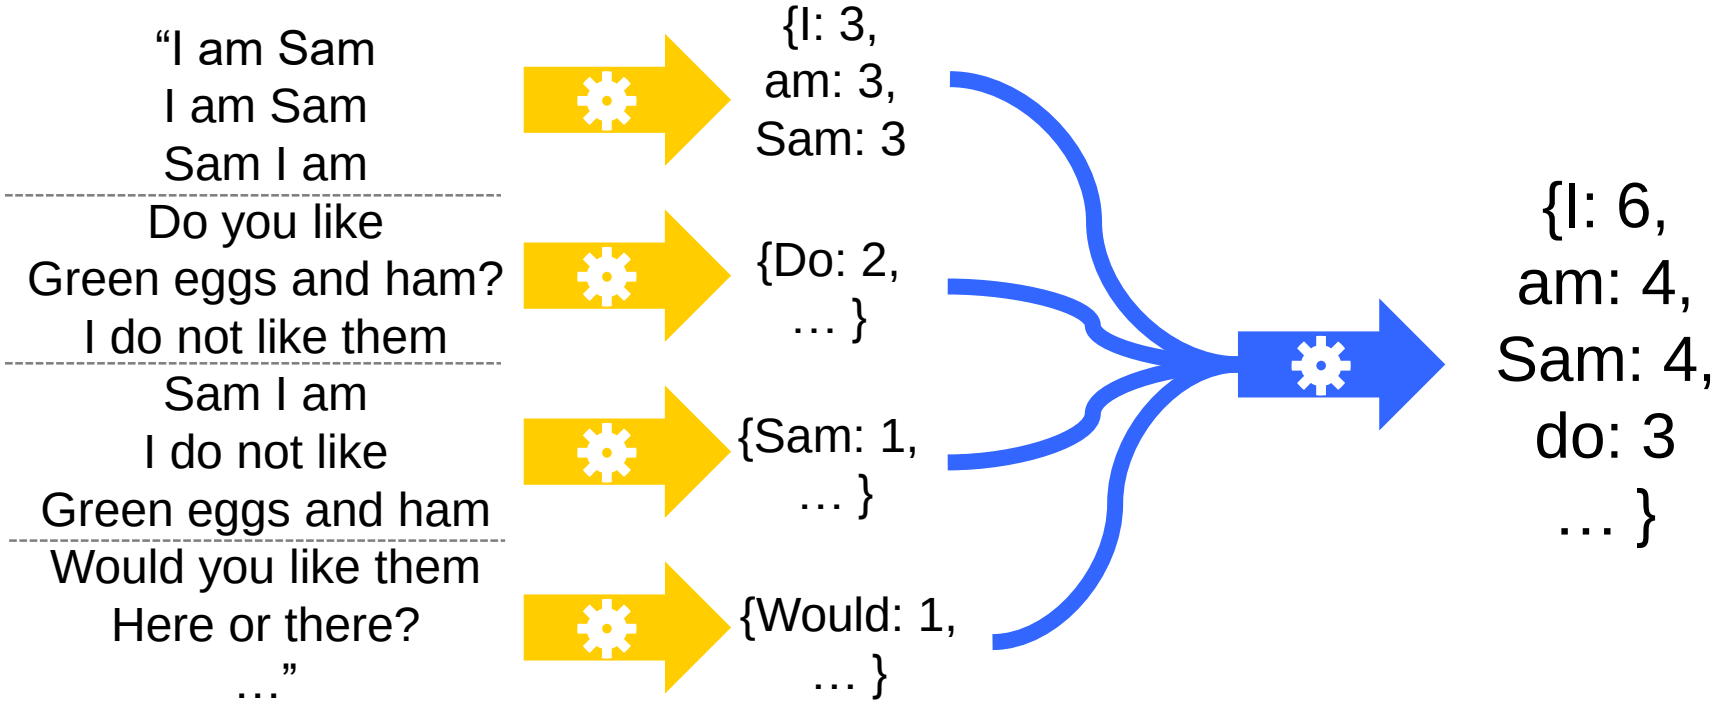

[K. Ousterhout - UCB 194-16]

COMPUTE IN PARALLEL

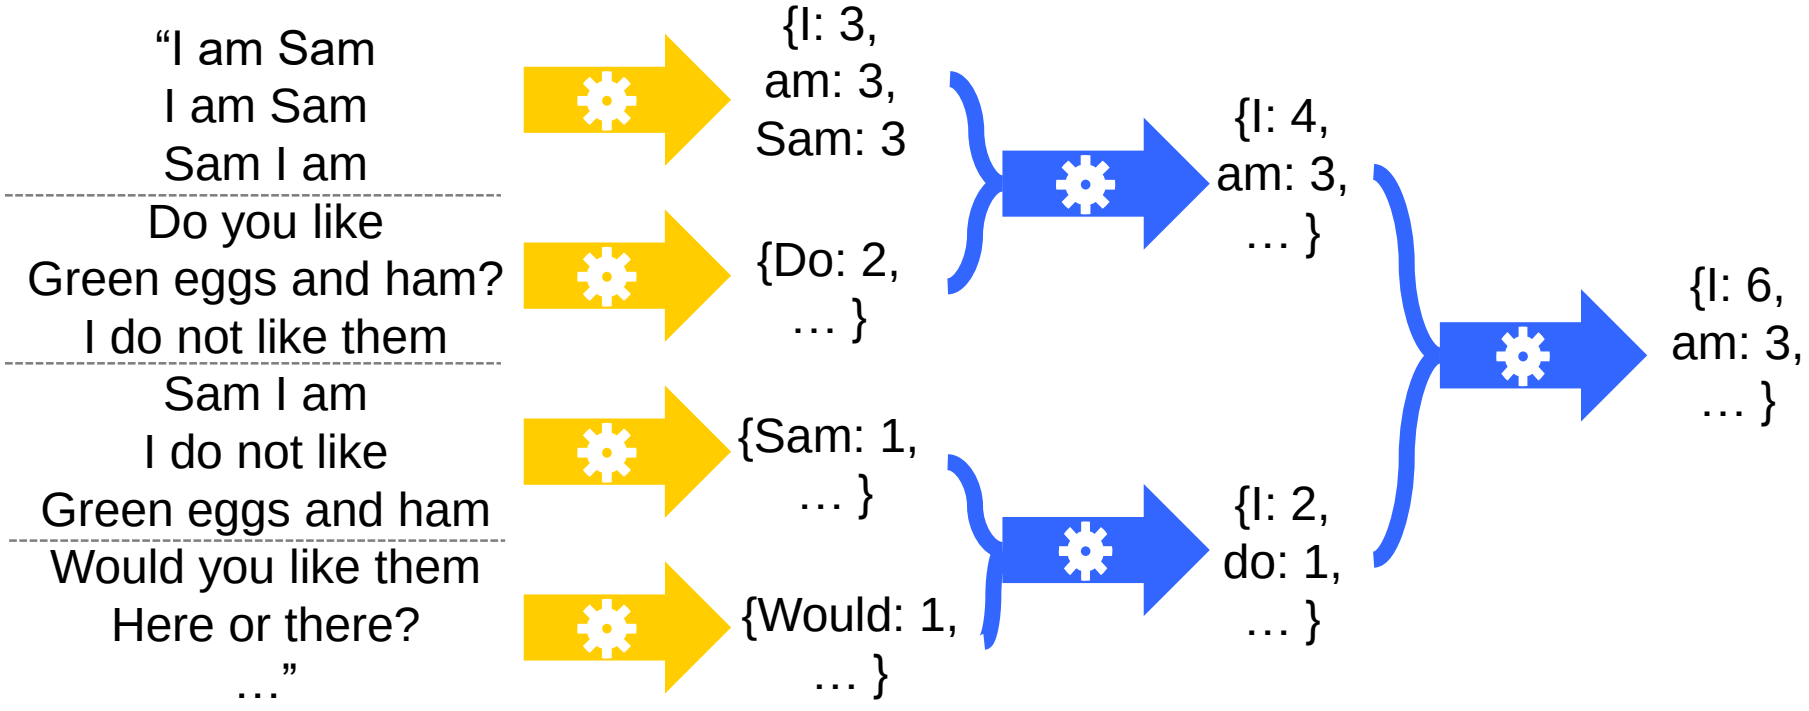

[K. Ousterhout - UCB 194-16]

COMPUTE IN PARALLEL

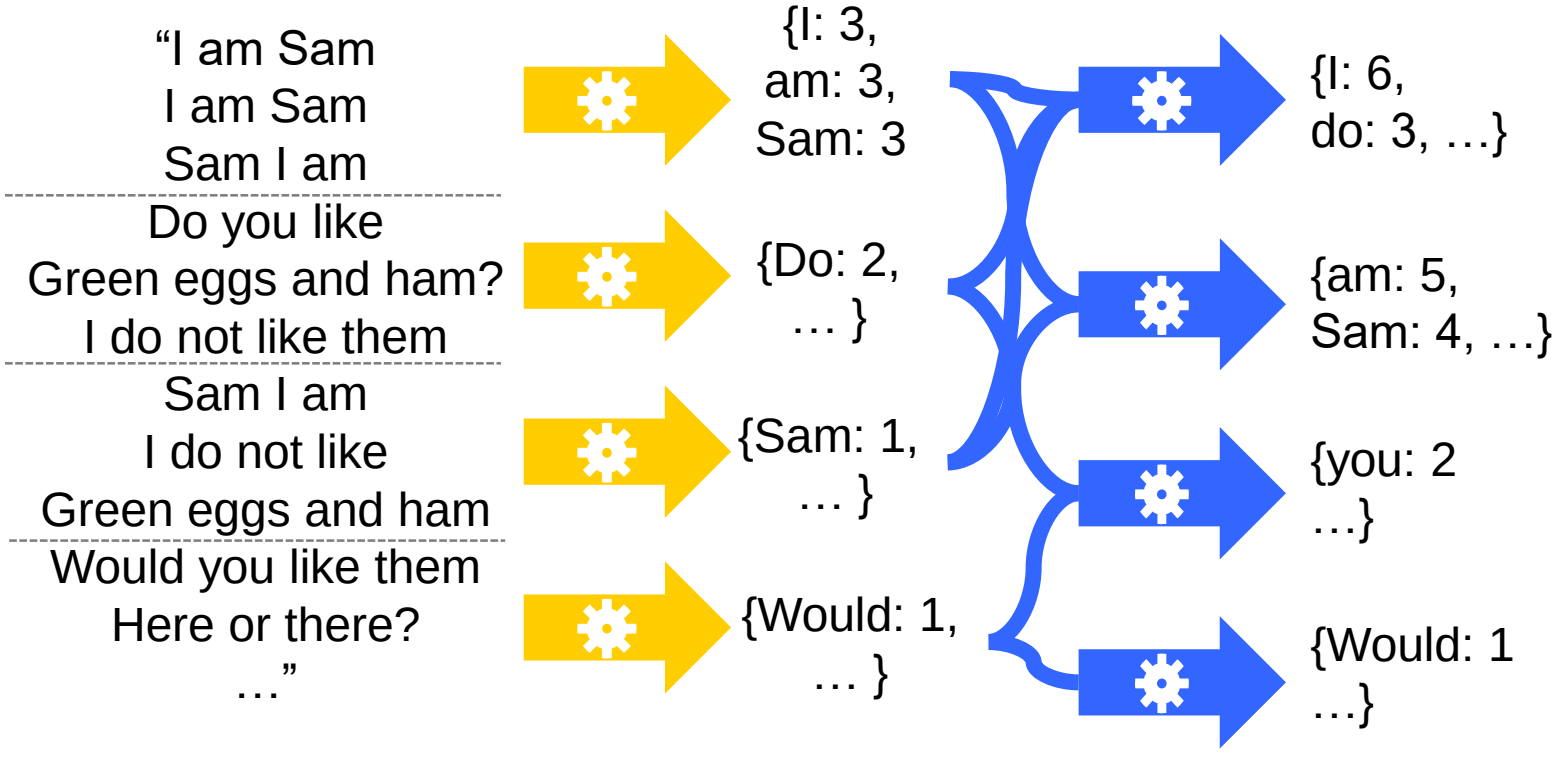

[K. Ousterhout - UCB 194-16]

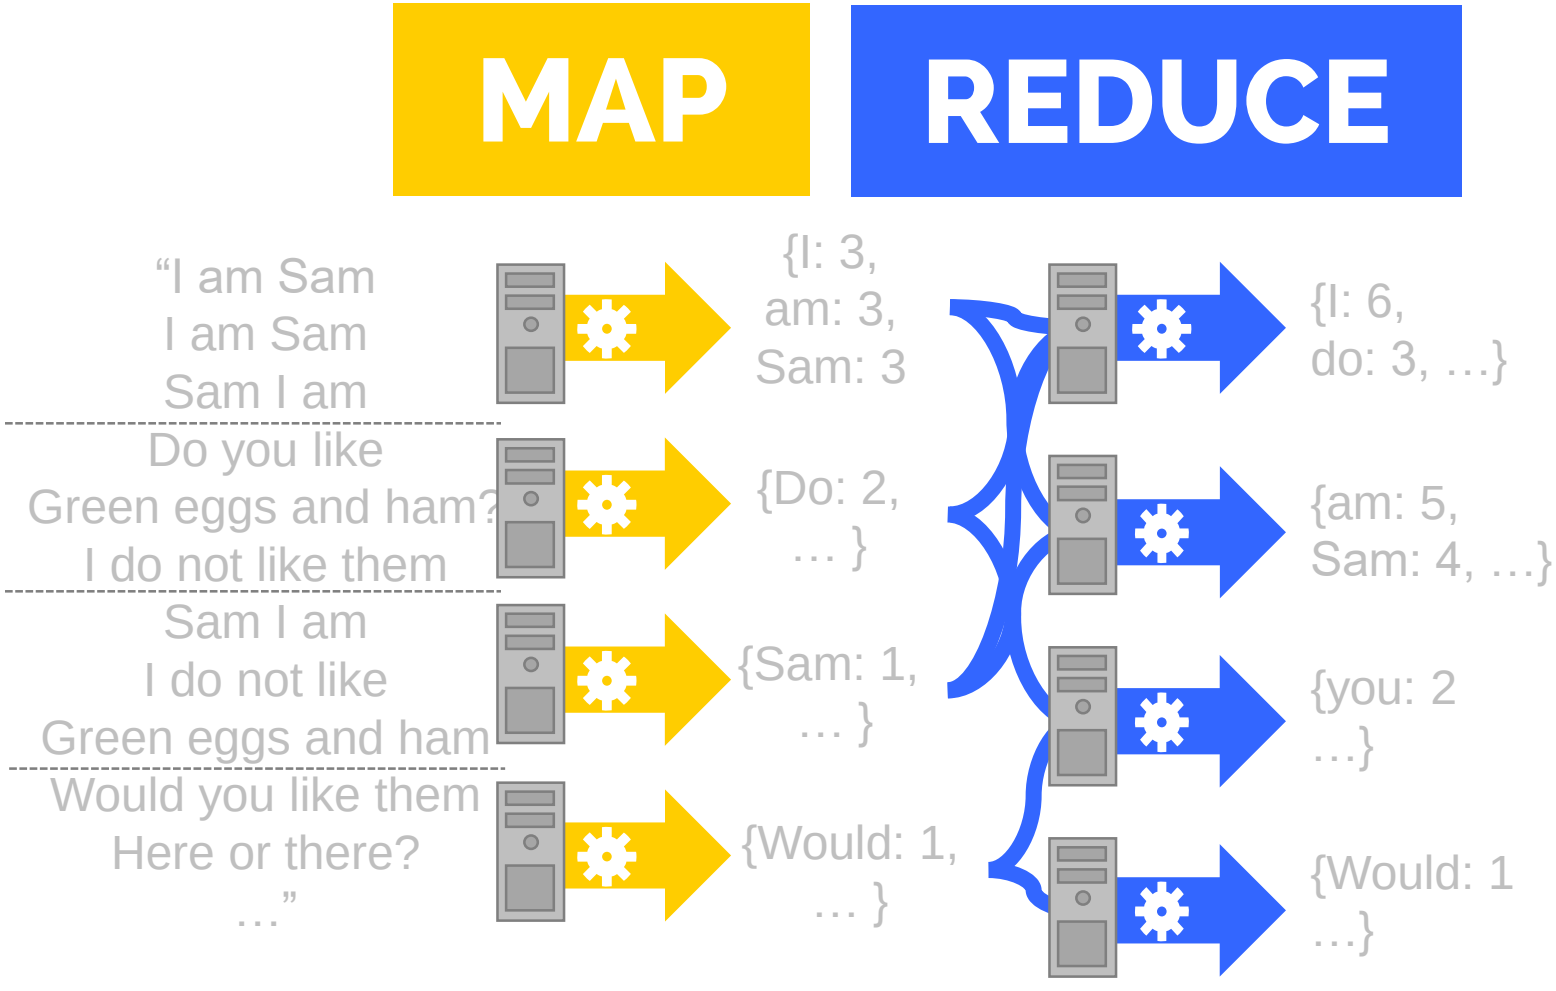

[K. Ousterhout - UCB 194-16]

MAP-REDUCE

SPLIT DATA & SEND TO MULTIPLE MACHINES (IF NOT ALREADY THERE)

MAP FILTER, SORT, AND PROCESS DATA LOCALLY

REDUCE CONSOLIDATE AND SUMMARIZE

MAP-REDUCE

CAN BE SHORT, SELF-CONTAINED FUNCTIONS

(HERE AS PYTHON-ESQUE PSEUDO CODE)

MAP

```
function Map(Document document):
    for each Word w in document:
        EmitIntermediate(w, 1)
```

REDUCE

```
function Reduce(Word w, Iterator intermediates):
    int count= 0
    for each int value in intermediates:
        count += value
    Emit(w, count)
```

MAP-REDUCE

BIG INSIGHT ISN'T
MAP / REDUCE METHODS,
BUT THEIR **SIMPLICITY**
AND THE **ARCHITECTURE**
AROUND THEM

PROVIDES **SCALABILITY**
AND **FAULT-TOLERANCE**
FOR BIG DATA
PROCESSING JOBS

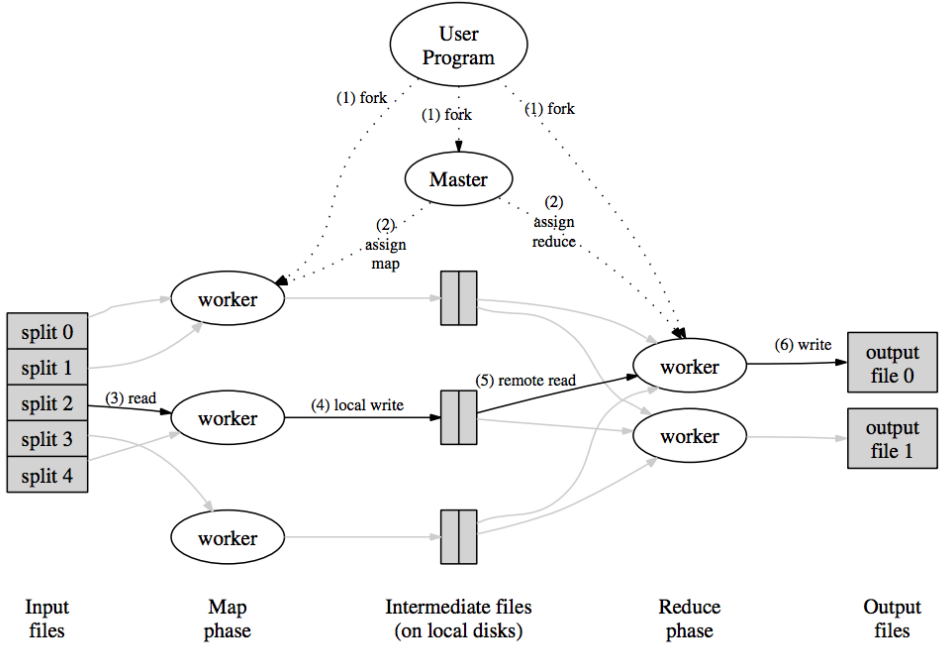

DEALING WITH ERRORS

SERVER FAILURE

1 server fails every 3 years

➔ 10K nodes see 10 faults/day

STRAGGLERS

Nodes are slow or unresponsive

[K. Ousterhout - UCB 194-16]

JUST LAUNCH A REPLACEMENT

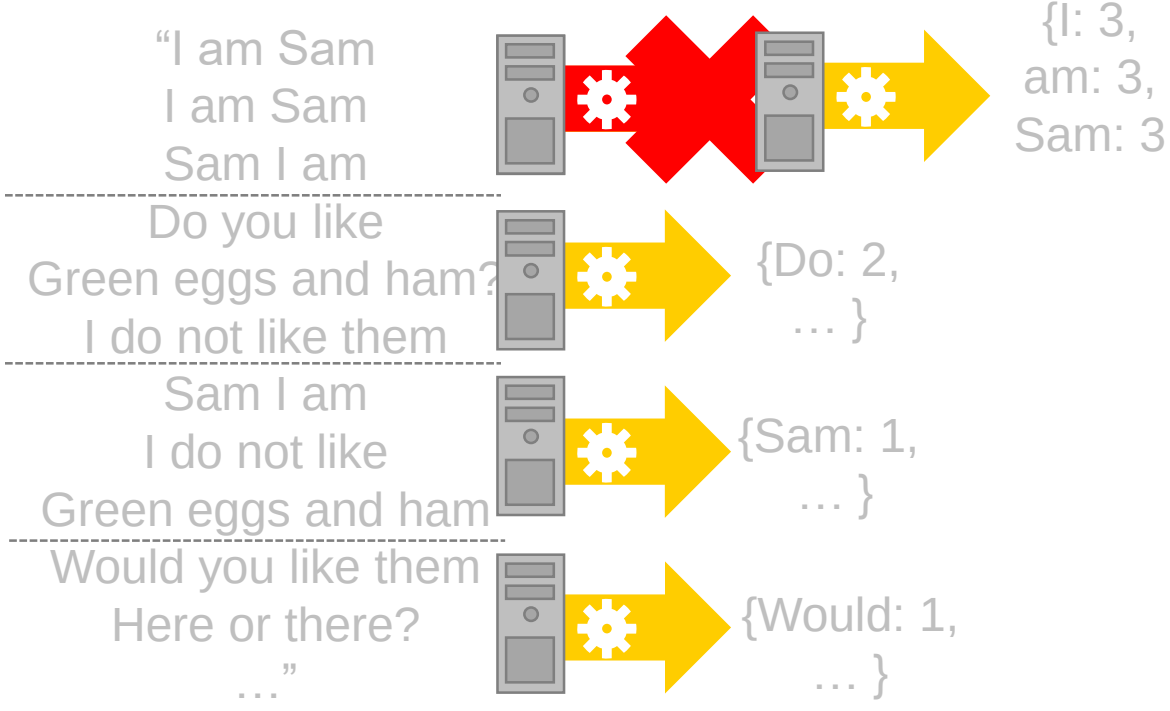

[K. Ousterhout - UCB 194-16]

APACHE HADOOP

OPEN-SOURCE DISTRIBUTED FILE SYSTEM
+ MAP REDUCE **AND MORE**

INSPIRED BY GOOGLE'S SYSTEMS

MANY DATA PROCESSING
PIPELINES NOW BUILT
ON HADOOP INFRASTRUCTURE

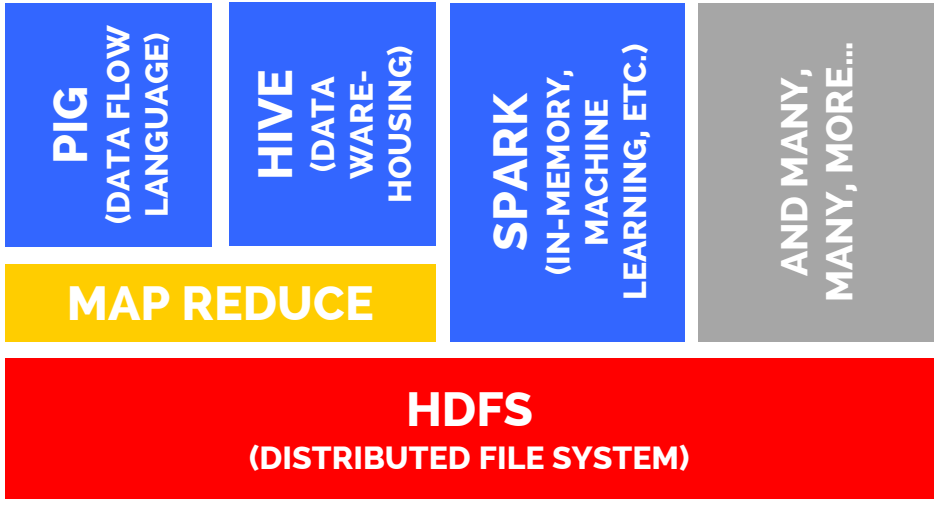

SOME OPTIONS FOR SPECIFYING BIG DATA PROCESSING OPERATIONS

WRITE YOUR OWN MAP-REDUCE METHODS

USE A QUERY LANGUAGE LIKE **APACHE PIG** THAT CAN COMPILE DOWN TO MAP REDUCE-STYLE DISTRIBUTED COMPUTATIONS

```
a = load '/documents';
b = foreach a generate flatten(TOKENIZE((chararray)$0)) as word;
c = group b by word;
d = foreach c generate COUNT(b), group;
store d into ’/pig_wordcount';
```

BENEFITS AND CHALLENGES

Data manipulation on clusters
is now a **big business.**

There is a **huge library of tools** for querying and processing distributed data.

