Interpretting Embeddings with Comparison

Michael Gleicher

Department of Computer Sciences

University of Wisconsin - Madison

Two talks (or five) in one!

How do we deal with embeddings of text data?

Alexander, et. al. **Serendip: Topic Model-Driven Visual Exploration of Text Corpora**. VAST '14.

Alexander & Gleicher. **Task-Driven Comparison of Topic Models**. VAST '15

Heimerl & Gleicher. **Interactive Analysis of Word Vector Embeddings**. EuroVis '18.

Heimerl, et al. **Interactive Visual Comparison of Object Embeddings**. (submitted)

How do we use comparison as a tool for hard data problems?

Gleicher. **Considerations for Visualizing Comparison**. InfoVis 2017.

How do we deal with embeddings of text data?

Alexander, et. al. **Serendip: Topic Model-Driven Visual Exploration of Text Corpora**. VAST '14.

Alexander & Gleicher. **Task-Driven Comparison of Topic Models**. VAST '15

Heimerl & Gleicher. **Interactive Analysis of Word Vector Embeddings**. EuroVis '18.

Heimerl, et al. **Interactive Visual Comparison of Object Embeddings**. [submitted]

Eric Alexander
Assistant Prof.
Carleton College

Florian Heimerl
Post-Doc
University of Wisconsin

`

Michael Gleicher

University of Wisconsin
Visual Computing Group

Human Graphics Interaction

authoring pictures, videos, animations

Human Robot Interaction

robots!

Human Data Interaction

visualization, visual analytics, interactive learning

~~Data Visualization~~

~~Visual Analytics~~

Human Data Interaction

Interpretting Embeddings with Comparison

Michael Gleicher

Department of Computer Sciences

University of Wisconsin - Madison

Two different stories…

Embeddings are a great way to analyze text collections!

But they are hard to interpret!

Use comparison as a strategy

Comparison is a great way to think about analysis problems!

But it's abstract – need examples!

Use embeddings as a case study

Interpretting **Embeddings** with Comparison

Michael Gleicher
Department of Computer Sciences
University of Wisconsin - Madison

What is an embedding?

General mathematics:

Place a smaller structure into a larger structure

Computer science:

Place a discrete set of objects into a vector space

Encode relationships between objects

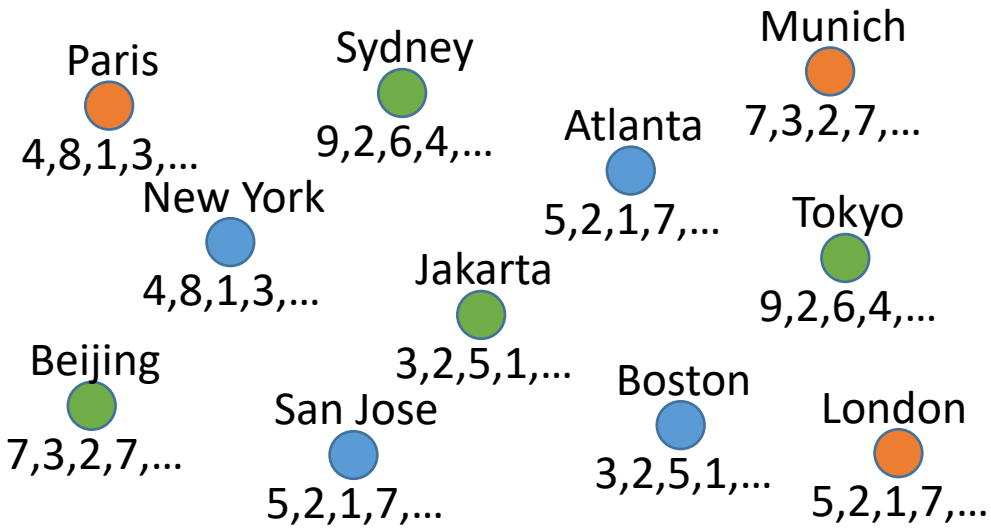

High Dimensional Data

Objects have associated **Vectors**

Kinds of relationships in embeddings

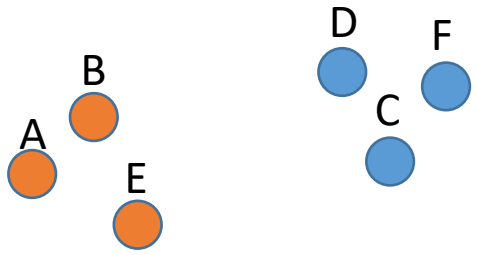

Distance

A is closer to B than to C

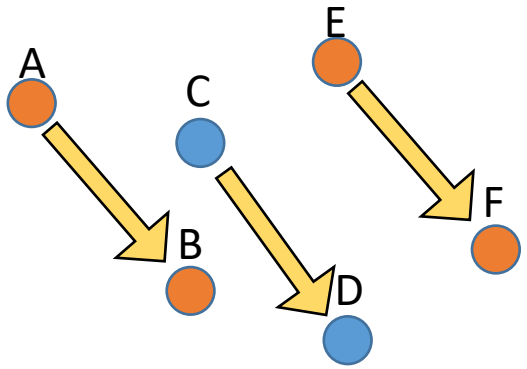

Linear Structure

A is to B as C is to D

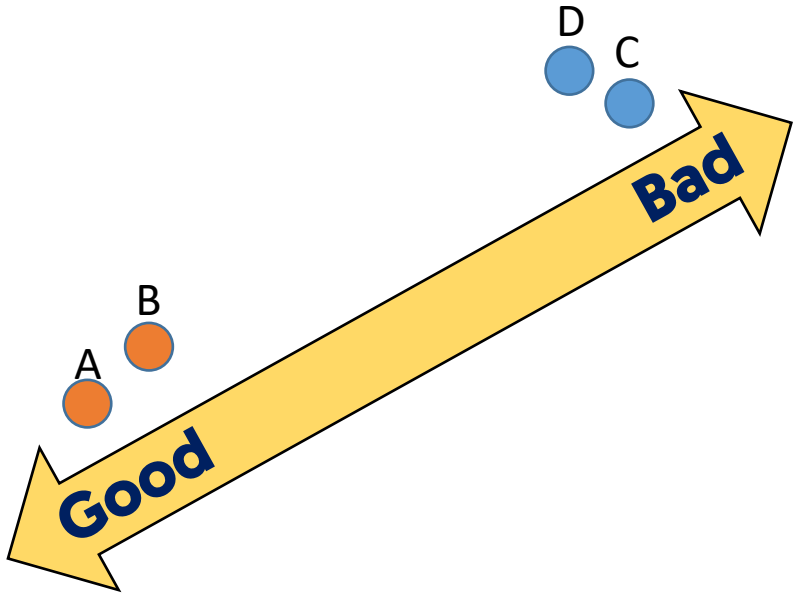

Semantic Directions

A is more X than C

Relationships are interesting even if global positions are not

Embeddings

We care about **relationships** not **values**

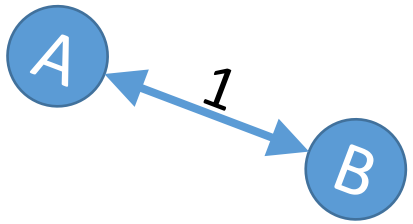

The coordinates (axes) **may** have no meaning

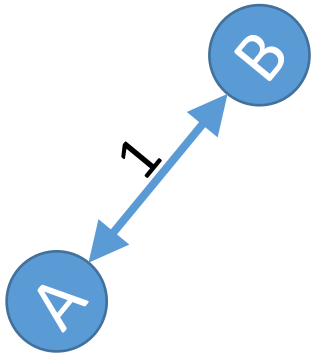

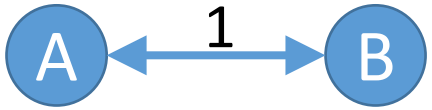

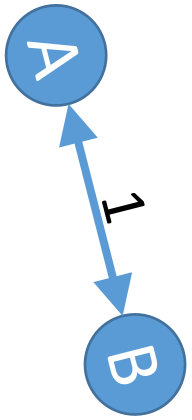

Very different shapes?