BUT... Most of these tools are **not** real-time or interactive.

WHAT IF YOU NEED TO INTERACTIVELY EXAMINE OR VISUALIZE A BIG DATASET?

DATA ANALYSIS AT SCALE

CHALLENGES

ANALYSIS AND CLUSTER COMPUTING

INTERACTING WITH BIG DATA

PARALLELIZING HUMAN INTELLIGENCE

STRATEGIES FOR PROVIDING INTERACTIVITY WITH BIG DATA

1. INTERACTIVITY VIA PRECOMPUTATION

(AGGREGATE AND THEN INTERACT)

2. VISUALIZATION AS QUERY SPECIFICATION

(LEAVE BIG DATA ON THE SERVERS)

3. SAMPLE INTERACTIVELY

(APPROXIMATE FIRST THEN REFINE)

STRATEGIES FOR PROVIDING INTERACTIVITY WITH BIG DATA

PARALLELIZE

1. INTERACTIVITY VIA PRECOMPUTATION

(AGGREGATE AND THEN INTERACT)

2. VISUALIZATION AS QUERY SPECIFICATION

(LEAVE BIG DATA ON THE SERVERS)

3. SAMPLE INTERACTIVELY

(APPROXIMATE FIRST THEN REFINE)

SAMPLING FOR INTERACTION

STANDARD QUERY

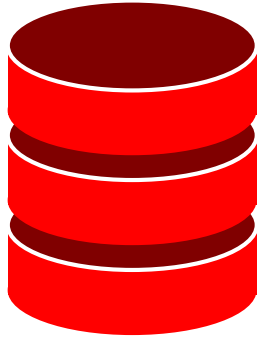

BIG DISTRIBUTED DATABASE

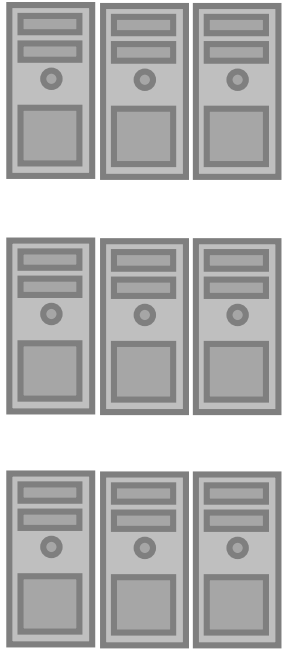

ALL RESULTS RETURNED AT THE END

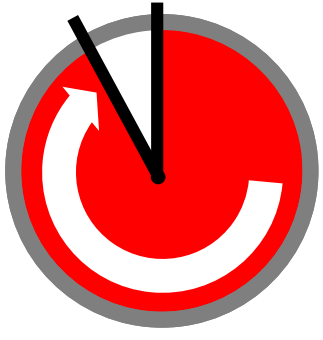

RESULTS

INTERACTIVE SAMPLING

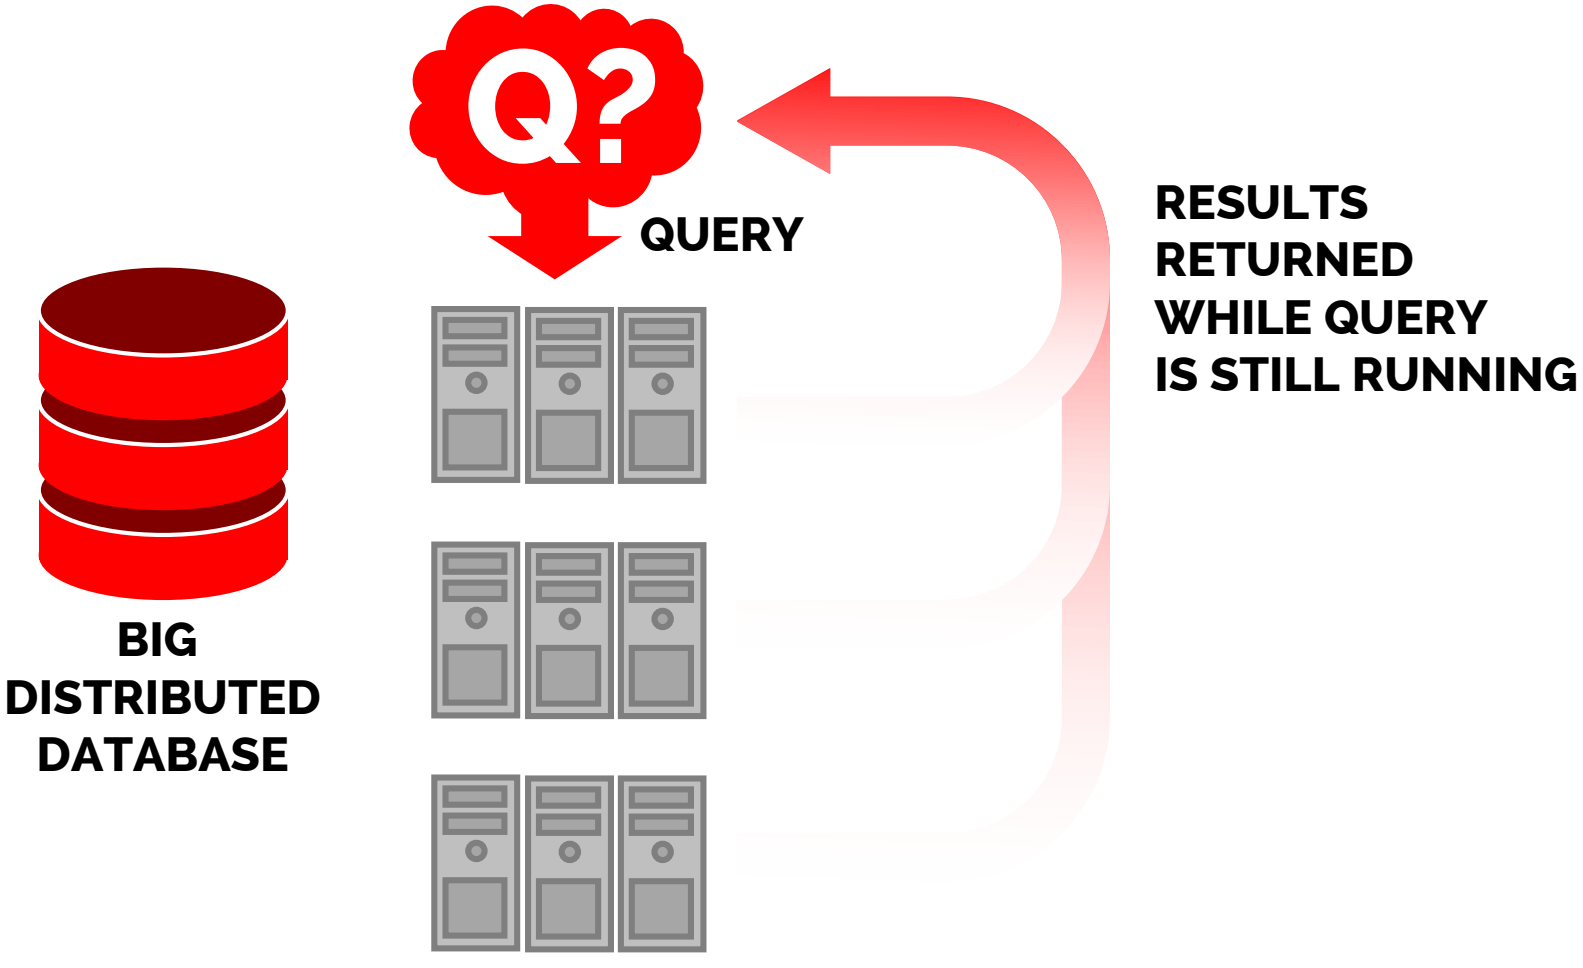

INTERACTIVE SAMPLING

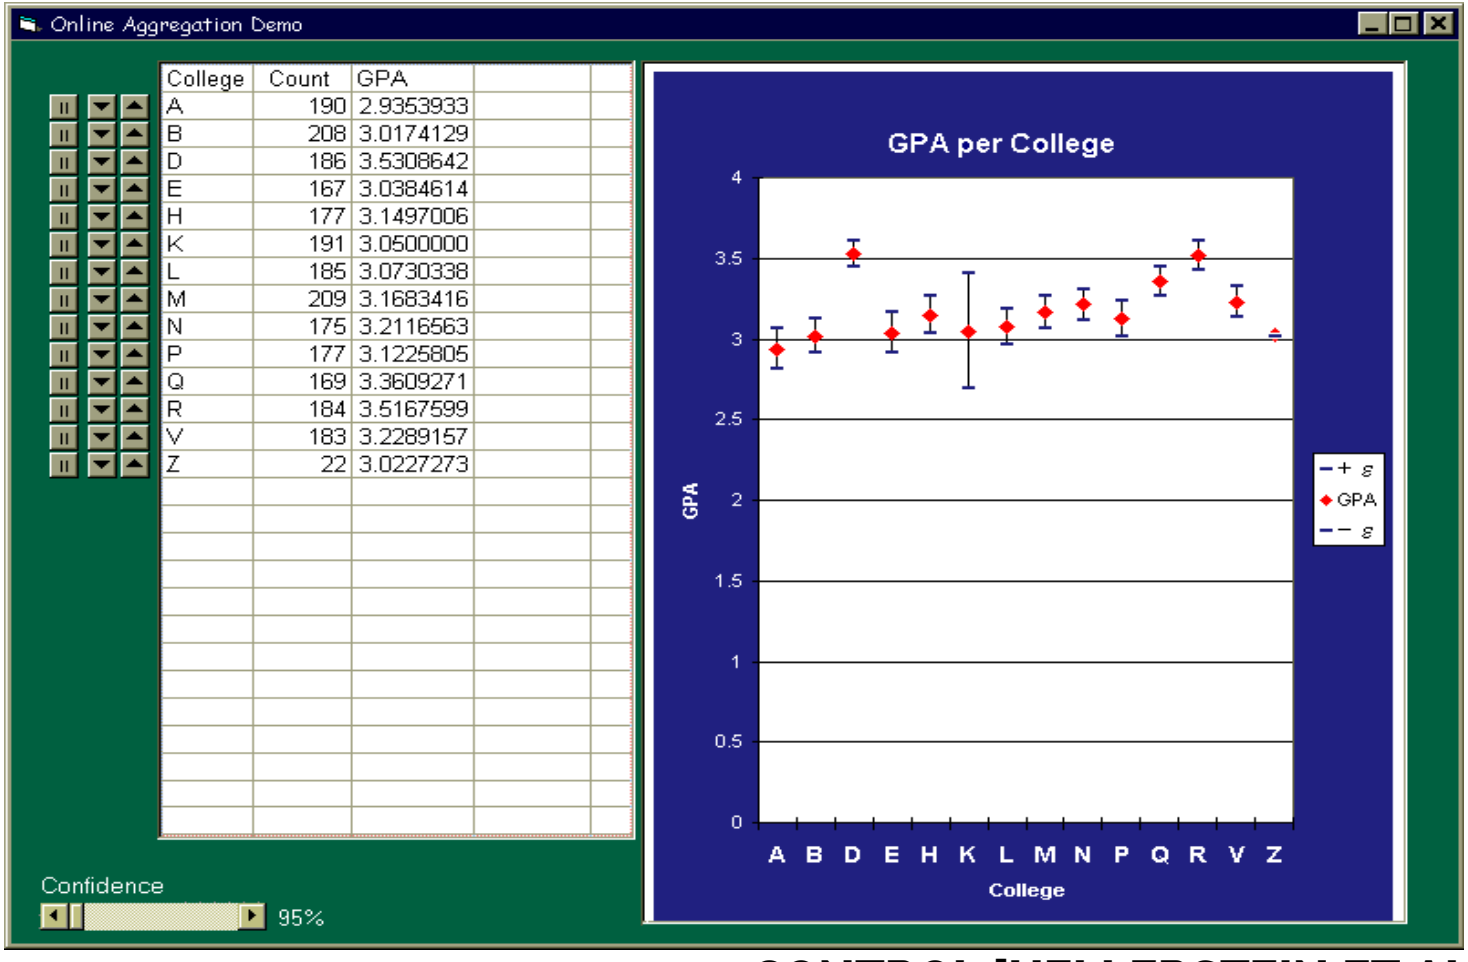

http://control.cs.berkeley.edu/papers.html

CONTROL [HELLERSTEIN ET AL. 1999]

INTERACTIVE SAMPLING

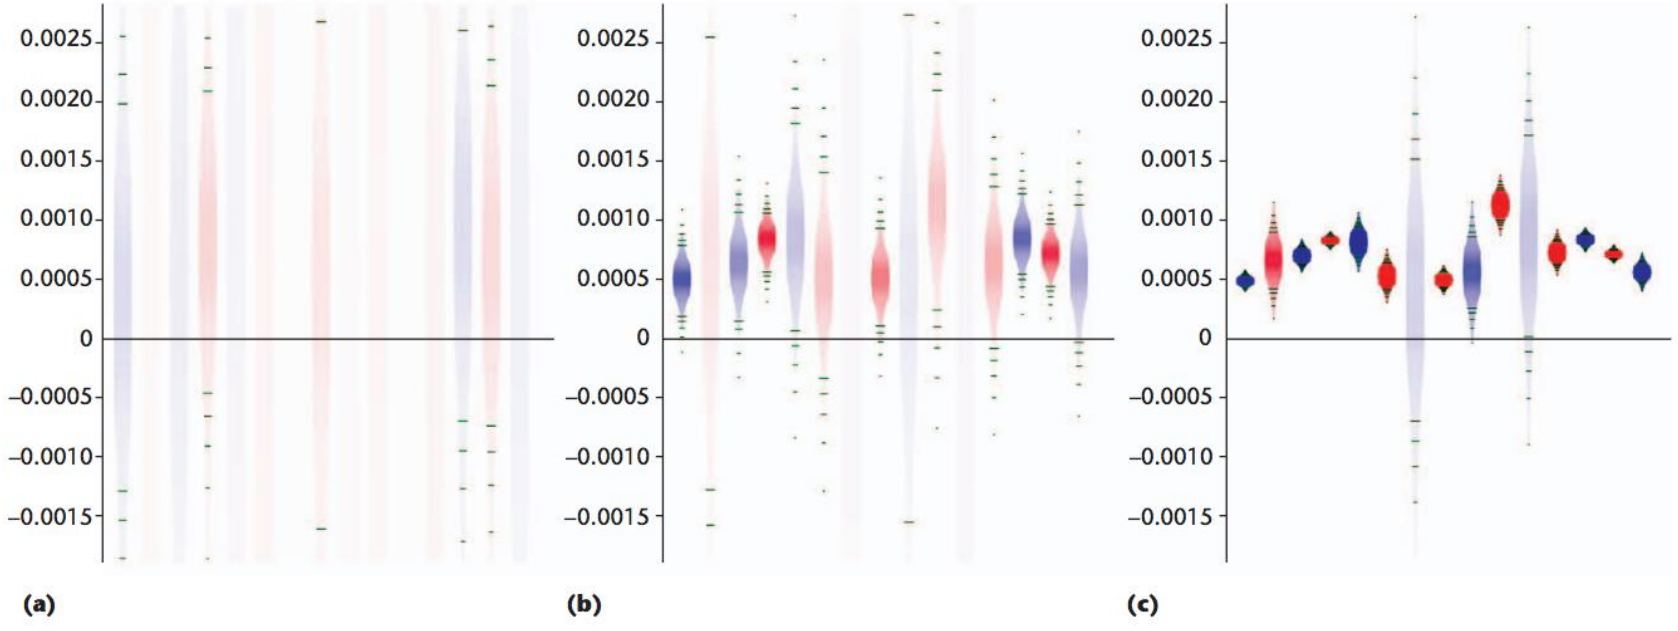

SAMPLEACTION [FISHER ET AL. 2012]

INTERACTIVE SAMPLING

BUT…

MOST BACKENDS AREN'T DESIGNED TO RETURN PROGRESSIVE RESULTS

WE NEED GOOD SAMPLING DISTRIBUTIONS FOR EACH FIELD TO PRODUCE MEANINGFUL INTERMEDIATE RESULTS

HOW BEST TO VISUALIZE UNCERTAINTY?

HOW WELL CAN PEOPLE INTERPRET PARTIAL RESULTS?

THIS IS STILL A VERY OPEN RESEARCH AREA!

HOW TO SHOW UNCERTAINTY?

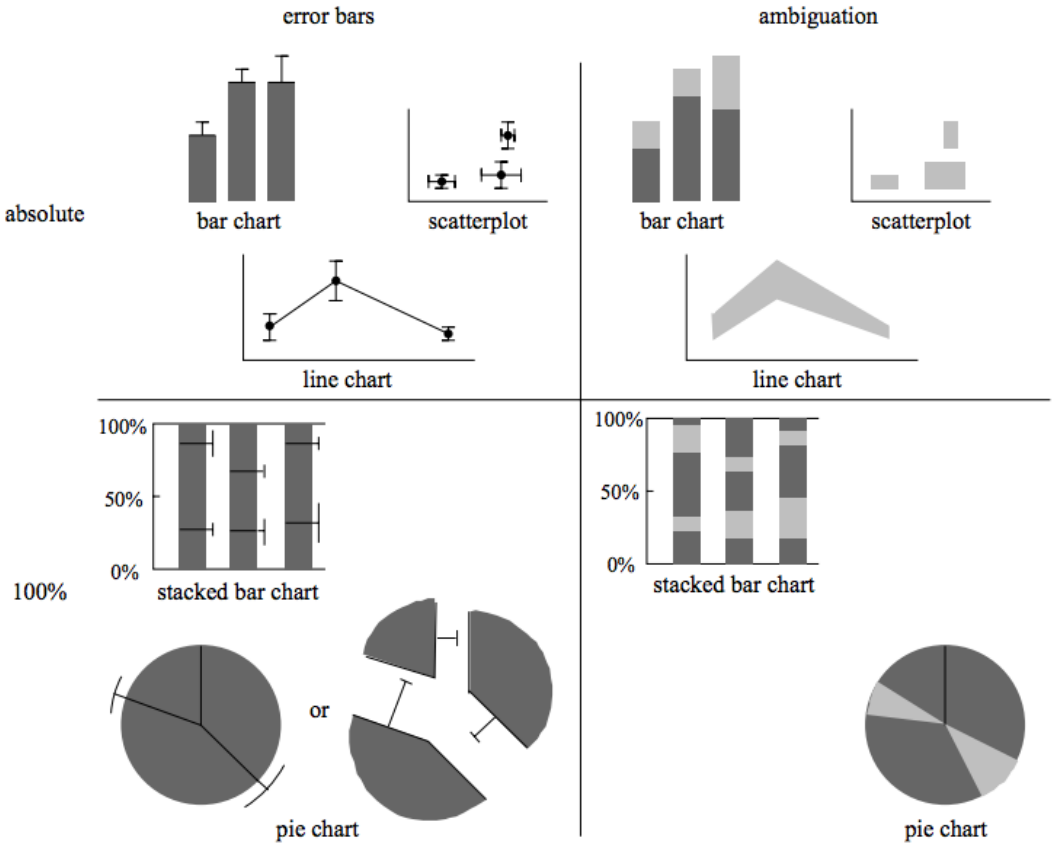

Figure 1: Error bars and ambiguation applied to some common chart types.

[Olston & Mackinlay, 2002]

HOW TO SHOW UNCERTAINTY?

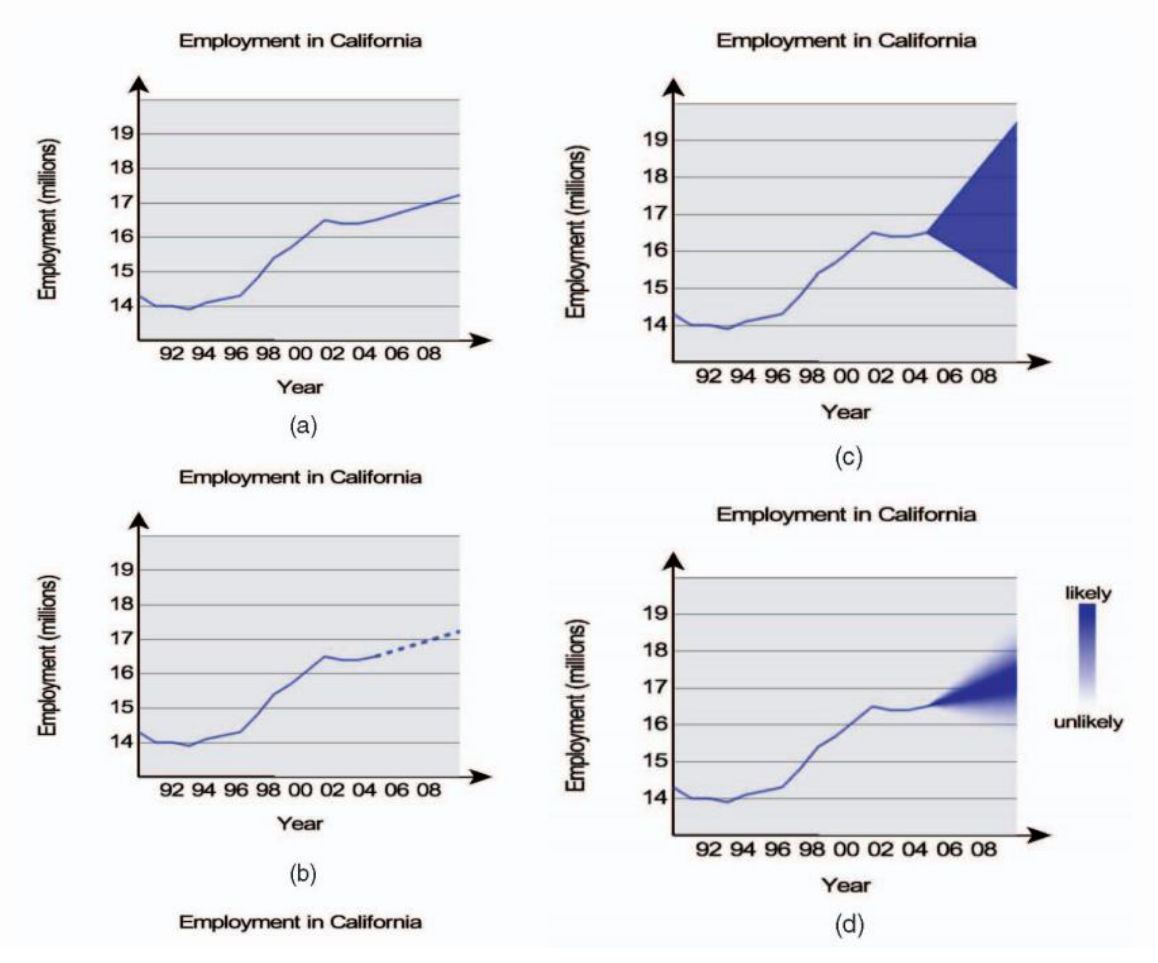

[Streit, Pham, & Brown 2008]

HOW TO SHOW UNCERTAINTY?

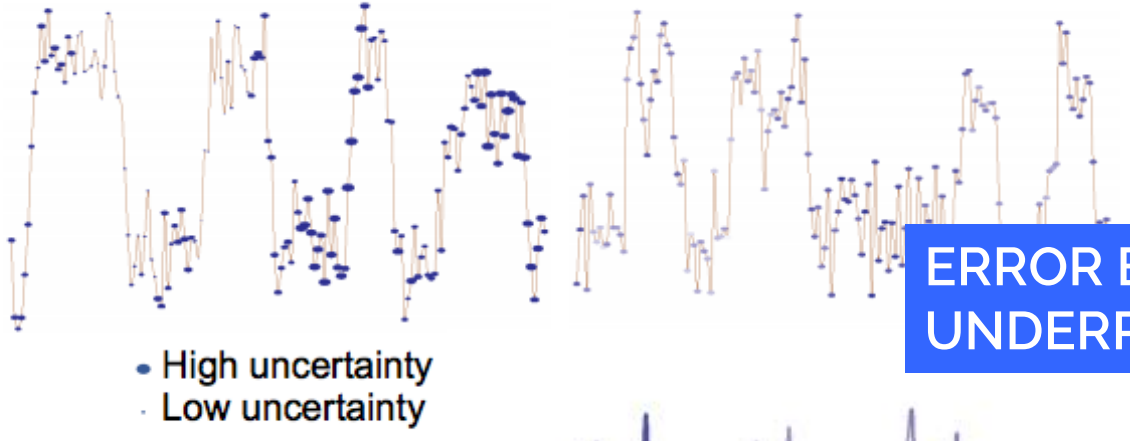

a.

ERROR BARS CONSISTENTLY UNDERPERFORMED

[Sanyal, et al. 2009]

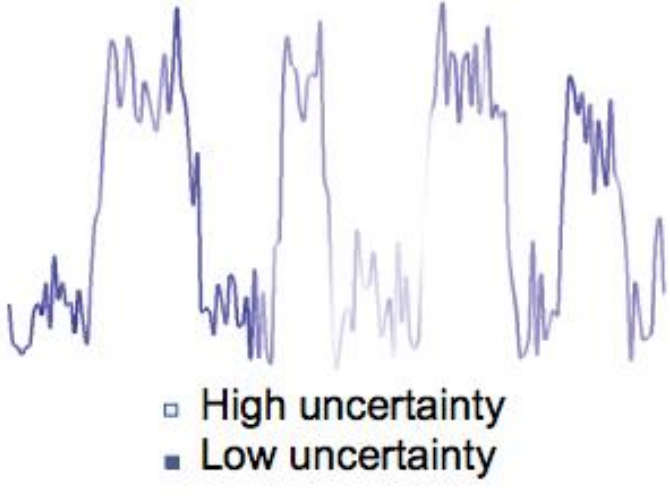

c.

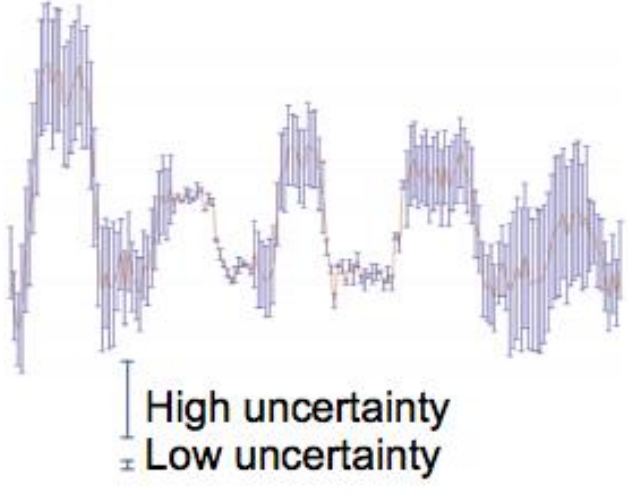

d.

HOW TO SHOW UNCERTAINTY?

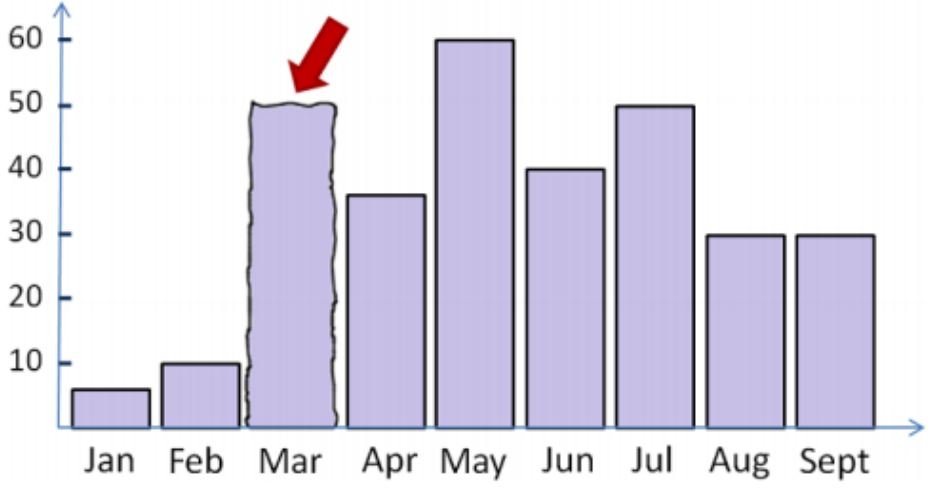

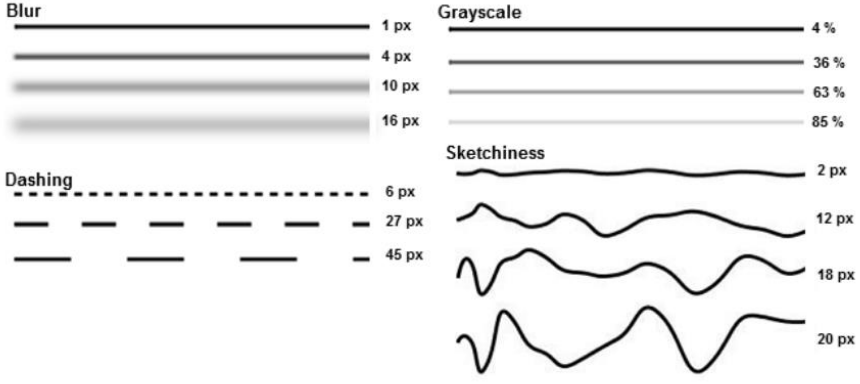

[Boukhelifa, et al. 2012]

PEOPLE DON'T ALWAYS INTERPRET THESE AS SHOWING UNCERTAINTY

A FEW INTERESTING RESEARCH PROTOTYPES

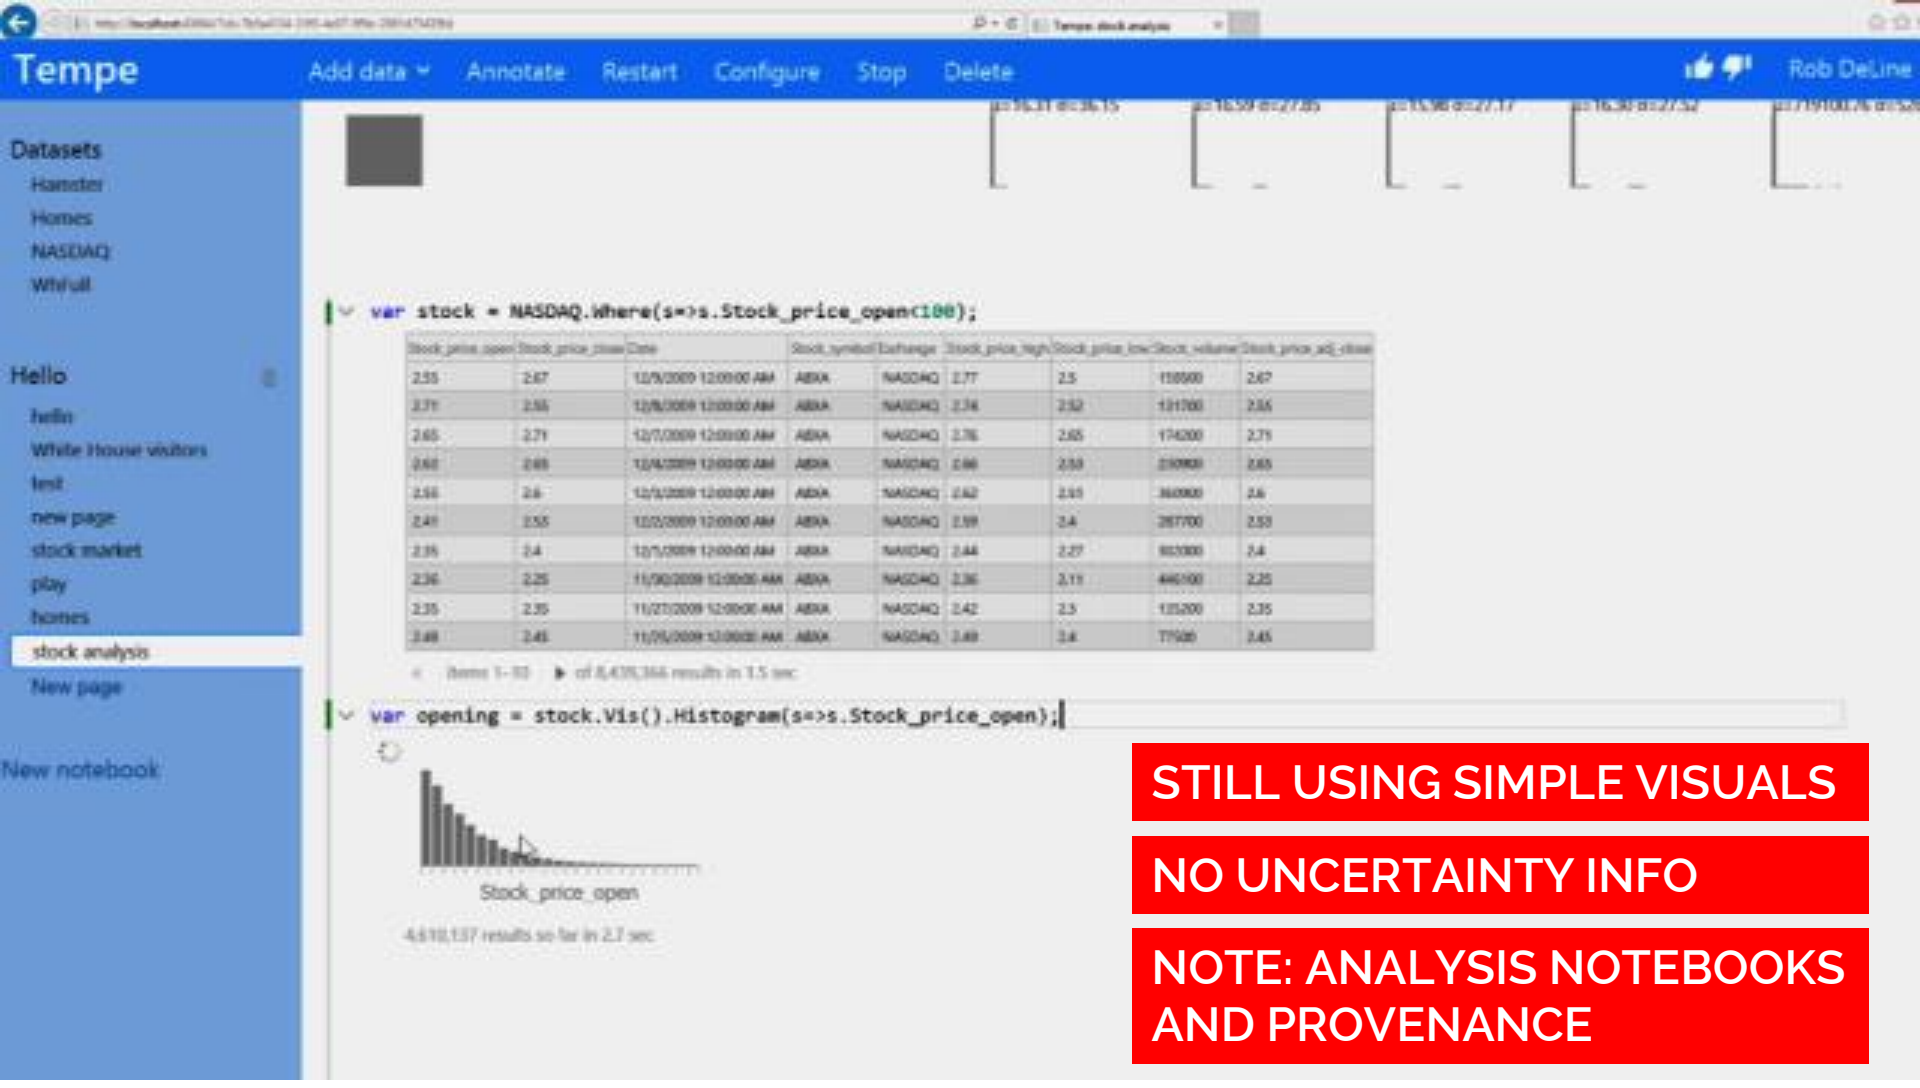

TEMPE [Microsoft Research 2014]

DATA ANALYSIS AT SCALE

CHALLENGES

ANALYSIS AND CLUSTER COMPUTING

INTERACTING WITH BIG DATA

PARALLELIZING HUMAN INTELLIGENCE

HOW CAN WE LEVERAGE MULTIPLE PEOPLE TO EXPEDITE ANALYSIS?

Analyst

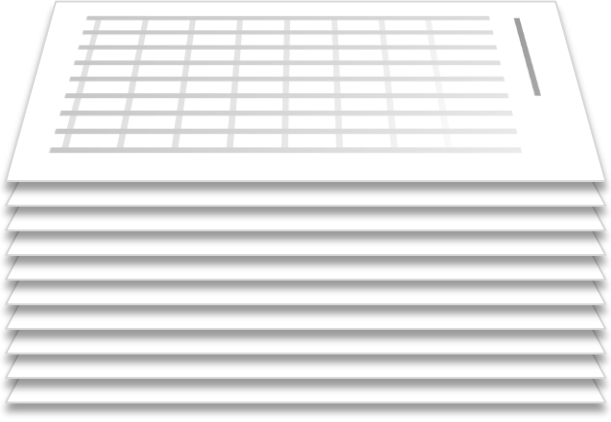

CollegeRankings2013.csv

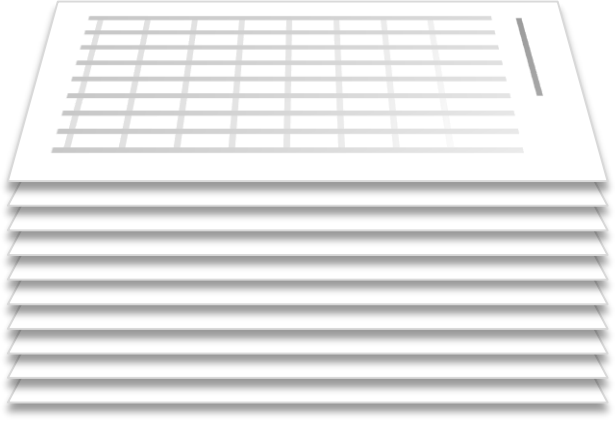

Analyst

"Can I enlist others to help make sense of my data?"