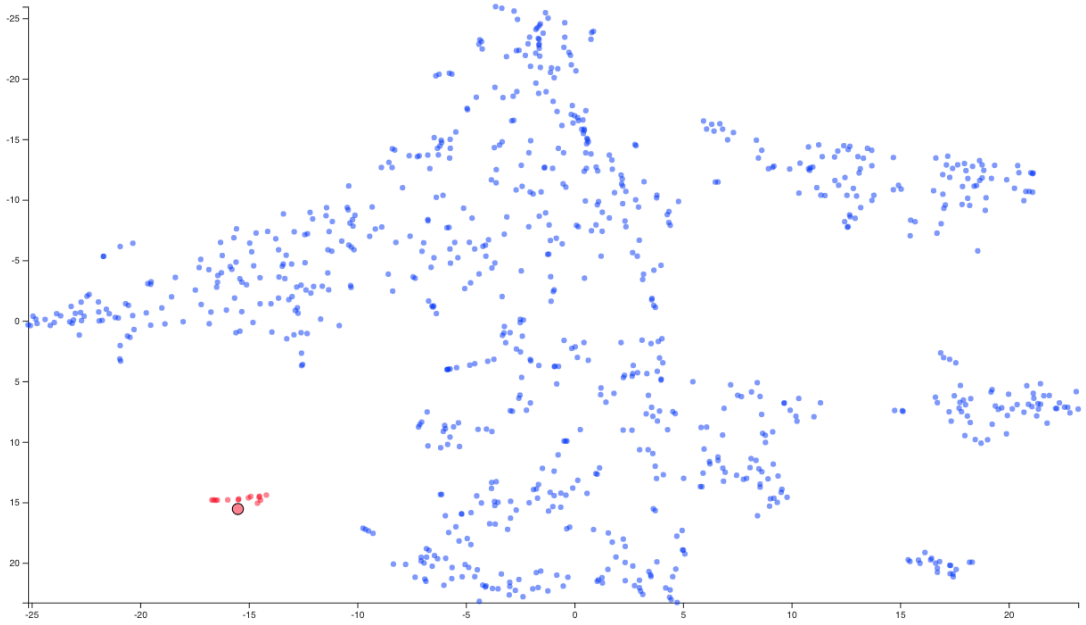

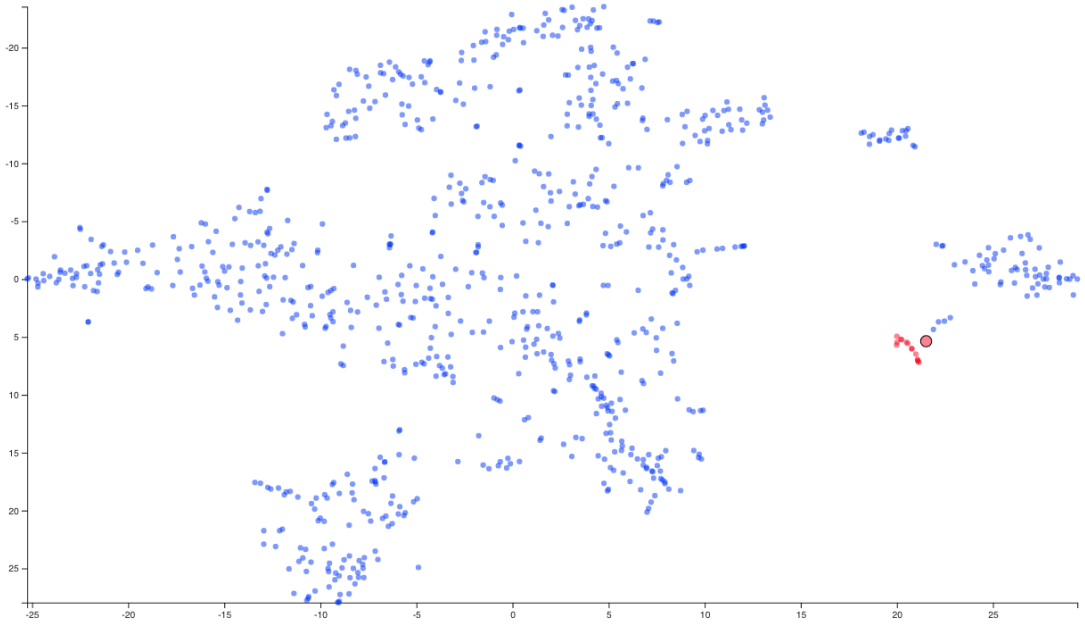

Similar neighbors

(points close in one are close in the other)