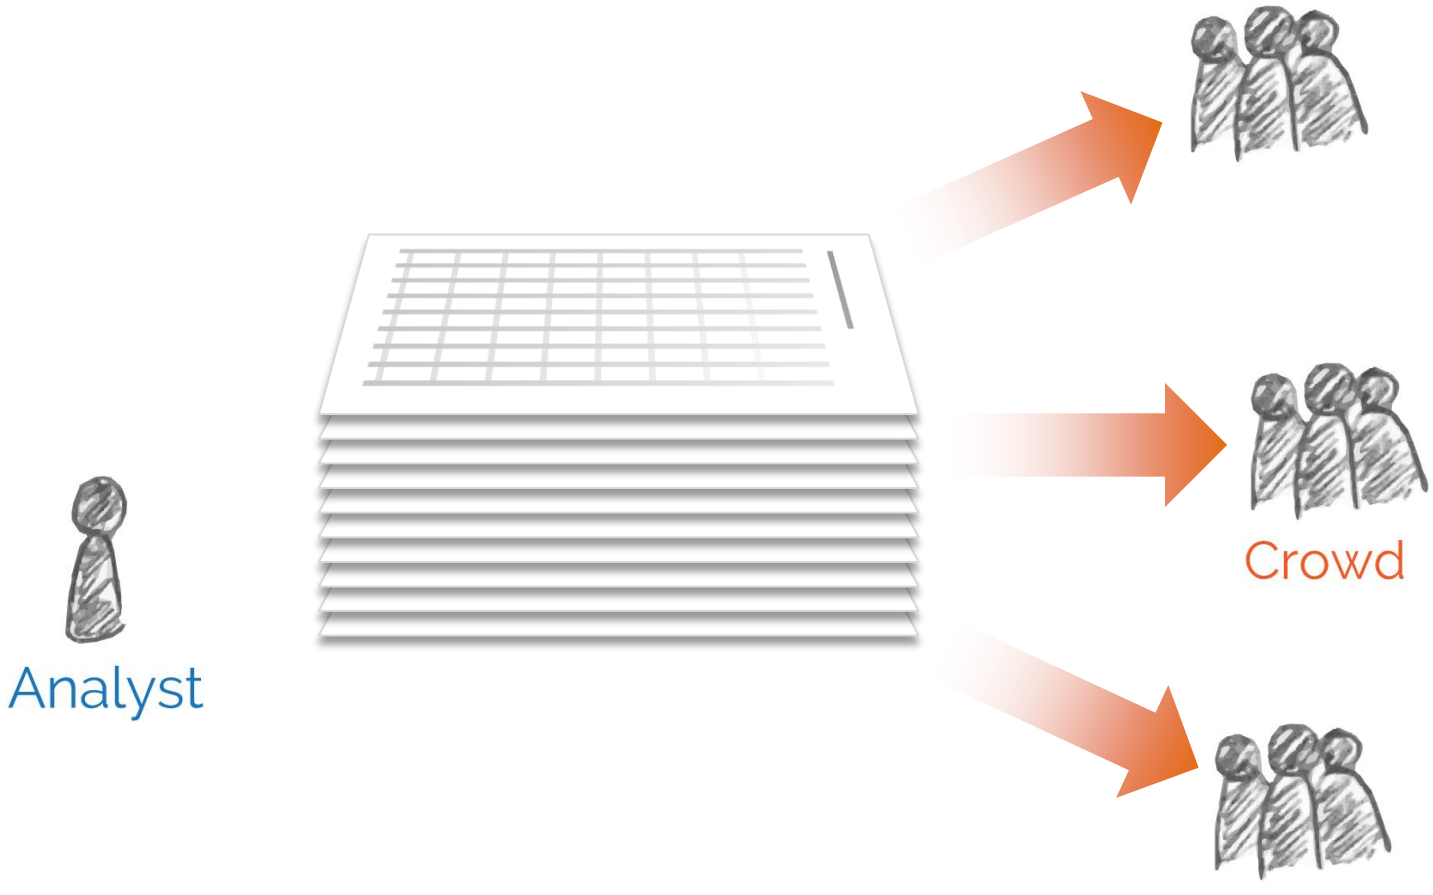
Analyst
Crowd

MANY IMPORTANT ANALYSIS TASKS REQUIRE HUMAN INTELLIGENCE BUT LEND THEMSELVES WELL TO PARALLELIZATION

MANY IMPORTANT ANALYSIS TASKS REQUIRE HUMAN INTELLIGENCE BUT LEND THEMSELVES WELL TO PARALLELIZATION

Sensemaking Loop

Foraging Loop

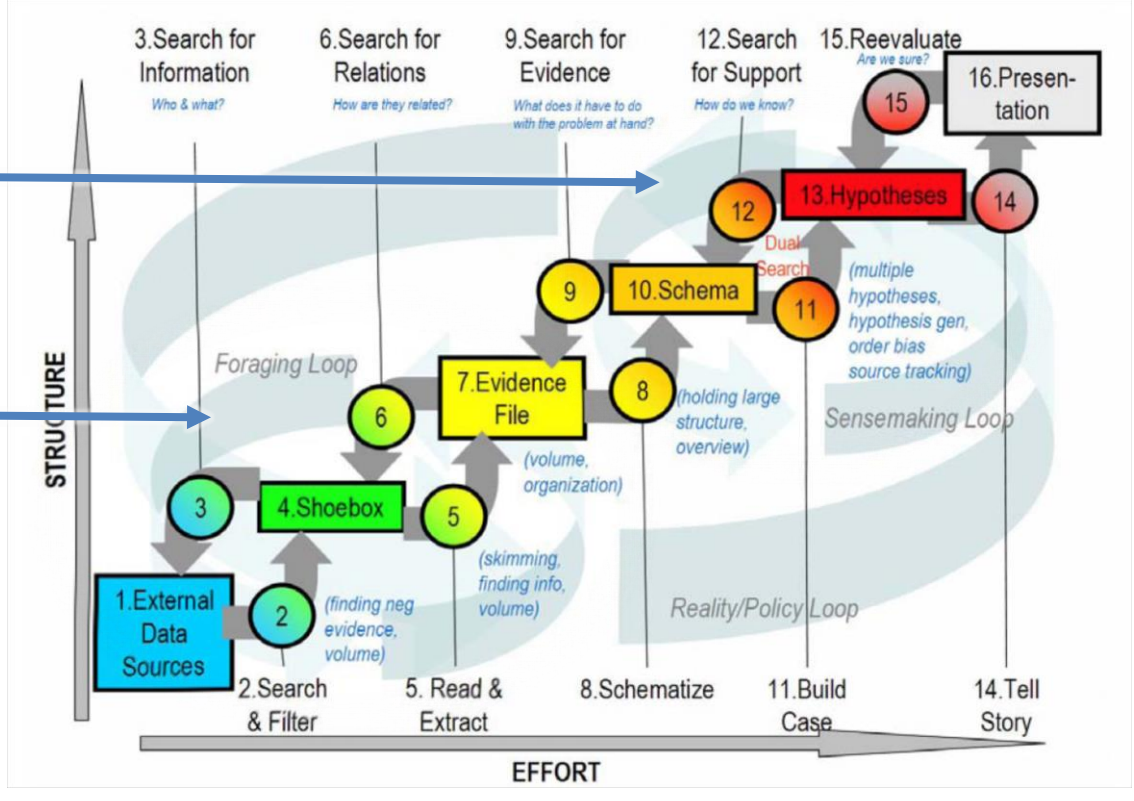

[Pirolli & Card 2005]

MANY EYES

[Viégas, et al. 2007, 2008]

GOOGLE BOOKS N-GRAMS

CROWDSOURCING DATA ANALYSIS

DATA COLLECTION & CITIZEN SCIENCE

ANALYSIS COMPETITIONS

"MICROWORK" AND TASK MARKETS

COLLABORATION TOOLS FOR ANALYSTS

CITIZEN SCIENCE

DATA COLLECTION

CREEK WATCH
[IBM]

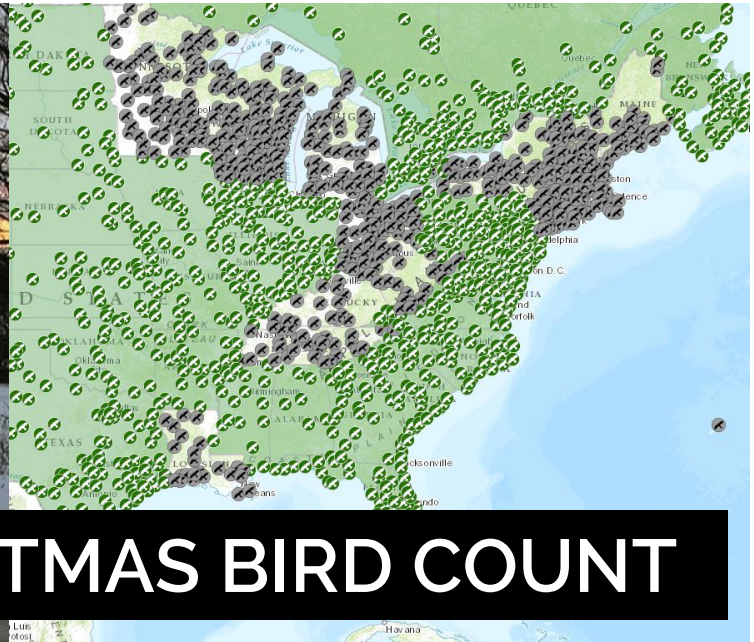

CHRISTMAS BIRD COUNT

CITIZEN SCIENCE

DATA PROCESSING

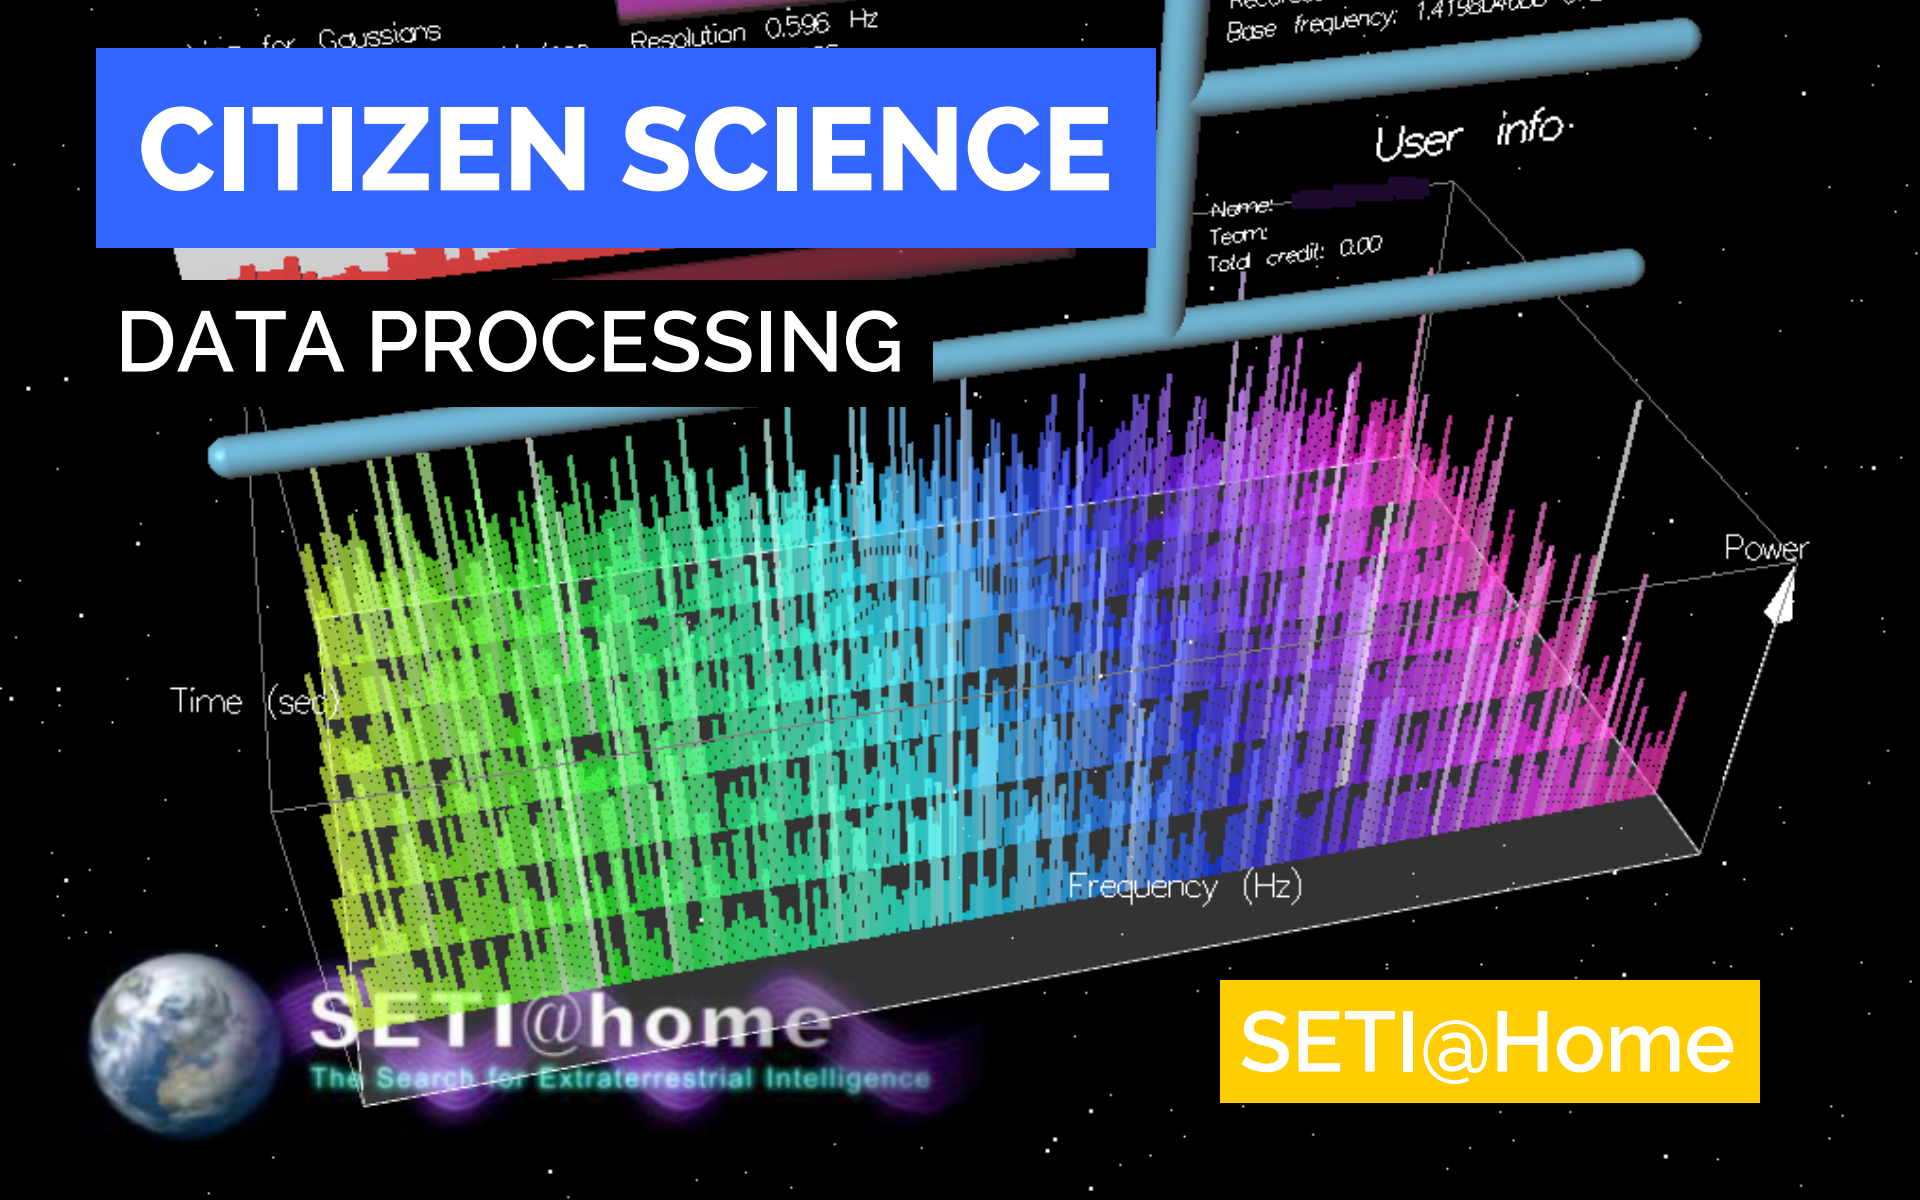

SETI@Home

CITIZEN SCIENCE

HUMAN VISION & PROBLEM SOLVING

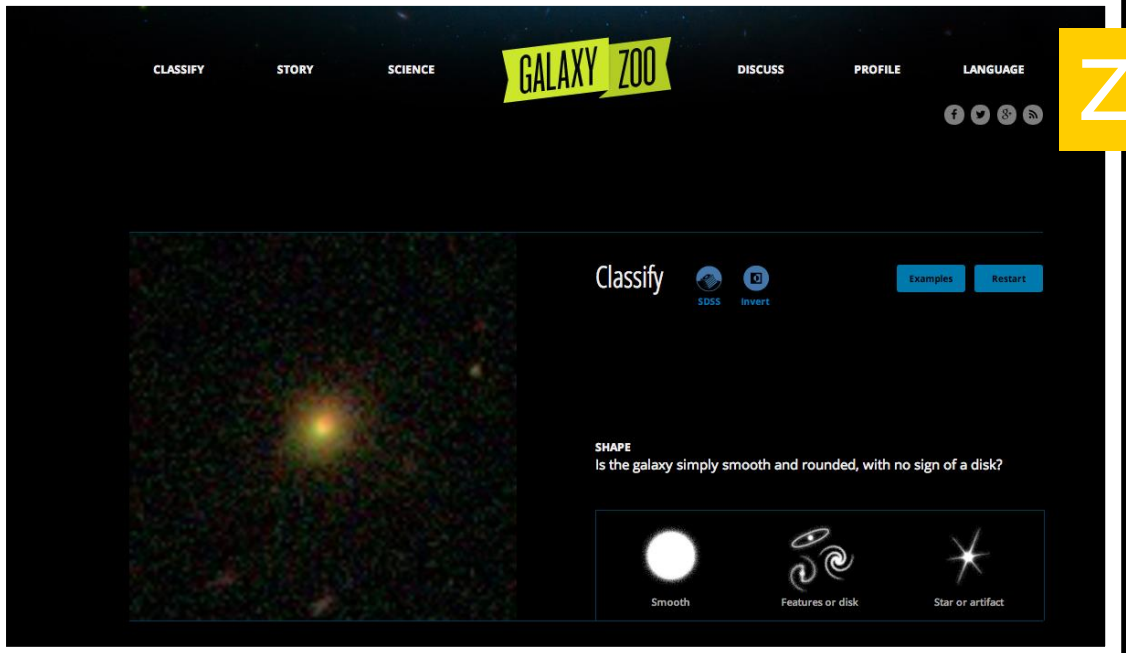

ZOONIVERSE

CITIZEN SCIENCE

HUMAN VISION & PROBLEM SOLVING

FOLD.IT

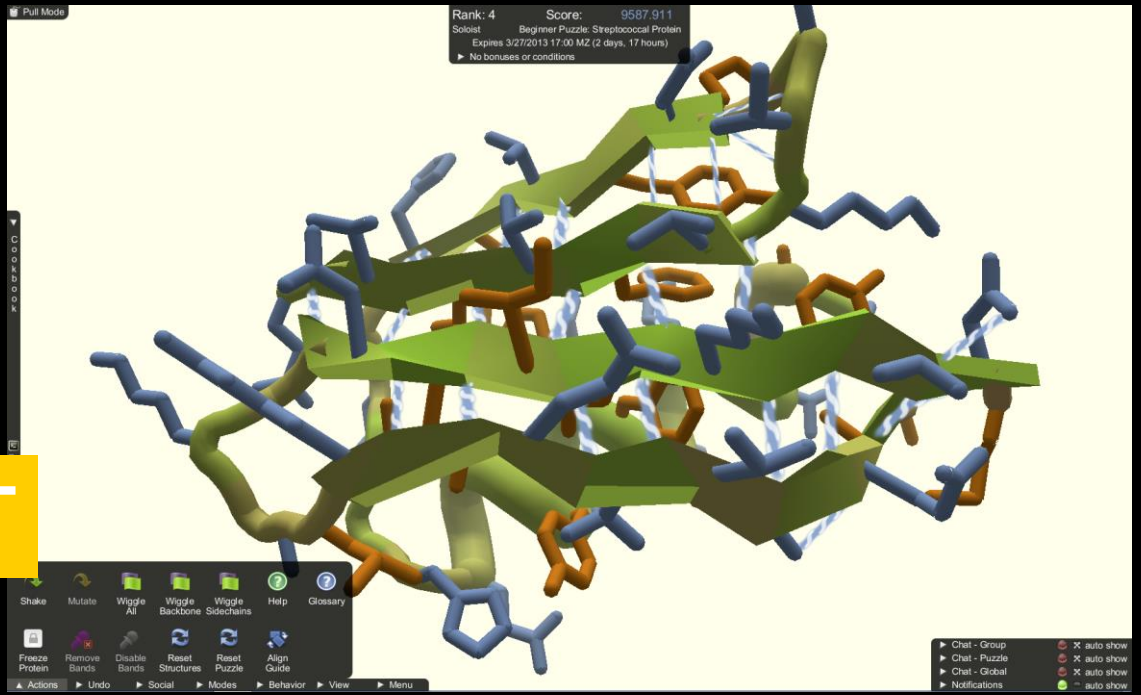

ANALYSIS COMPETITIONS

NETFLIX PRIZE

KAGGLE

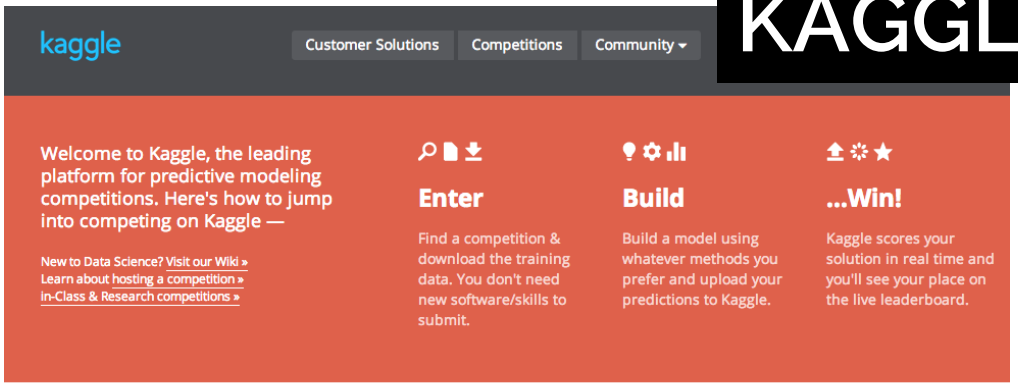

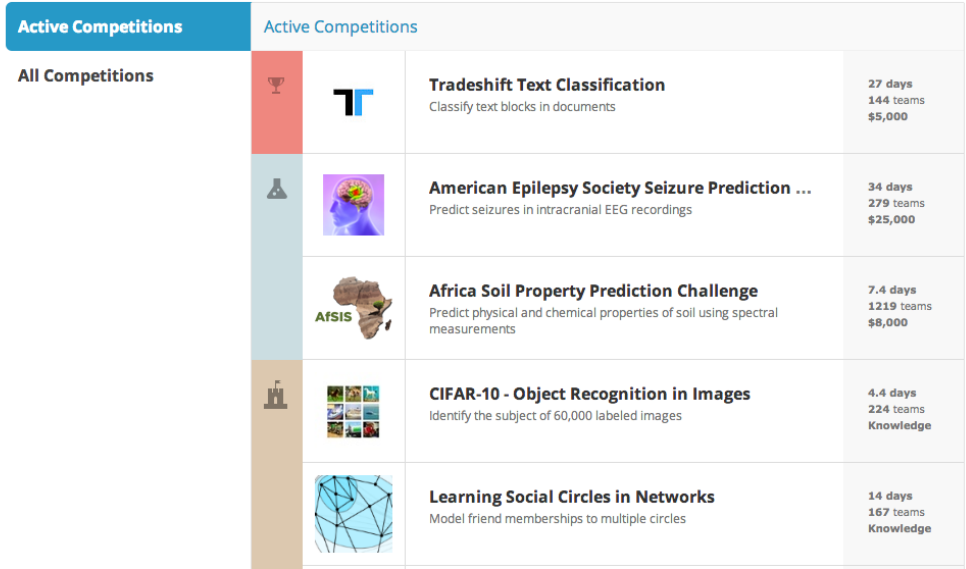

MICROWORK PLATFORMS

SITES WHERE WORKERS PERFORM SMALL PIECES OF WORK ("TASKS") - USUALLY IN EXCHANGE FOR SMALL FINANCIAL REWARDS.

amazon mechanical turk™
Artificial Artificial Intelligence

CrowdFlower

mobileworks

MICROWORK

USING APIS – DEVELOPERS CAN WRITE PROGRAMS THAT INCORPORATE HUMAN JUDGEMENT

"HUMAN COMPUTATION"

APPLYING MICROWORK TO DATA ANALYSIS

CROWDSOURCING LOW-LEVEL ANALYSIS

DATA COLLECTION AND DATA ENTRY

LABELING

DATA CLEANING

SENTIMENT ANALYSIS

MANY IMPORTANT ANALYSIS TASKS REQUIRE HUMAN INTELLIGENCE BUT LEND THEMSELVES WELL TO PARALLELIZATION

Sensemaking Loop

Foraging Loop

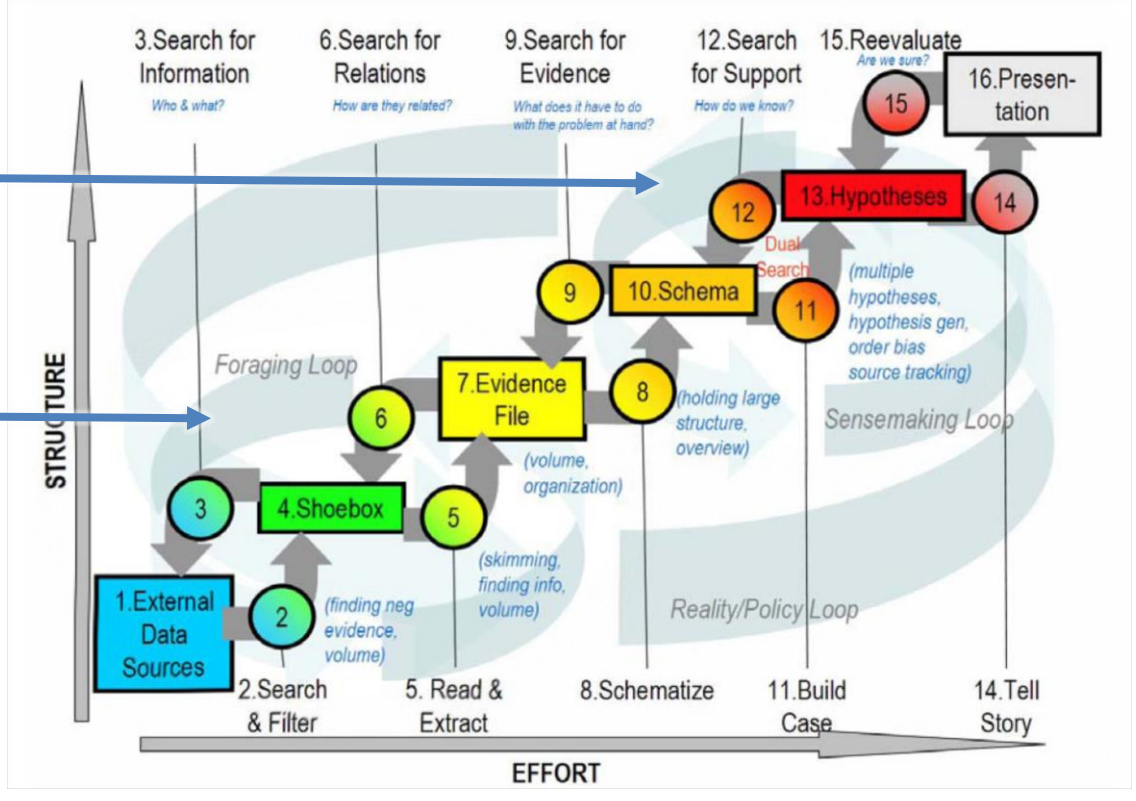

[Pirolli & Card 2005]

CROWDSOURCING HIGHER-LEVEL ANALYSIS TASKS

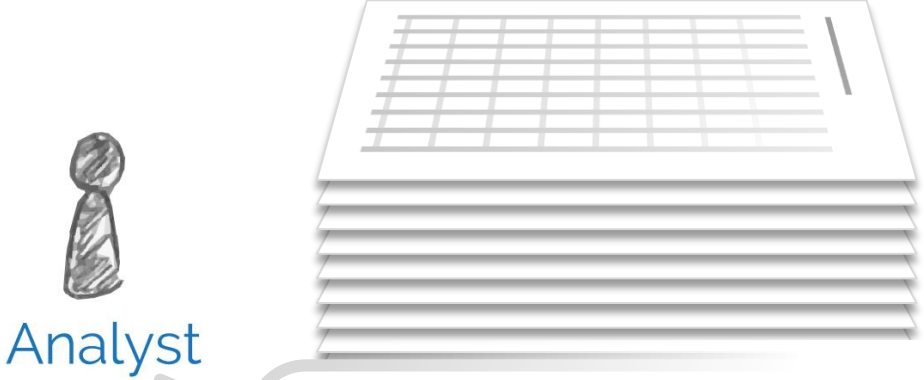

A WORKFLOW FOR CROWDSOURCING DATA ANALYSIS

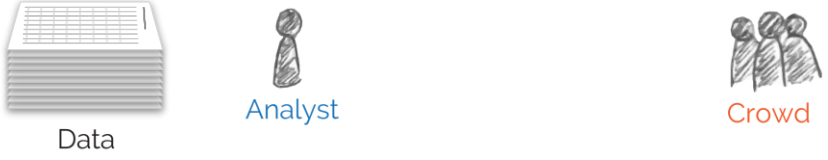

[Willett et al. CHI 2012, VAST 2013]

A WORKFLOW FOR CROWDSOURCING DATA ANALYSIS

Data

Analyst

Crowd

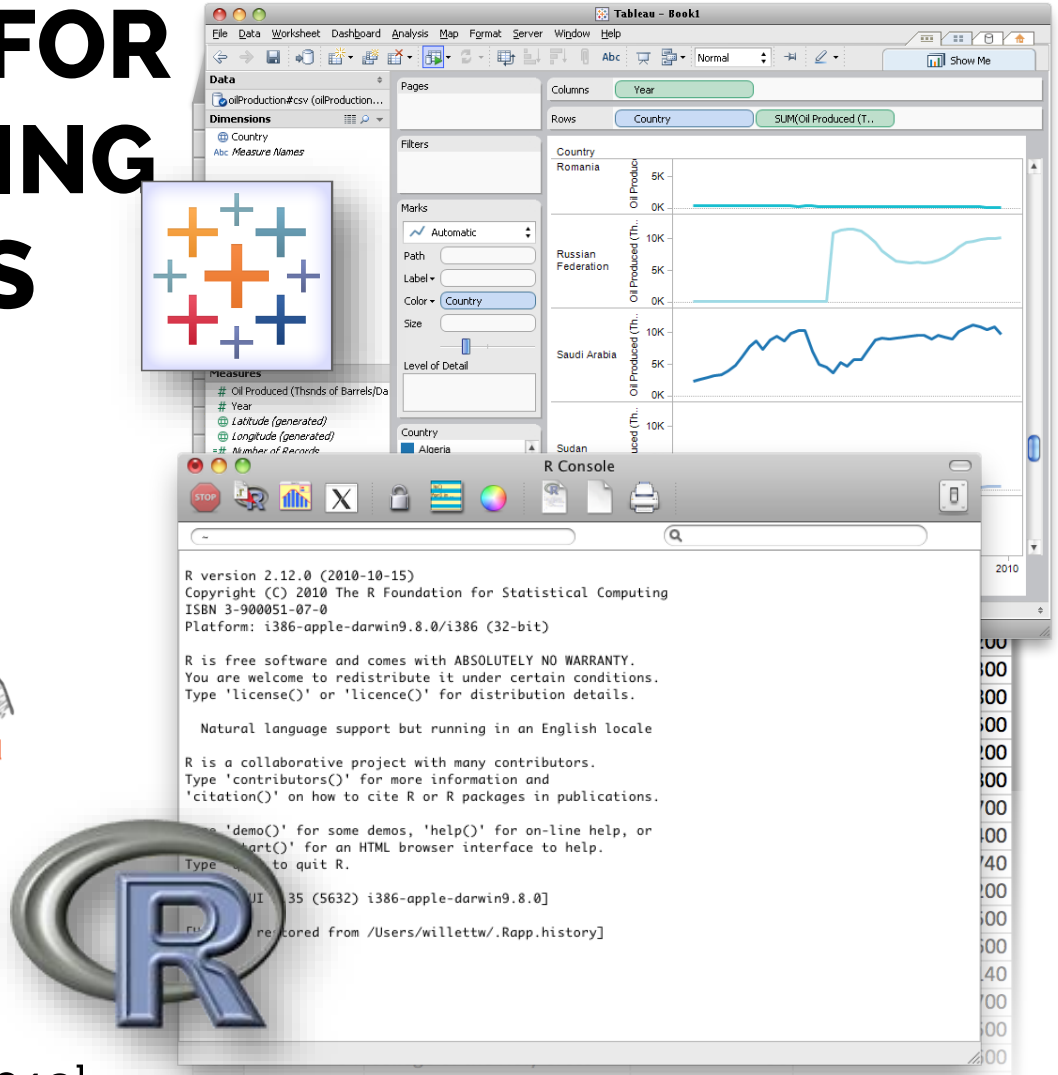

[Willett et al. CHI 2012, VAST 2013]

A WORKFLOW FOR CROWDSOURCING DATA ANALYSIS

Data

Analyst

Crowd

[Willett et al. CHI 2012, VAST 2013]

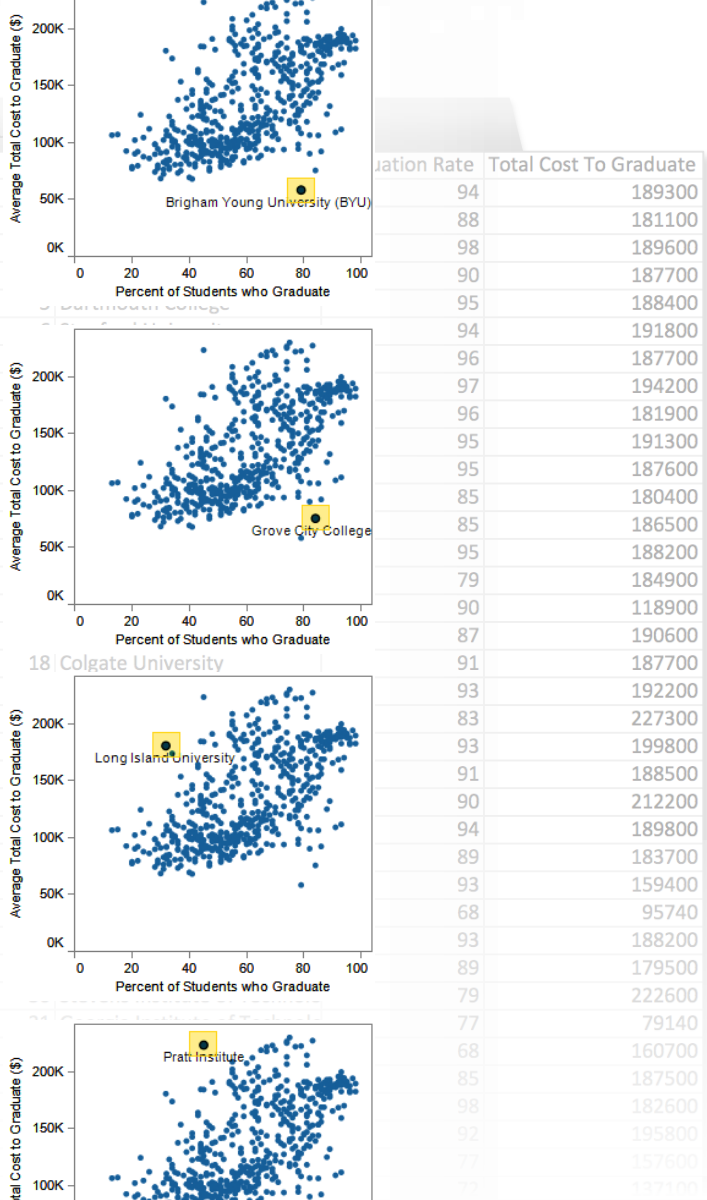

A WORKFLOW FOR CROWDSOURCING DATA ANALYSIS

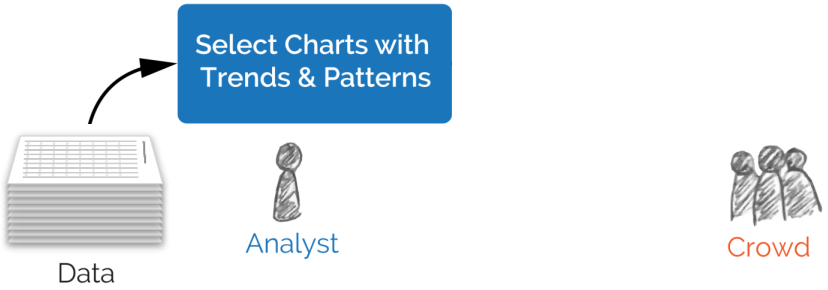

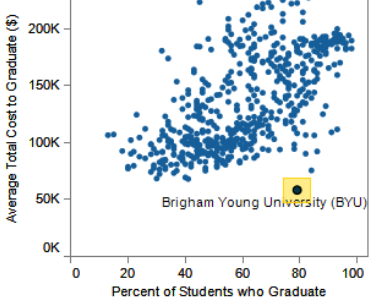

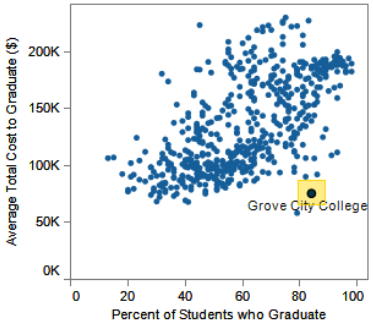

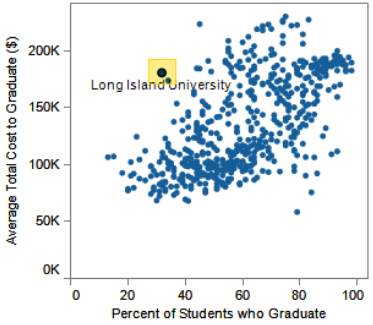

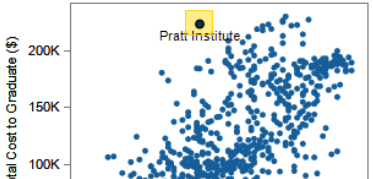

[Willett et al. CHI 2012, VAST 2013]

Explanation Task

proxy.commentspace.net/explainTask?studyName=Oil-Demo&assignmentId=TestAI&workerId=Exper...

Each of the charts in this HIT shows the **average amount of oil produced per day** by one or more countries over the past 50 years

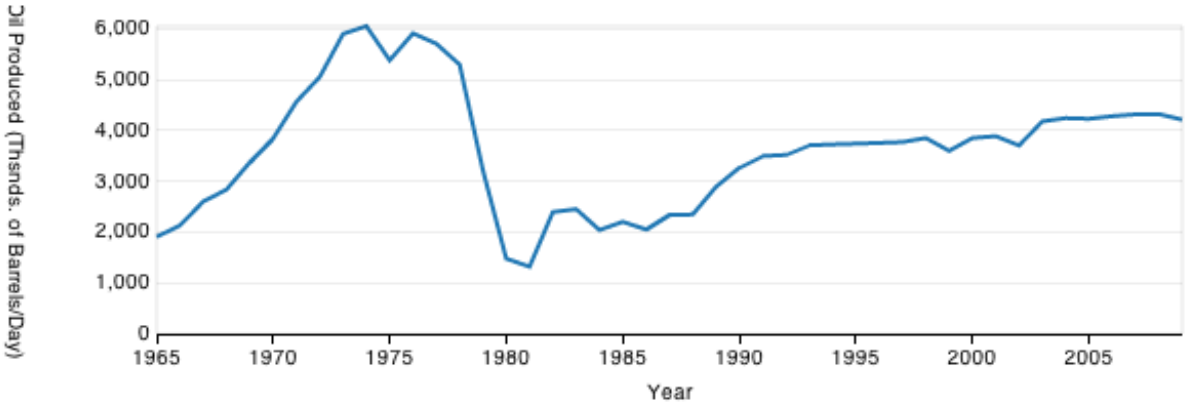

This chart shows **Oil Produced (Thsnds. of Barrels/Day)** by **Year**. The view is filtered by **Country** to show only **"Iran"**.

1. Explain **why** the strong ***peak or valley*** highlighted in the chart might have occurred.

Submit Task

http://proxy.commentspace.net/media/html/Combinator/Combinator.html?

A WORKFLOW FOR CROWDSOURCING DATA ANALYSIS

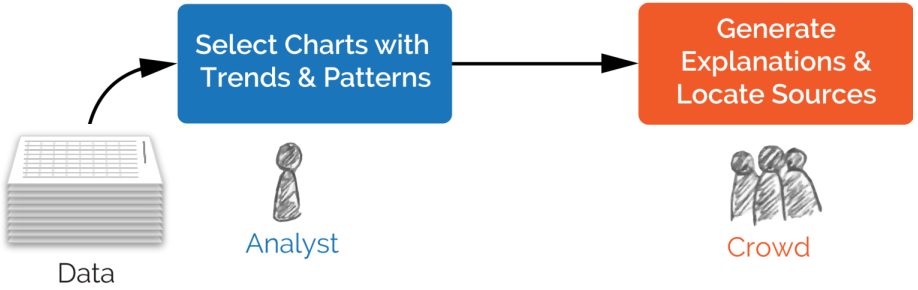

[Willett et al. CHI 2012, VAST 2013]

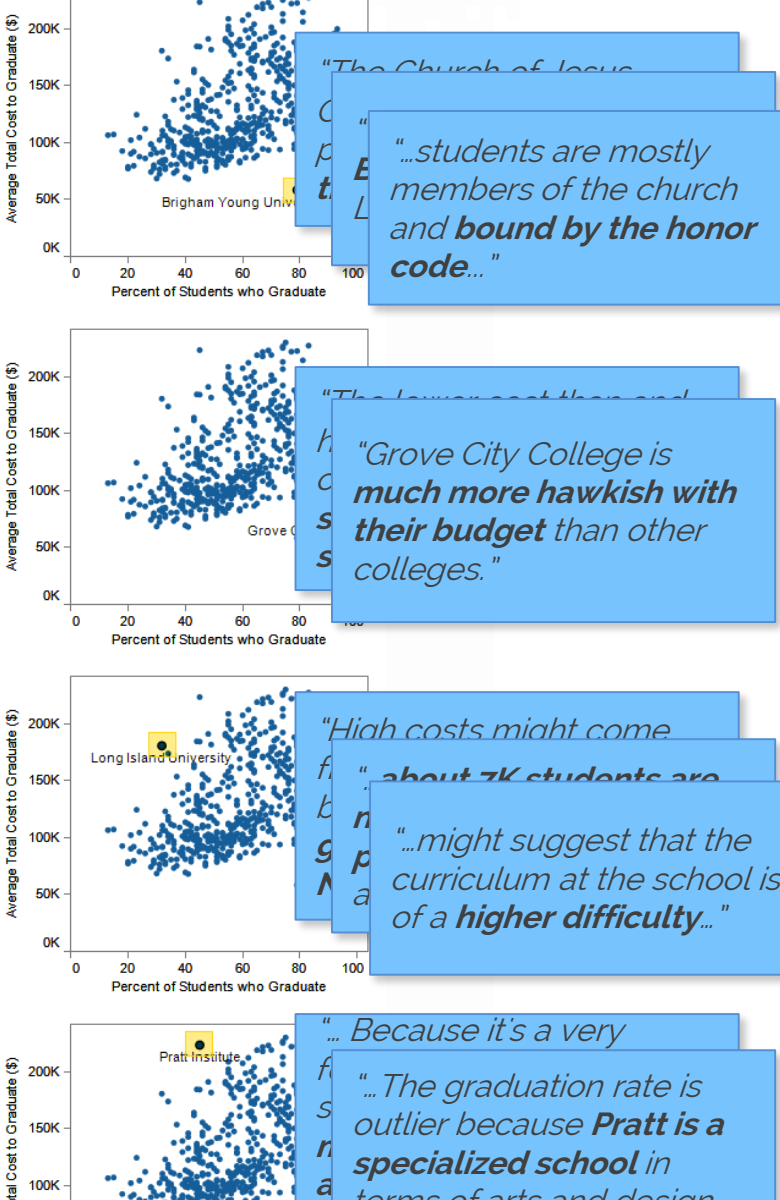

“COULD THIS CREATE MORE WORK FOR THE ANALYST?”

"COULD THIS CREATE MORE WORK FOR THE ANALYST?"

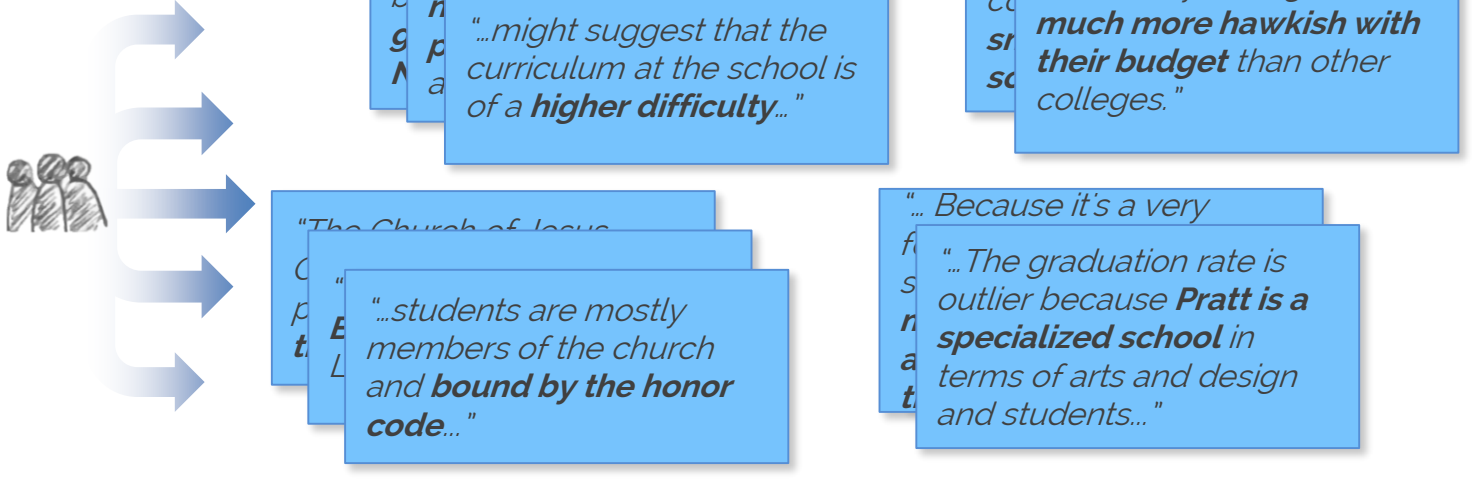

A WORKFLOW FOR CROWDSOURCING DATA ANALYSIS

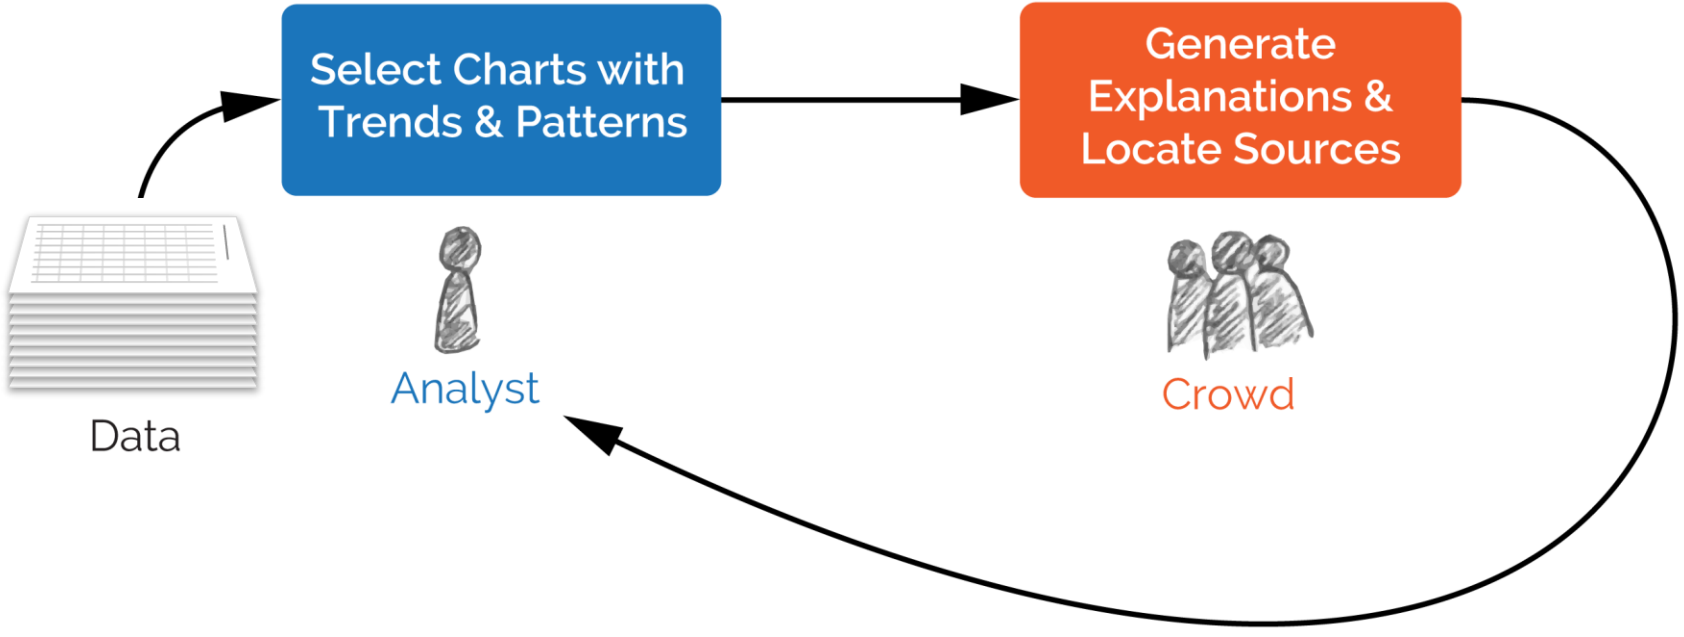

CROWD-ENABLED EXTENSIONS FOR PROCESSING AND MANAGING RESULTS

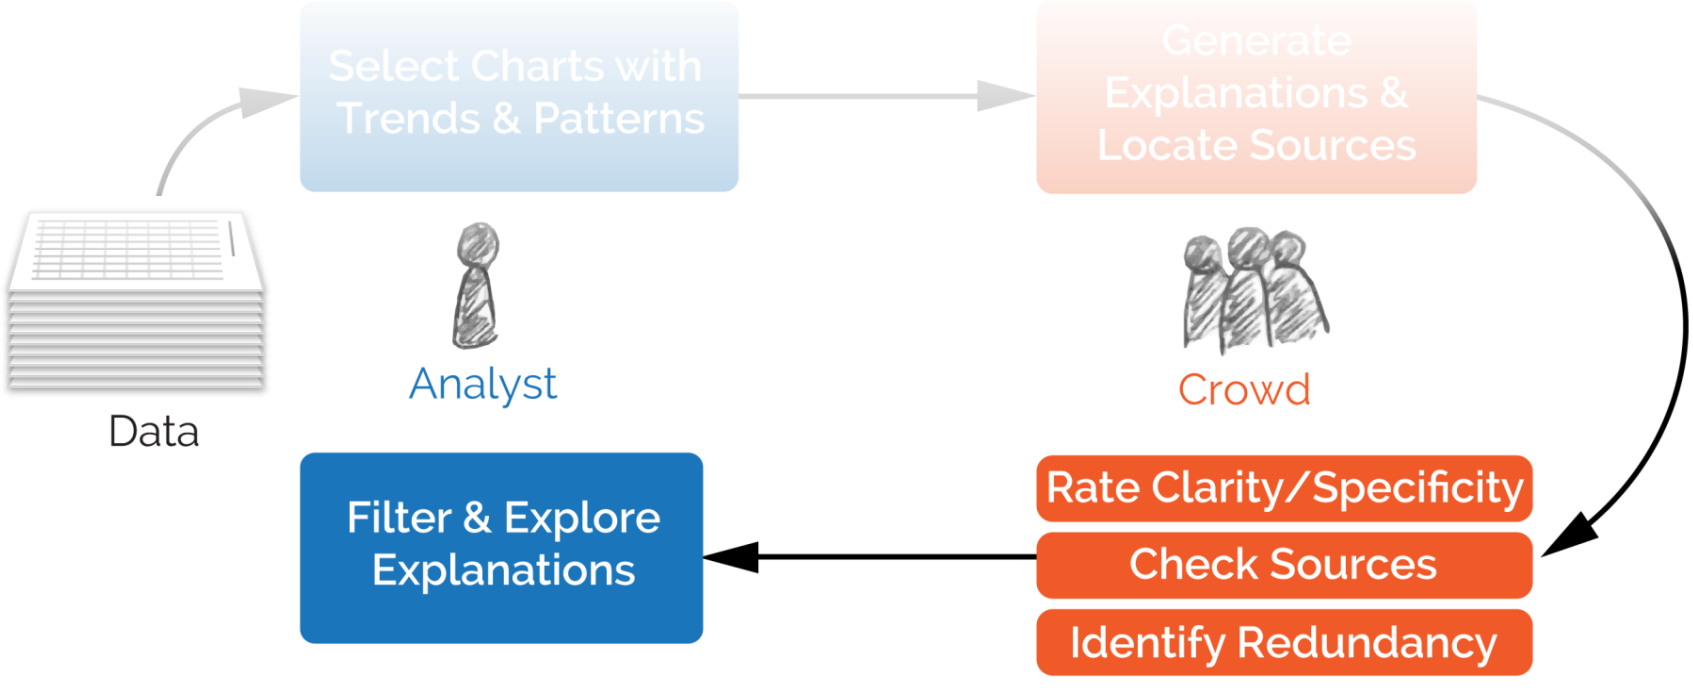

THREE CRITERIA FOR PLAUSIBLE EXPLANATIONS

CLARITY AND SPECIFICITY

PROVENANCE

REDUNDANCY

+ AN INTERFACE FOR MANAGING CROWDSOURCED EXPLANATIONS

CLARITY & SPECIFICITY

CLARITY AND SPECIFICITY

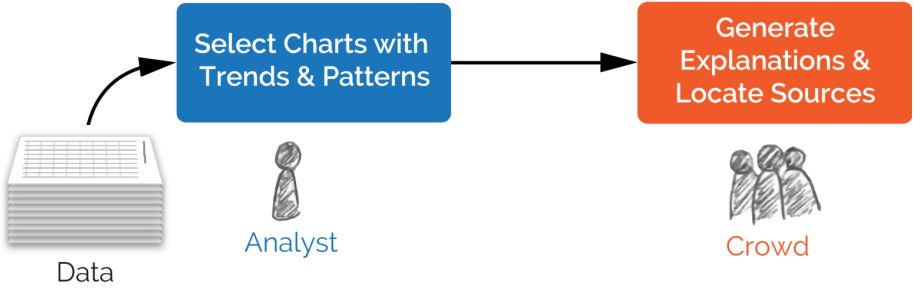

Rating Task

Show Instructions

Each of the charts in this hit compares the graduation rate (x-axis) and the total cost to graduate (y-axis) for 554 top US colleges and universities (as ranked by Bloomberg Businessweek in 2010). Each point represents a single college or university.

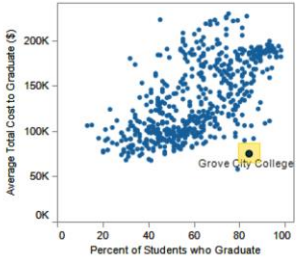

Prompt: Explain **why** the ***outlier*** highlighted in the chart might be different from the other items. (Give **one** specific, well-justified answer.)

Response R2: *" Grove City College is a private Christian college. The College maintains a strict Christian affiliation, in contrast to many institutions whose religions affiliations have become merely historical in nature. This Christian identity, as well as a heavily politically Conservative identity, on campus may likely attract superior students who would not choose to attend otherwise comparable institutions lacking this culture."*(Reference: http://www.discoverthenetworks.org/Articles/Conservative%20Colleges.htm)

1. Does this response provide an explanation for **why** the highlighted outlier in the chart might have occurred? ○ **Yes** ○ **No** ○ **None Present**
2. How **clear** and **specific** is the response? Clear/Specific) ← ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 → (Very Clear/Specific)

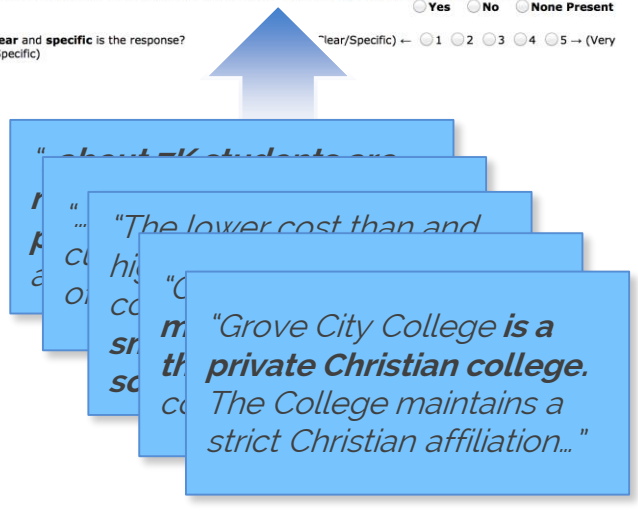

Each of the charts in this hit compares the graduation rate (x-axis) and the total cost to graduate (y-axis) for 554 top US colleges and universities (as ranked by Bloomberg Businessweek in 2010). Each point represents a single college or university.

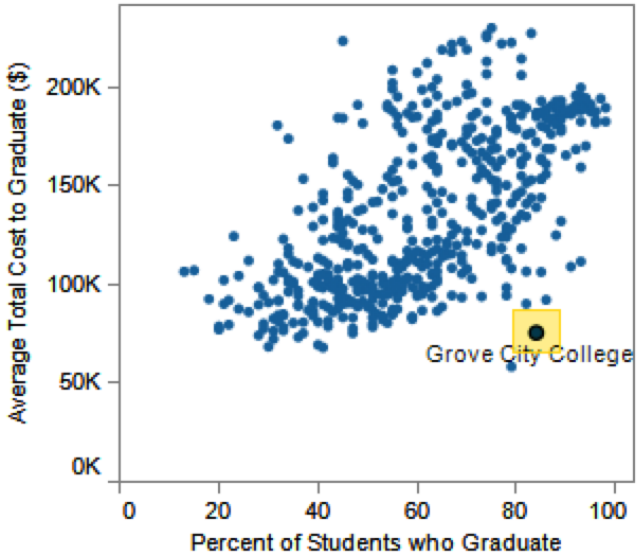

Prompt: Explain **why** the ***outlier*** highlighted in the chart might be different from the other items. (Give **one** specific, well-justified answer.)

Response R2: *" Grove City College is a private Christian college. The College maintains a strict Christian affiliation, in contrast to many institutions whose religions affiliations have become merely historical in nature. This Christian identity, as well as a heavily politically Conservative identity, on campus may likely attract superior students who would not choose to attend otherwise comparable institutions lacking this culture."*(Reference: http://www.discoverthenetworks.org/Articles/Conservative%20Colleges.htm)

1. Does this response provide an explanation for **why** the highlighted outlier in the chart might have occurred?
   ◯ **Yes** ◯ **No** ◯ **None Present**

2. How **clear** and **specific** is the response? (Not Clear/Specific) ← ◯ 1 ◯ 2 ◯ 3 ◯ 4 ◯ 5 → (Very Clear/Specific)
   Clear/Specific)

PROVENANCE

PROVENANCE

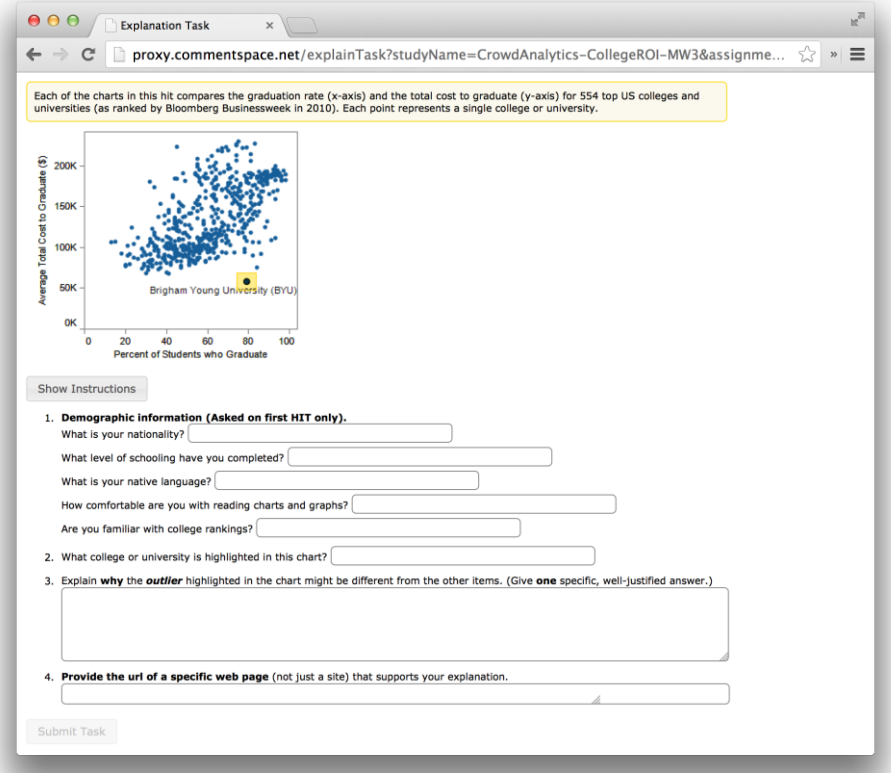

Explanation Task

What are our

workers doing?

PROVENANCE

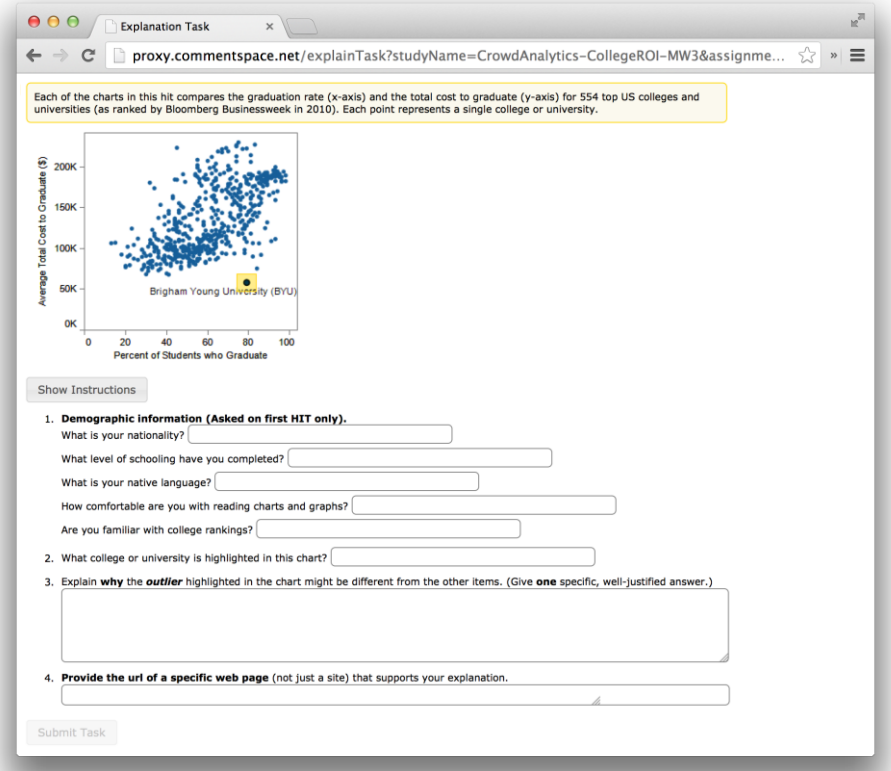

Explanation Task

INSTRUMENTING EXPLANATION TASKS

Examine a line chart showing employment change in a US city and briefly explain it.

Requester: visualizationlab.ucb **Reward:** $0.40 per HIT **HITs Available:** 10 **Duration:** 30 minutes

Qualifications Required: Location is US

Each of the charts in this hit compares the graduation rate (x-axis) and the total cost to graduate (y-axis) for 554 top US colleges and universities (as ranked by Bloomberg Businessweek in 2010). Each point represents a single college or university.

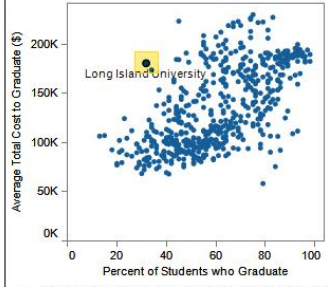

1. What college or university is highlighted in this chart?

2. Explain **why** the strong ***outlier*** highlighted in the chart might be different from the other items. (Try to give **one** specific, well-justified answer per text box.)

   If there are multiple explanations, enter each one in a separate text box.

   Using the browser to the right, find text on a web page that justifies each explanation. Select the text and click the "mark as source" button to add it.

Explanation 1

Source:

+ Add Another Explanation | -

Finished with this HIT? Let someone else do it?

Submit HIT | Return HIT

PROVENANCE

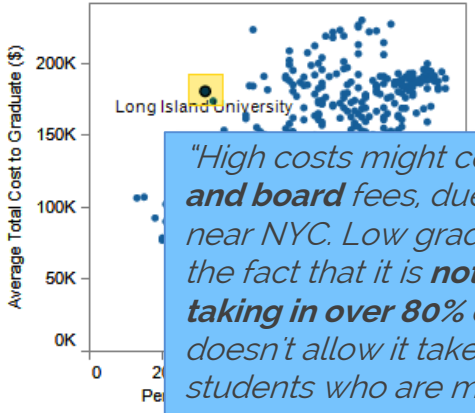

*"High costs might come from it's **high room and board** fees, due to its geographic location near NYC. Low graduation rates come from the fact that it is **not a very selective school, taking in over 80% of applicants**, which doesn't allow it take many top ranked students who are more academically motivated."*

Visitation logs

2011-12-11 09:22:04 google.com
2011-12-11 09:22:04 sqr:helo
2011-12-11 09:23:08 google.com/search?hl=en&source=hp
2011-12-11 09:23:11 google.com/search?hl=en&q=Long Isl
2011-12-11 09:23:13 google.com/search?q=Long Island Un
2011-12-11 09:23:31 google.com/search?q=Long Island Un
2011-12-11 09:23:38 google.com/search?q=Long Island Un
2011-12-11 09:23:43 google.com/search?q=Long Island Un
2011-12-11 09:23:54 google.com/search?q=Long Island Un
2011-12-11 09:24:09 colleges.usnews.rankingsandreviews.c

Paragraph-level citations

#123 Regional Universities (North)

Summary

LIU Post is a private institution that was founded in 1954. It has a total undergraduate enrollment of 8,315, its setting is suburban, and the campus size is 308 acres. It utilizes a semester-based academic calendar. LIU Post's ranking in the 2014 edition of Best Colleges is Regional Universities (North), 123. Its tuition and fees are $34,070 (2013-14).

2014 Quick Stats

720 Northern Boulevard
Brookville, NY 11548-1300
[map]
Phone: (516) 299-2000

2013-2014 Tuition
$34,070 tuition and fees

Students
8,315 enrolled
25% male / 75% female

Admissions
rolling admission
78.8% accepted

More Information

PROVENANCE

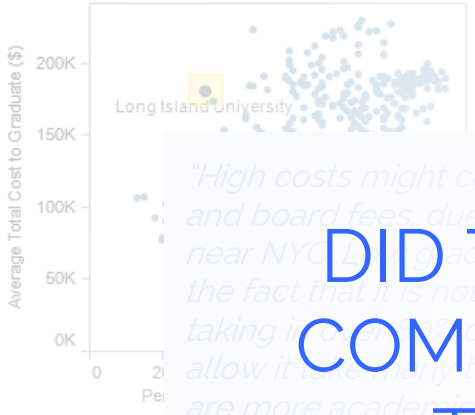

"High costs might come from it's high room and board fees due to its geographic location near NYC. Academics could be affected by the fact that it is not a very selective school taking in many students. Social life could allow it to be pretty lively. Finally, students are more academ[illegible]

DID THE **FACTS** AND **INFERENCE** COME FROM THE **SOURCE** OR DID THE **WORKER** ADD THEM?

Paragraph-level citations

LIU Post is a private institution that was founded in 1954. It has a total undergraduate enrollment of 8,315, its setting is suburban, and the campus size is 308 acres. It utilizes a semester-based academic calendar. LIU Post's ranking in the 2014 edition of Best Colleges is Regional Universities (North), 123. Its tuition and fees are $34,070 (2013-14).

2014 Quick Stats

720 Northern Boulevard
Brookville, NY 11548-1300
[map]
Phone: (516) 299-2000

2013-2014 Tuition
$34,070 tuition and fees

Students
8,315 enrolled
25% male / 75% female

Admissions
rolling admission
78.8% accepted

More Information

Visitation logs

```
2011-12-11 09:22:04 google.com
2011-12-11 09:22:04 sqr:helo
2011-12-11 09:23:08 google.com/search?hl=en&source=h
2011-12-11 09:23:11 google.com/search?hl=en&q=Long I
2011-12-11 09:23:13 google.com/search?q=Long Island U
2011-12-11 09:23:31 google.com/search?q=Long Island U
2011-12-11 09:23:38 google.com/search?q=Long Island U
2011-12-11 09:23:43 google.com/search?q=Long Island U
2011-12-11 09:23:54 google.com/search?q=Long Island U
2011-12-11 09:24:09 colleges.usnews.rankingsandreviews.
```

SOURCE-CHECKING MICROTASKS

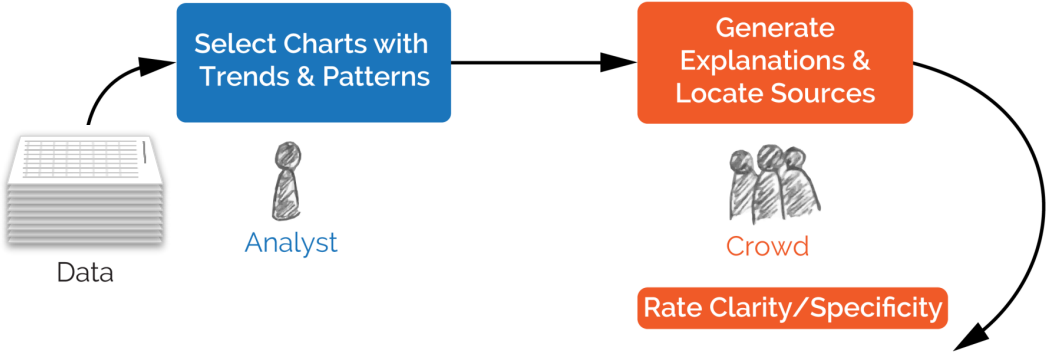

A second group of workers verifies links and attributes explanations to the source or the worker.(75% accurate in our preliminary tests)

REDUNDANCY

REDUNDANCY

Many explanations provided by workers are redundant.

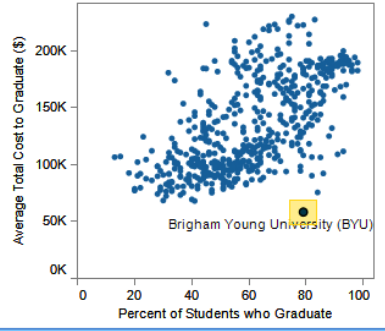

"The Church of Jesus Christ of Latter Day Saints pays a significant part of the tuition costs…"

"The cost of attendance at BYU is subsidized by the LDS church."

"98% of their students are members of LDS and they have lowered tuition…"

REDUNDANCY

Many explanations provided by workers are redundant.

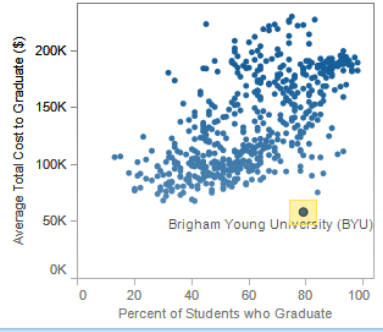

"The Church of Jesus Christ of Latter Day Saints pays a significant part of

"The cost of attendance at BYU is subsidized by the LDS church."

"98% of their students are members of LDS and they have lowered tuition."

- Duplicate results for analysts to examine.

+ Redundancy can signal high support and corroborating sources.

REDUNDANCY

Automated text similarity methods don't deal well with these kinds of content.

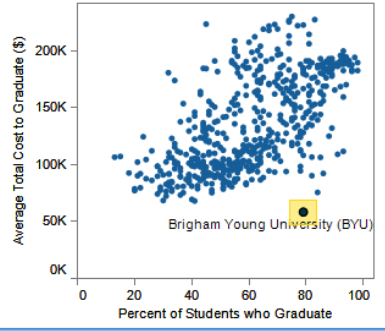

"The Church of Jesus Christ of Latter Day Saints pays a significant part of the tuition costs…"

"The cost of attendance at BYU is subsidized by the LDS church."

"98% of their students are members of LDS and they have lowered tuition…"

REDUNDANCY

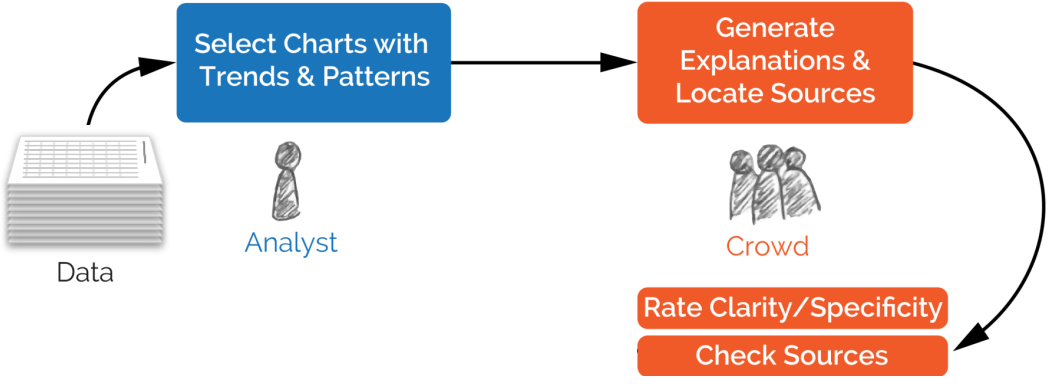

Can we crowdsource redundancy detection?

CLUSTERING VIA DISTRIBUTED COMPARISON

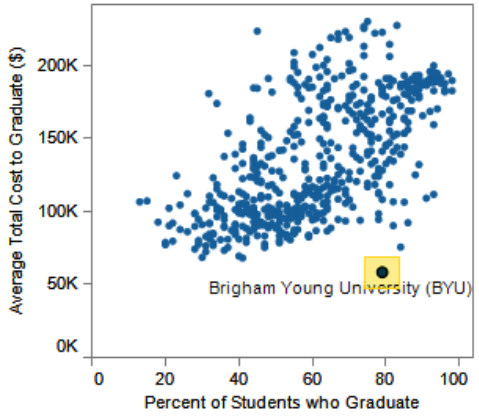

"98% of their students are members of LDS and they have lowered tuition…"

"The cost of attendance at BYU is subsidized by the LDS church."

"…students are mostly members of the church and bound by the honor code…"

"The Church of Jesus Christ of Latter Day

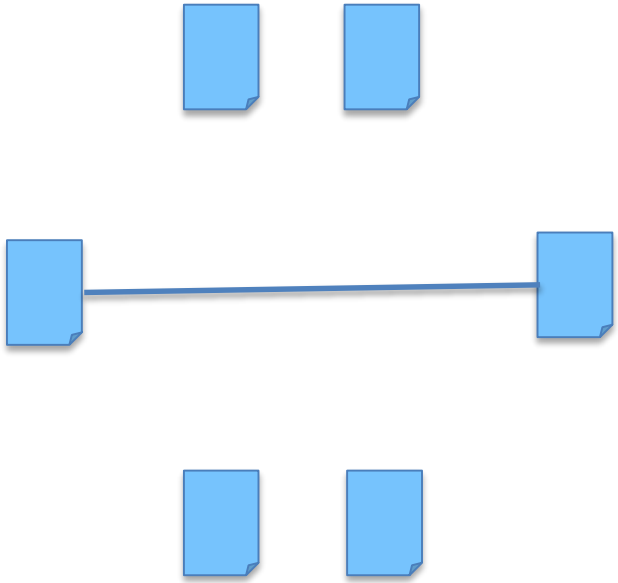

CLUSTERING VIA DISTRIBUTED COMPARISON

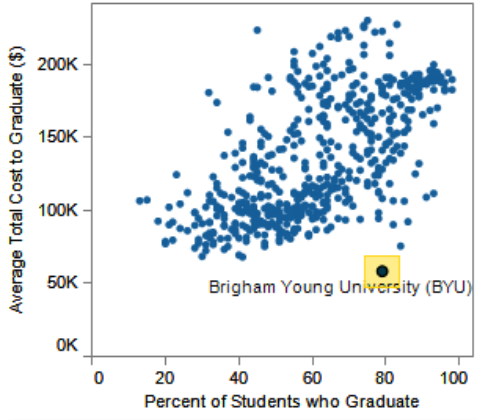

"98% of their students are members of LDS and they have lowered tuition…"

"The cost of attendance at BYU is subsidized by the LDS church."

Do these two responses give the same general explanation for the peaks and valleys in the chart?

- ○ **Yes. Both responses give the same general explanation.**
- ○ **No. The responses do not give the same explanation.**

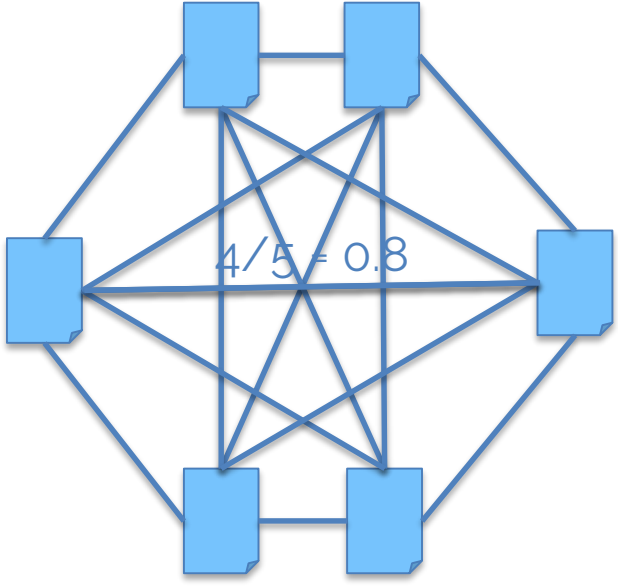

CLUSTERING VIA DISTRIBUTED COMPARISON

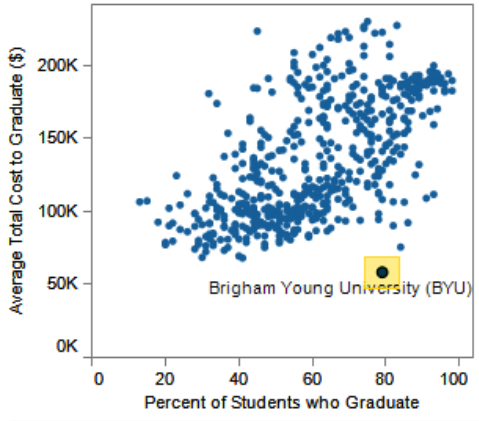

"98% of their students are members of LDS and they have lowered tuition..."

"The cost of attendance at BYU is subsidized by the LDS church."

Do these two responses give the same general explanation for the peaks and valleys in the chart?

- **Yes. Both responses give the same general explanation.**
- **No. The responses do not give the same explanation.**

\+ Simple tasks for workers

\+ Scales p[illegible]rly

\- Sensitive to clustering method

\- Workers have little context

CLUSTERING VIA COLOR-CODING

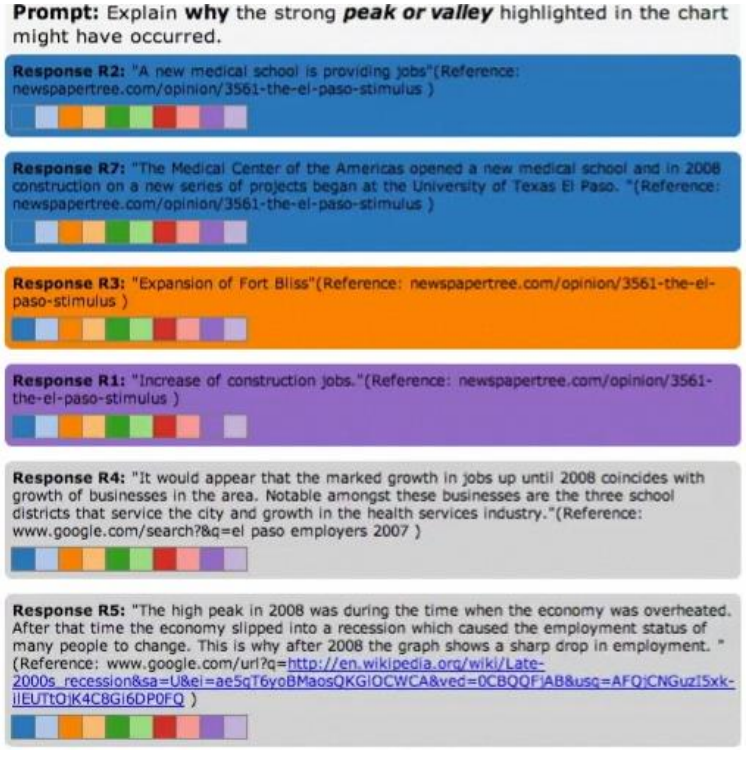

MULTIPLE WORKERS **INDEPENDENTLY CLUSTER** THE WHOLE SET.

USE **COMPUTATIONAL SIMILARITY METRICS** TO SELECT THE BEST, CONSISTENT CLUSTERING.

FINDING THE RIGHT BALANCE OF HUMAN AND AUTOMATED EFFORT

MANAGING THE CROWD'S WORK

MANAGING THE CROWD'S WORK

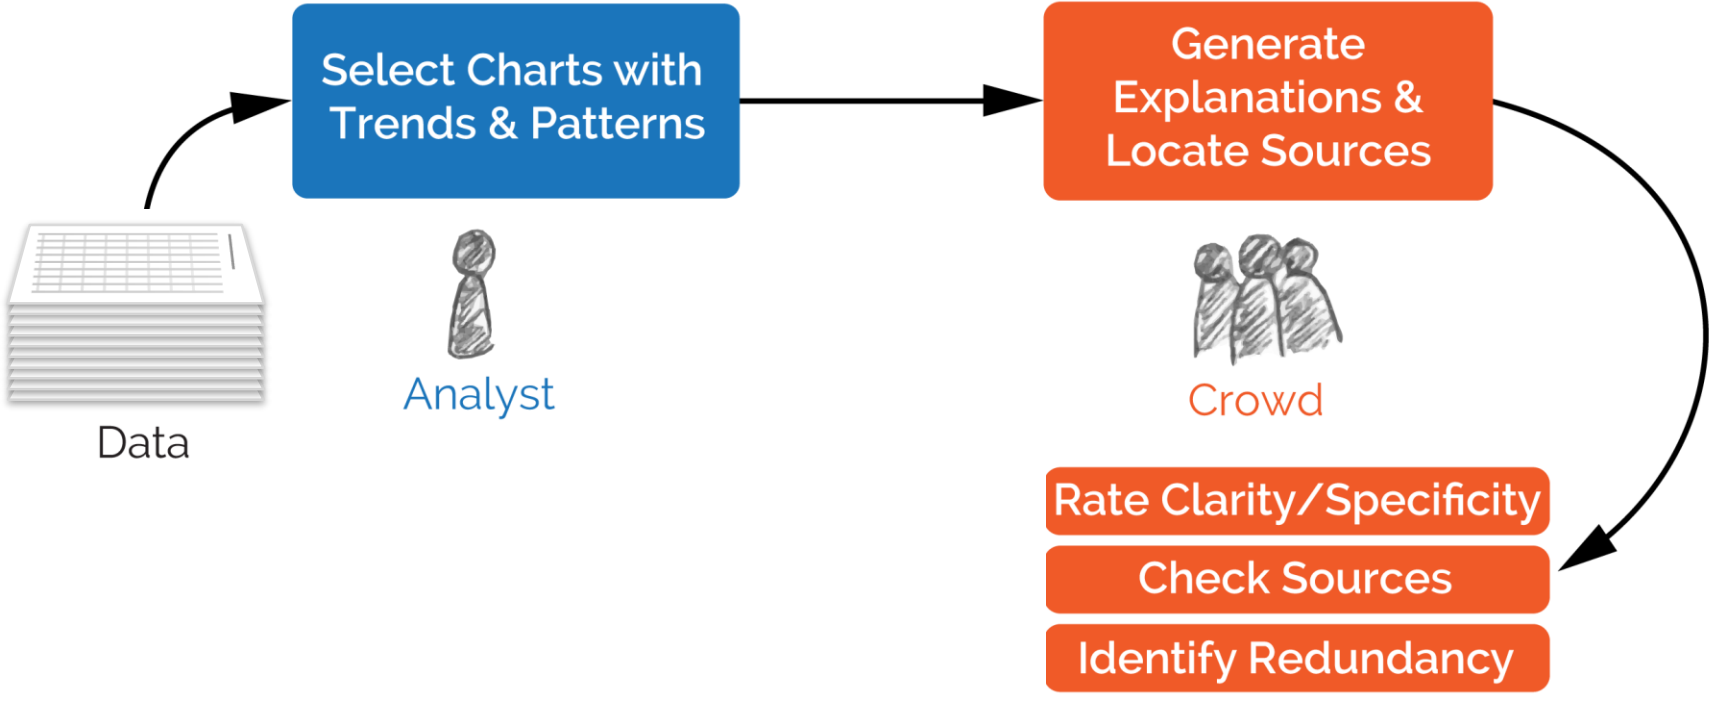

EXPLANATION MANAGEMENT INTERFACE

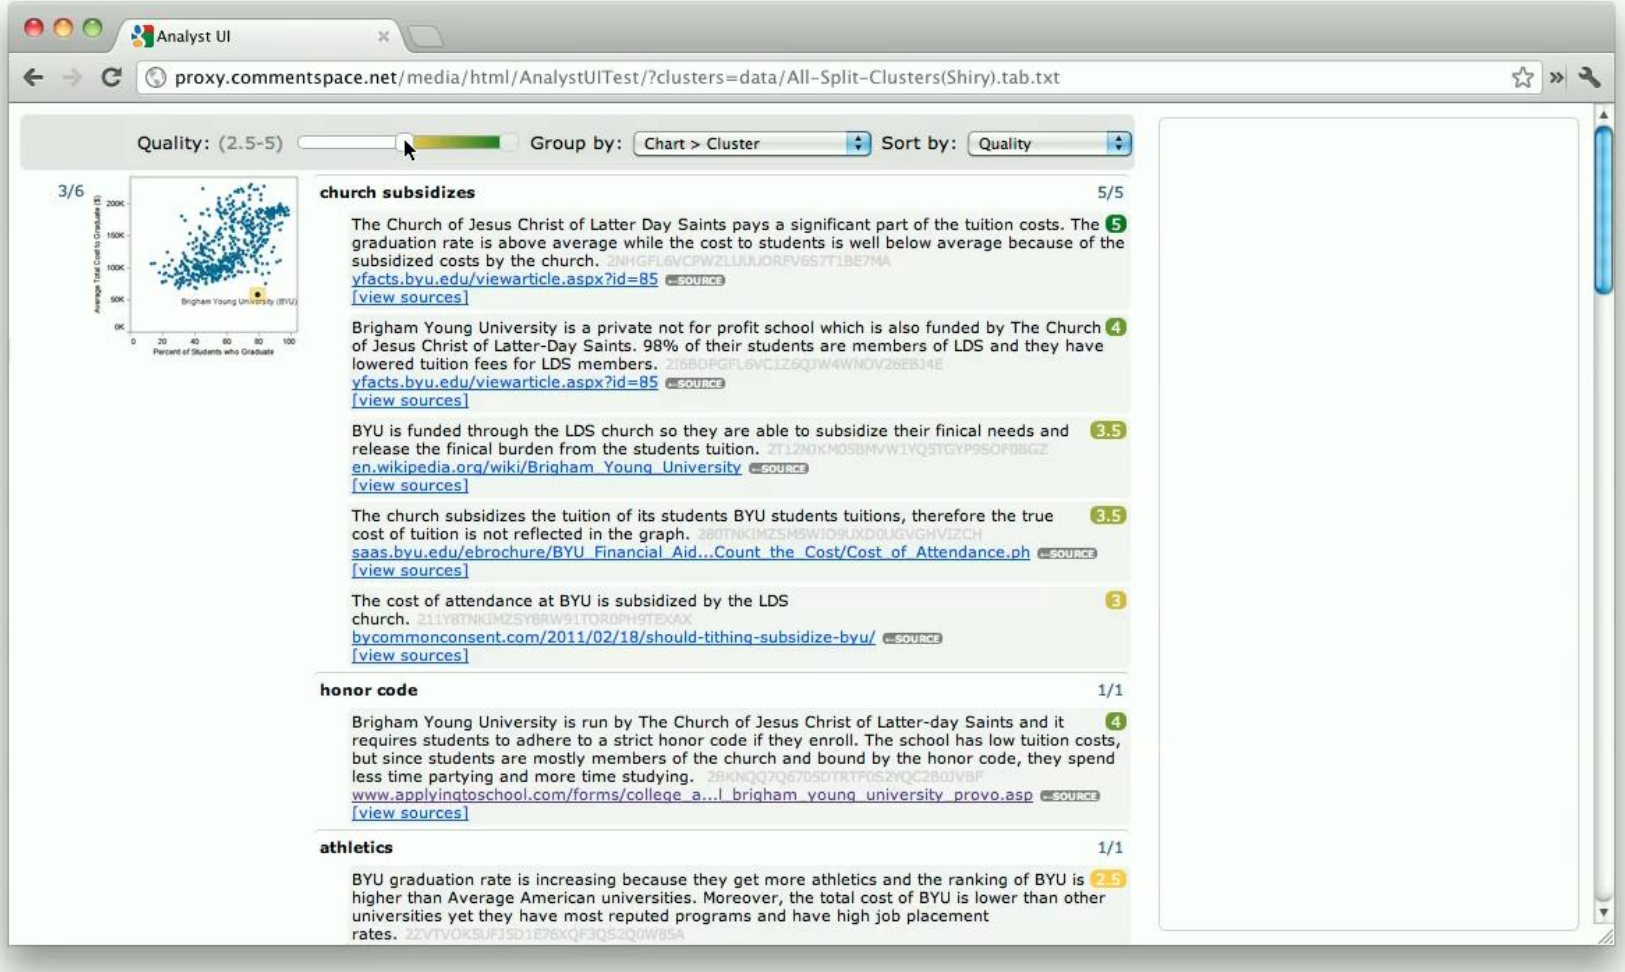

CROWDSOURCING HIGH-LEVEL ANALYSIS

HUMAN COMPUTATION CAN BE A USEFUL COMPLEMENT TO AUTOMATED PROCESSING

EVEN MORE INTERESTING WITH EXPERTISE

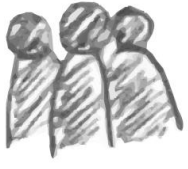

cheap low-skill crowds

VS.

more knowledgeable trusted ones

UNDERSTANDING HOW TO PARALLELIZE **ANALYSIS PROCESSES** MAY BE AS IMPORTANT AS PARALLELIZING COMPUTATION HAS BEEN.

DATA ANALYSIS AT SCALE

CHALLENGES

ANALYSIS AND CLUSTER COMPUTING

INTERACTING WITH BIG DATA

PARALLELIZING HUMAN INTELLIGENCE

UP NEXT

AFTER THE BREAK
APPLICATION AREAS

THIS AFTERNOON
DINO FUN WORLD PRESENTATIONS
(OPEN LAB)

DECEMBER 8th-19th
INFORMATION VISUALIZATION LECTURES
AT UNIVERSITÉ PARIS SUD

BONUS MATERIAL

MORE DETAILS ON CROWDSOURCED DATA ANALYSIS

CLUSTERING VIA COLOR-CODING

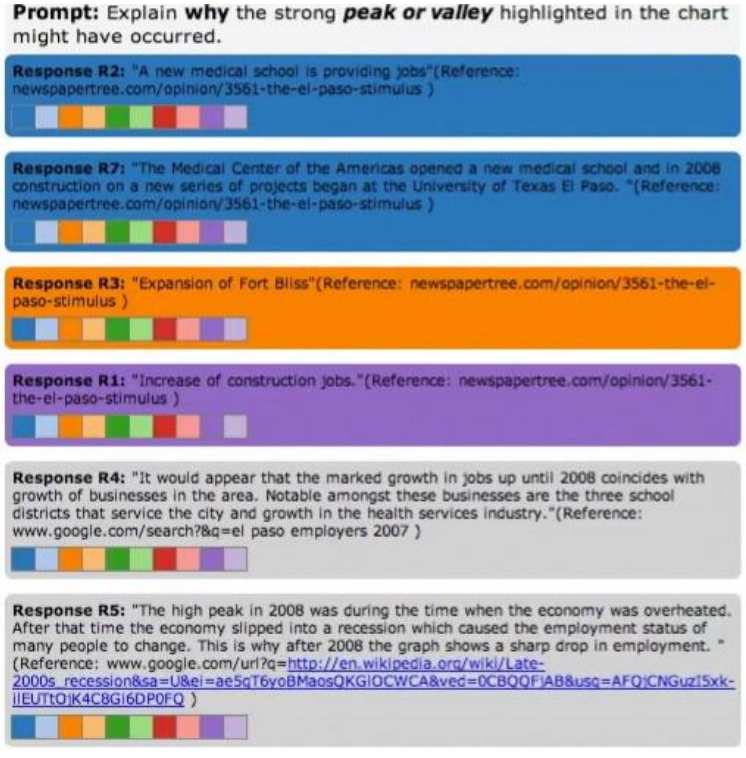

Individual workers cluster the whole set.

\+ Workers have **complete context**

\- **Individual workers** can cluster badly

Hard to integrate clusterings from multiple workers

HOW TO INTEGRATE COLOR-CLUSTERINGS?

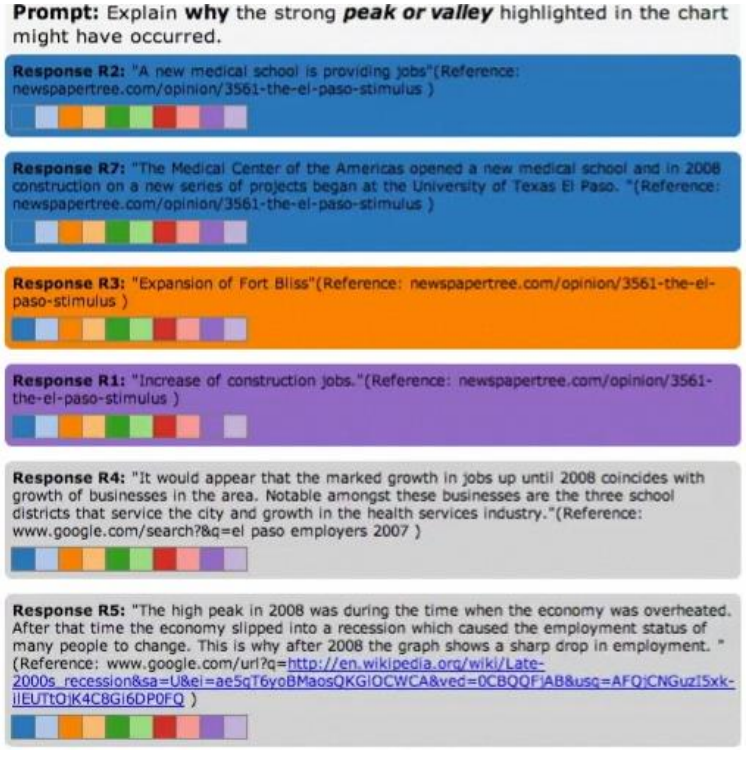

- A single worker's clustering is preferable to a combination of multiple clusterings.
- Clusters reproduced by multiple independent workers are likely to reflect actual redundancy.
- Errors tend to be either noisy or easy to catch.

HOW TO INTEGRATE COLOR-CLUSTERINGS?

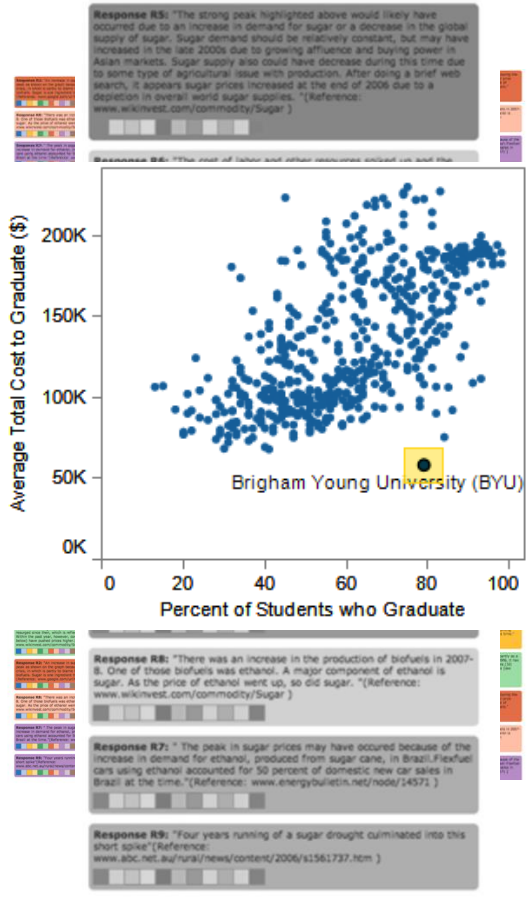

Selecting the Most-Representative Clustering

HOW TO INTEGRATE COLOR-CLUSTERINGS?

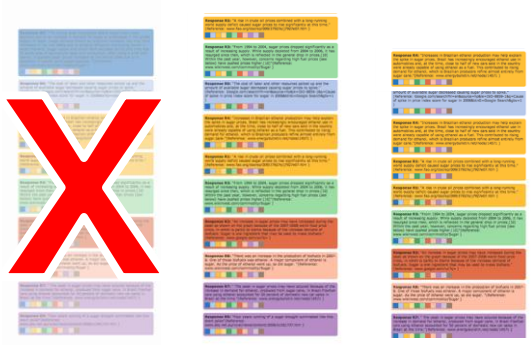

SELECTING THE MOST-REPRESENTATIVE CLUSTERING

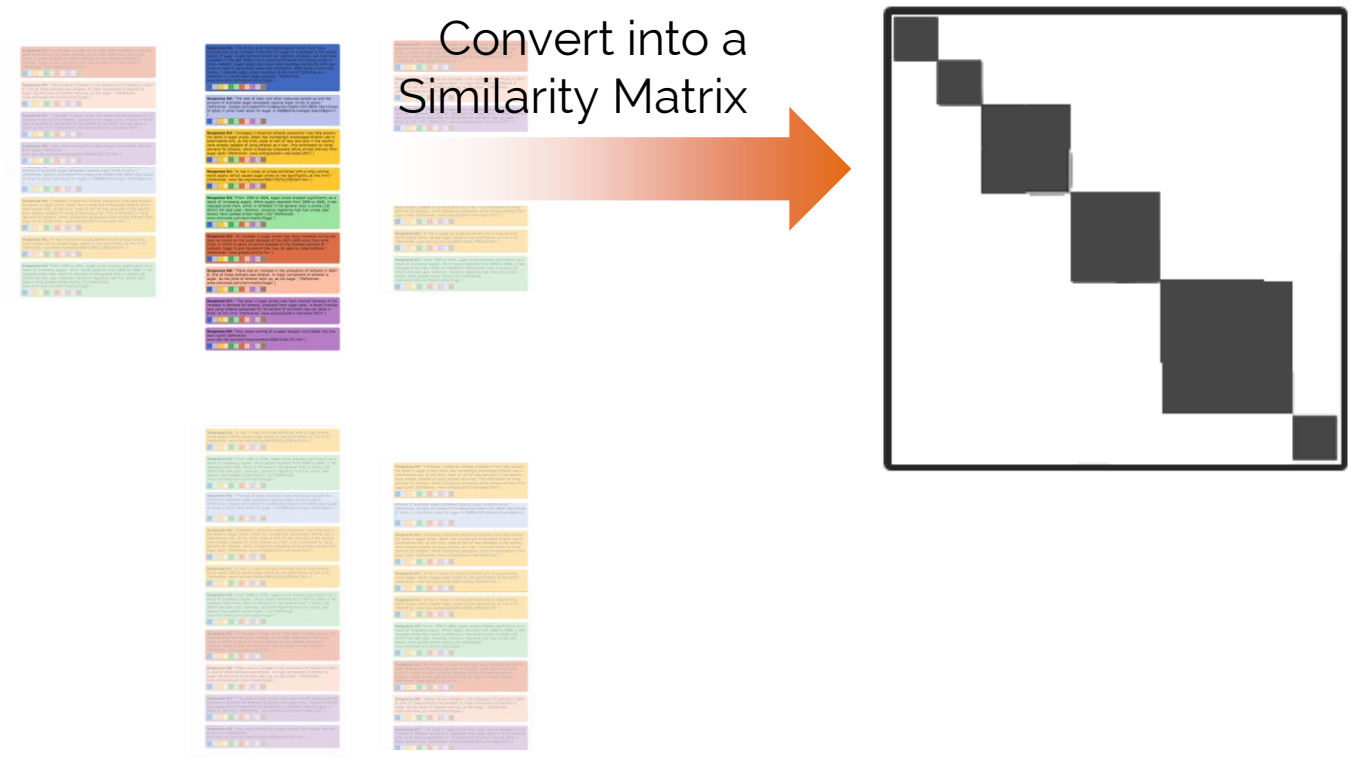

SELECTING THE MOST-REPRESENTATIVE CLUSTERING

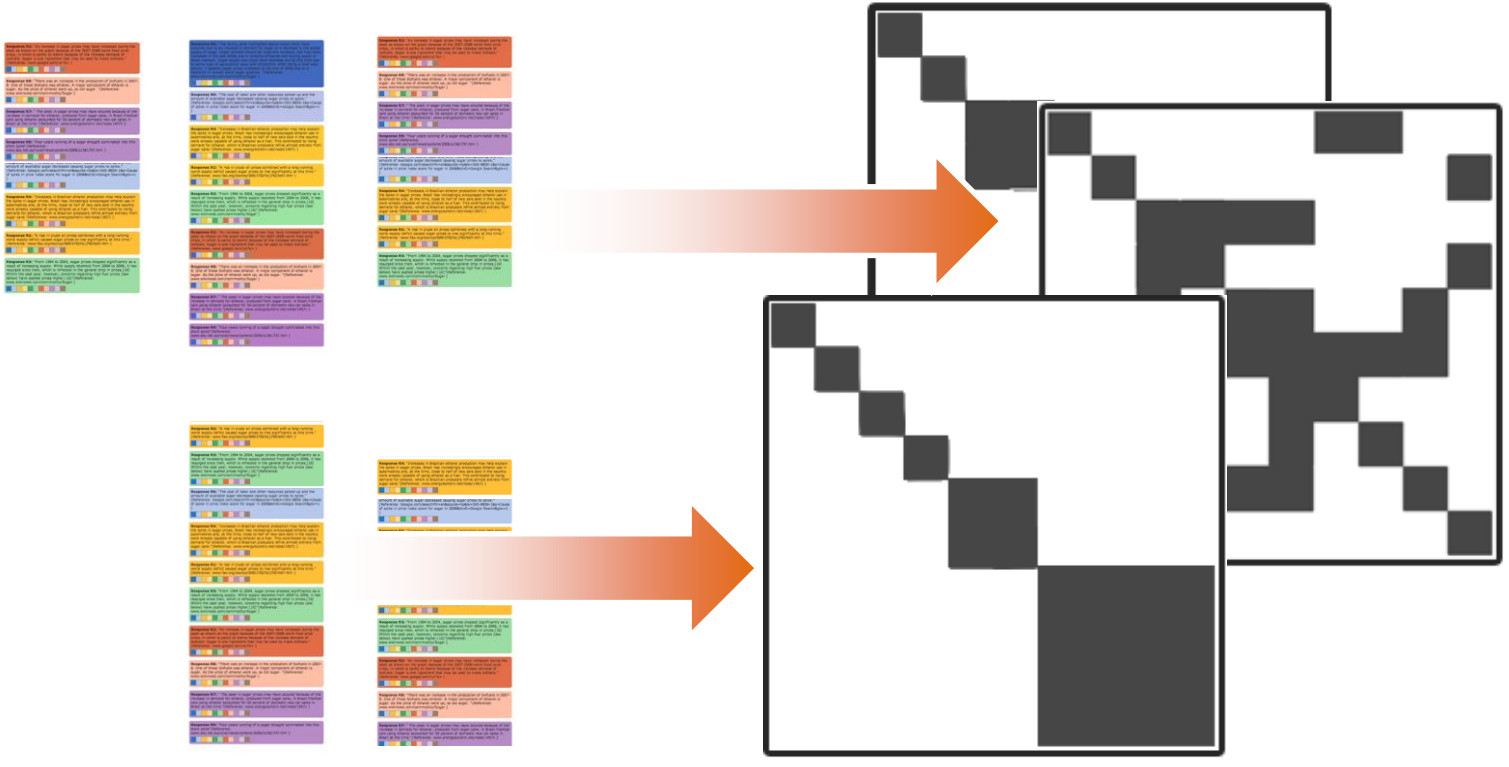

SELECTING THE MOST-REPRESENTATIVE CLUSTERING

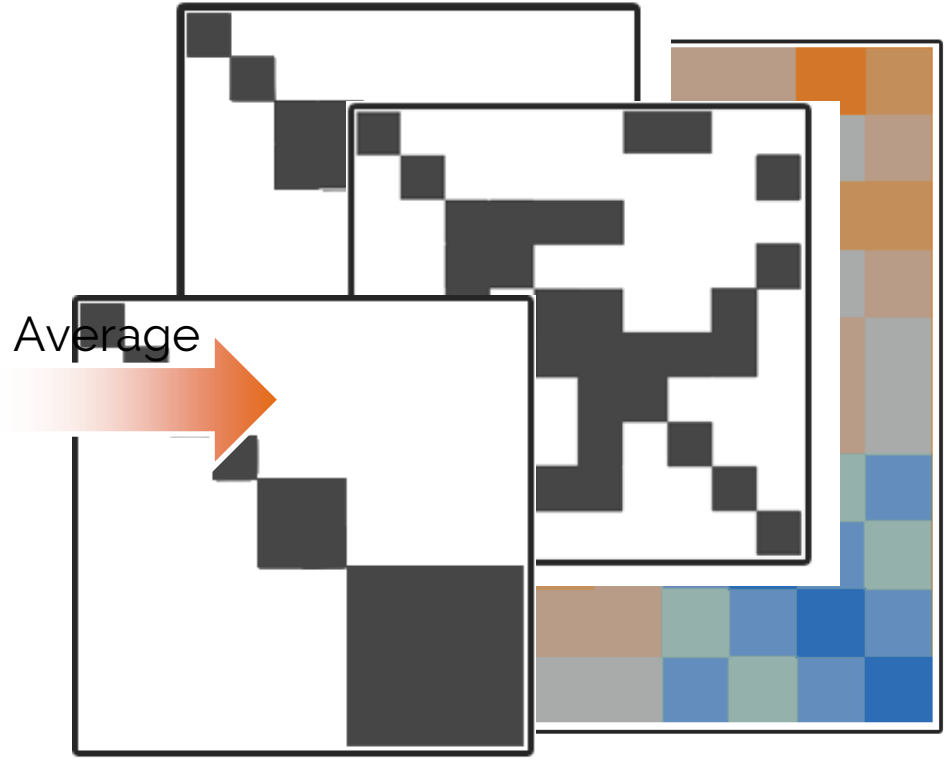

SELECTING THE MOST-REPRESENTATIVE CLUSTERING

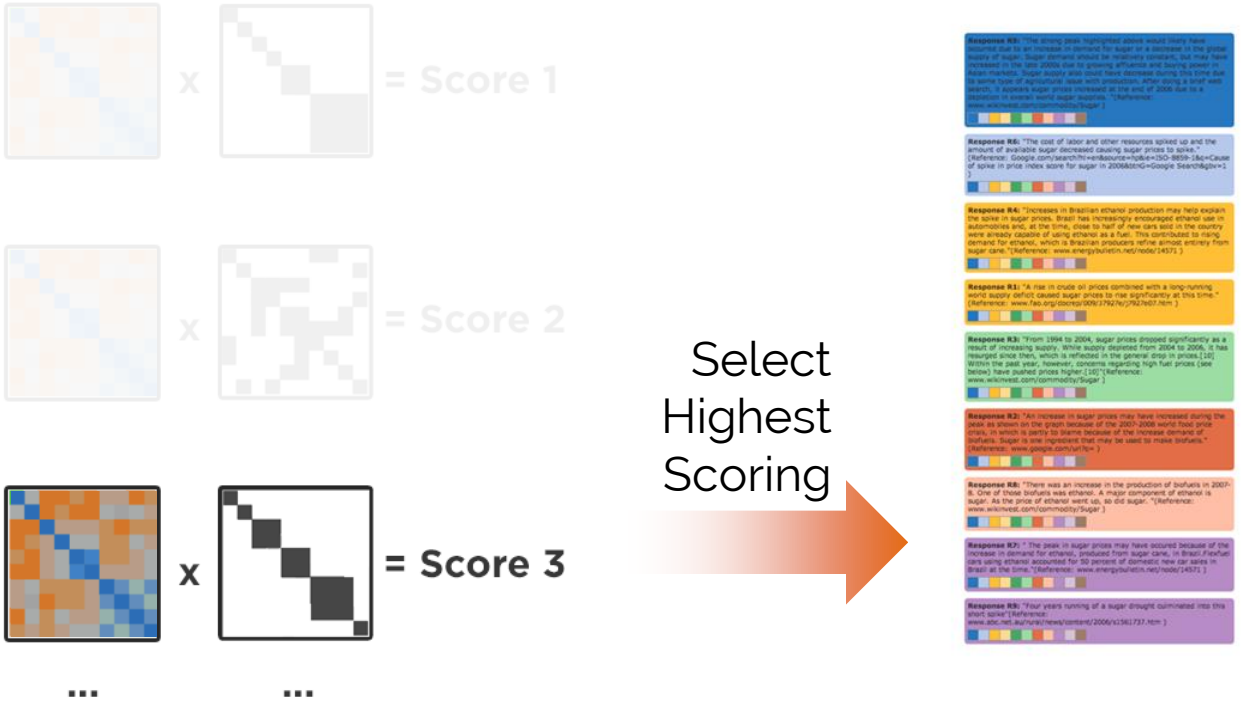

EVALUATING REDUNDANCY-DETECTION

Does color clustering with most-representative selection produce good clusterings?

Our Explanation Dataset

12 charts (4 each from 3 different data sets)

10 workers explained each chart

➡ 93 Workers produced 156 explanations (Avg=13 per chart)

EVALUATING REDUNDANCY-DETECTION

Does color clustering with most-representative selection produce good clusterings?

10 Workers used color clustering to group the explanations for each chart. (120 total clusterings)

We used most-representative selection to pick the best clustering for each chart. (12 clusterings)

EVALUATING REDUNDANCY-DETECTION

Baseline - Expert clustering (x 3)

To score a clustering, we use the F-measure to compute similarity to each expert, then average.

(completely dissimilar) [0 ⟷ 1] (identical)

EVALUATING REDUNDANCY-DETECTION

Unclustered Results F=0.68

Color Clustering F=0.73

Most-Representative Selection F=0.86

Experts vs. One Another F=0.84

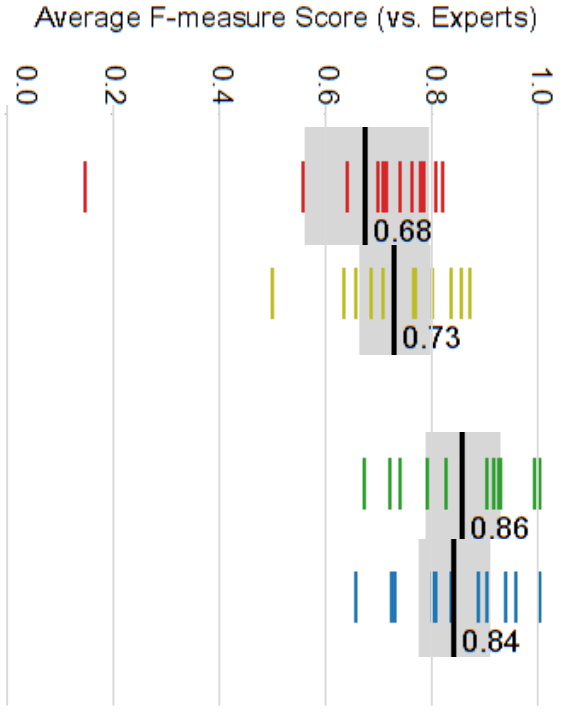

EVALUATING REDUNDANCY-DETECTION

Unclustered Results F=0.68

Color Clustering F=0.73

Most-Representative Selection F=0.86

Experts vs. One Another F=0.84

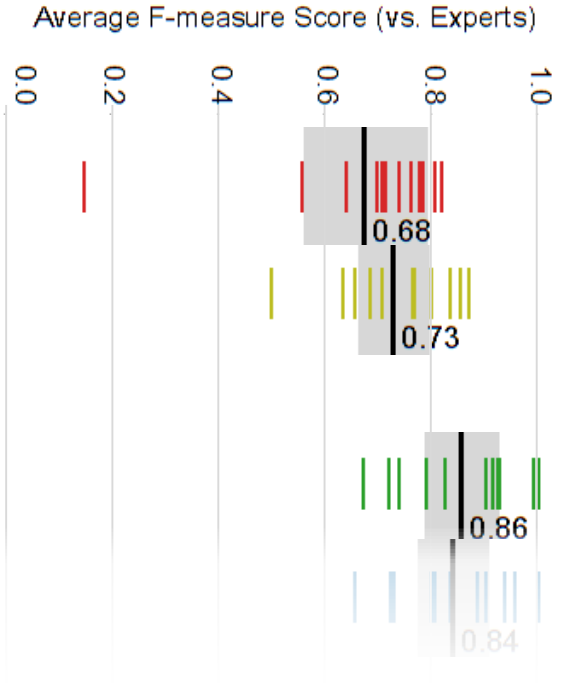

T-tests showed our most-representative results were significantly closer to experts than color clustering or unclustered were. (both $p < 0.01$)