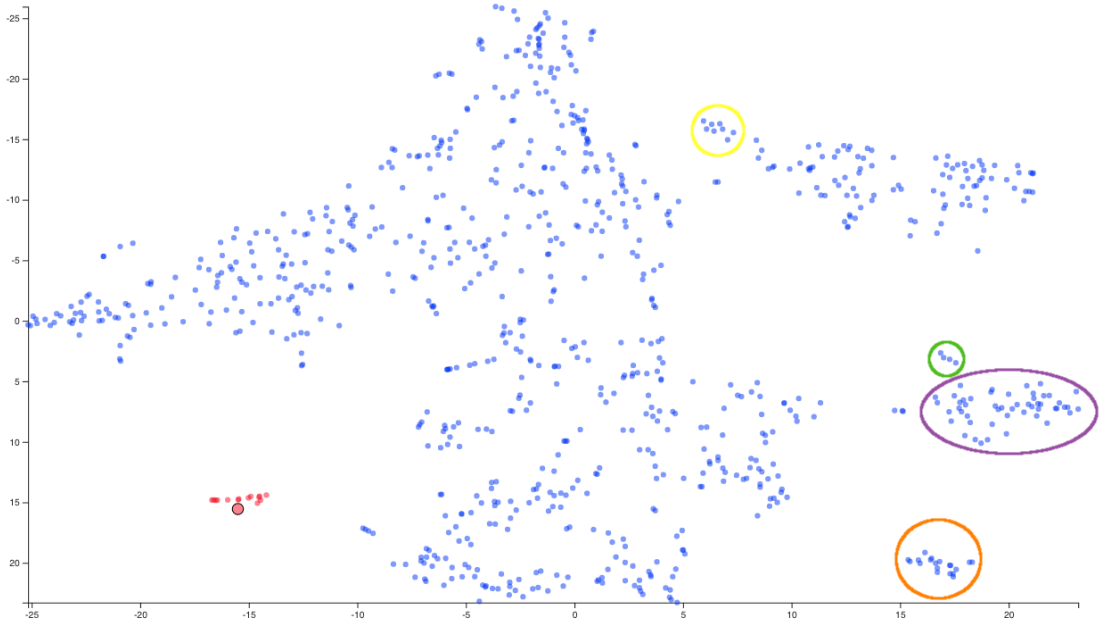

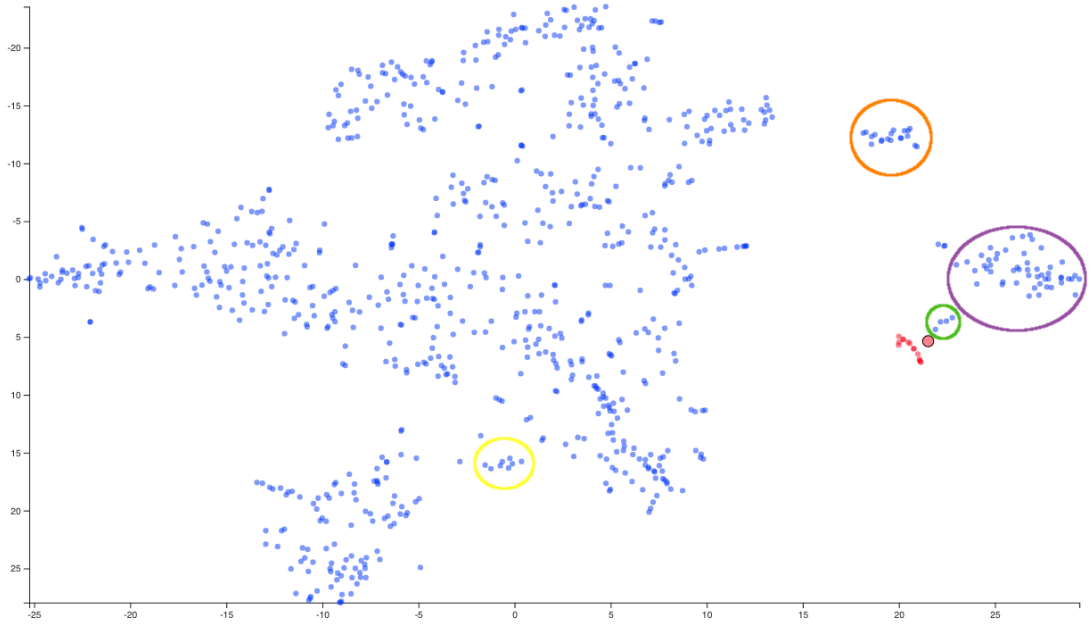

Embeddings for analyzing Text Corpora

Embed Documents

Similar documents are close

Estimate similar meaning by word usage statistics in a corpus

Embed Words

Similar words are close

Estimate similar meaning by word usage statistics in a corpus

Document Similarity

Documents that use similar words (probably) have similar content

Idea from the 1960s

OBITUARIES The New York Times

Overlooked No More: Karen Sparck Jones, Who Established the Basis for Search Engines

A pioneer of computer science for work combining statistics and linguistics, and an advocate for women in the field.

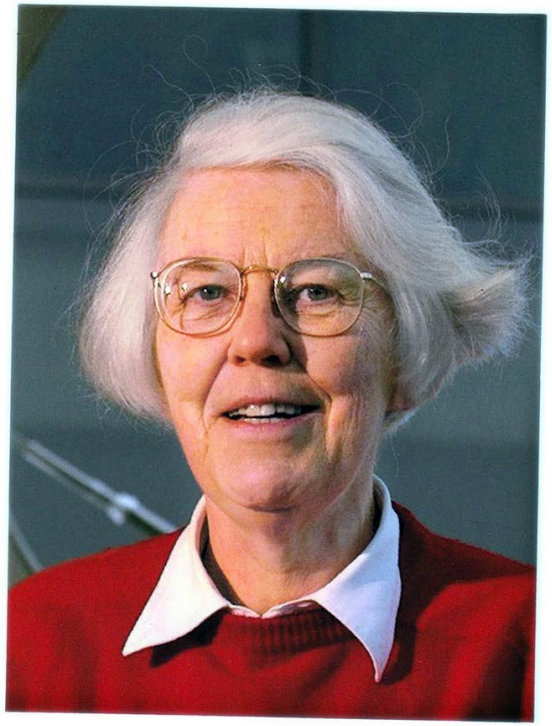

Karen Sparck Jones in 2002. She created a system that allowed people to search for data using words instead of code, establishing the basis of search engines.
Computer Laboratory/University of Cambridge

Jan. 2, 2019 15

Since 1851, obituaries in The New York Times have been dominated by white men. With Overlooked, we're adding the stories of remarkable people whose deaths went unreported in The Times.

https://www.nytimes.com/2019/01/02/obituaries/karen-sparck-jones-overlooked.html

Similarity based on word counts?

Need to deal with similar words (synonyms)

Need to reduce the number of words (too high dimensional)

Find **groups** of words

Pre-defined dictionaries (word types)

word usage analysis

Statistical groupings (words that tend to go together)

Topic Modeling

Topic modeling

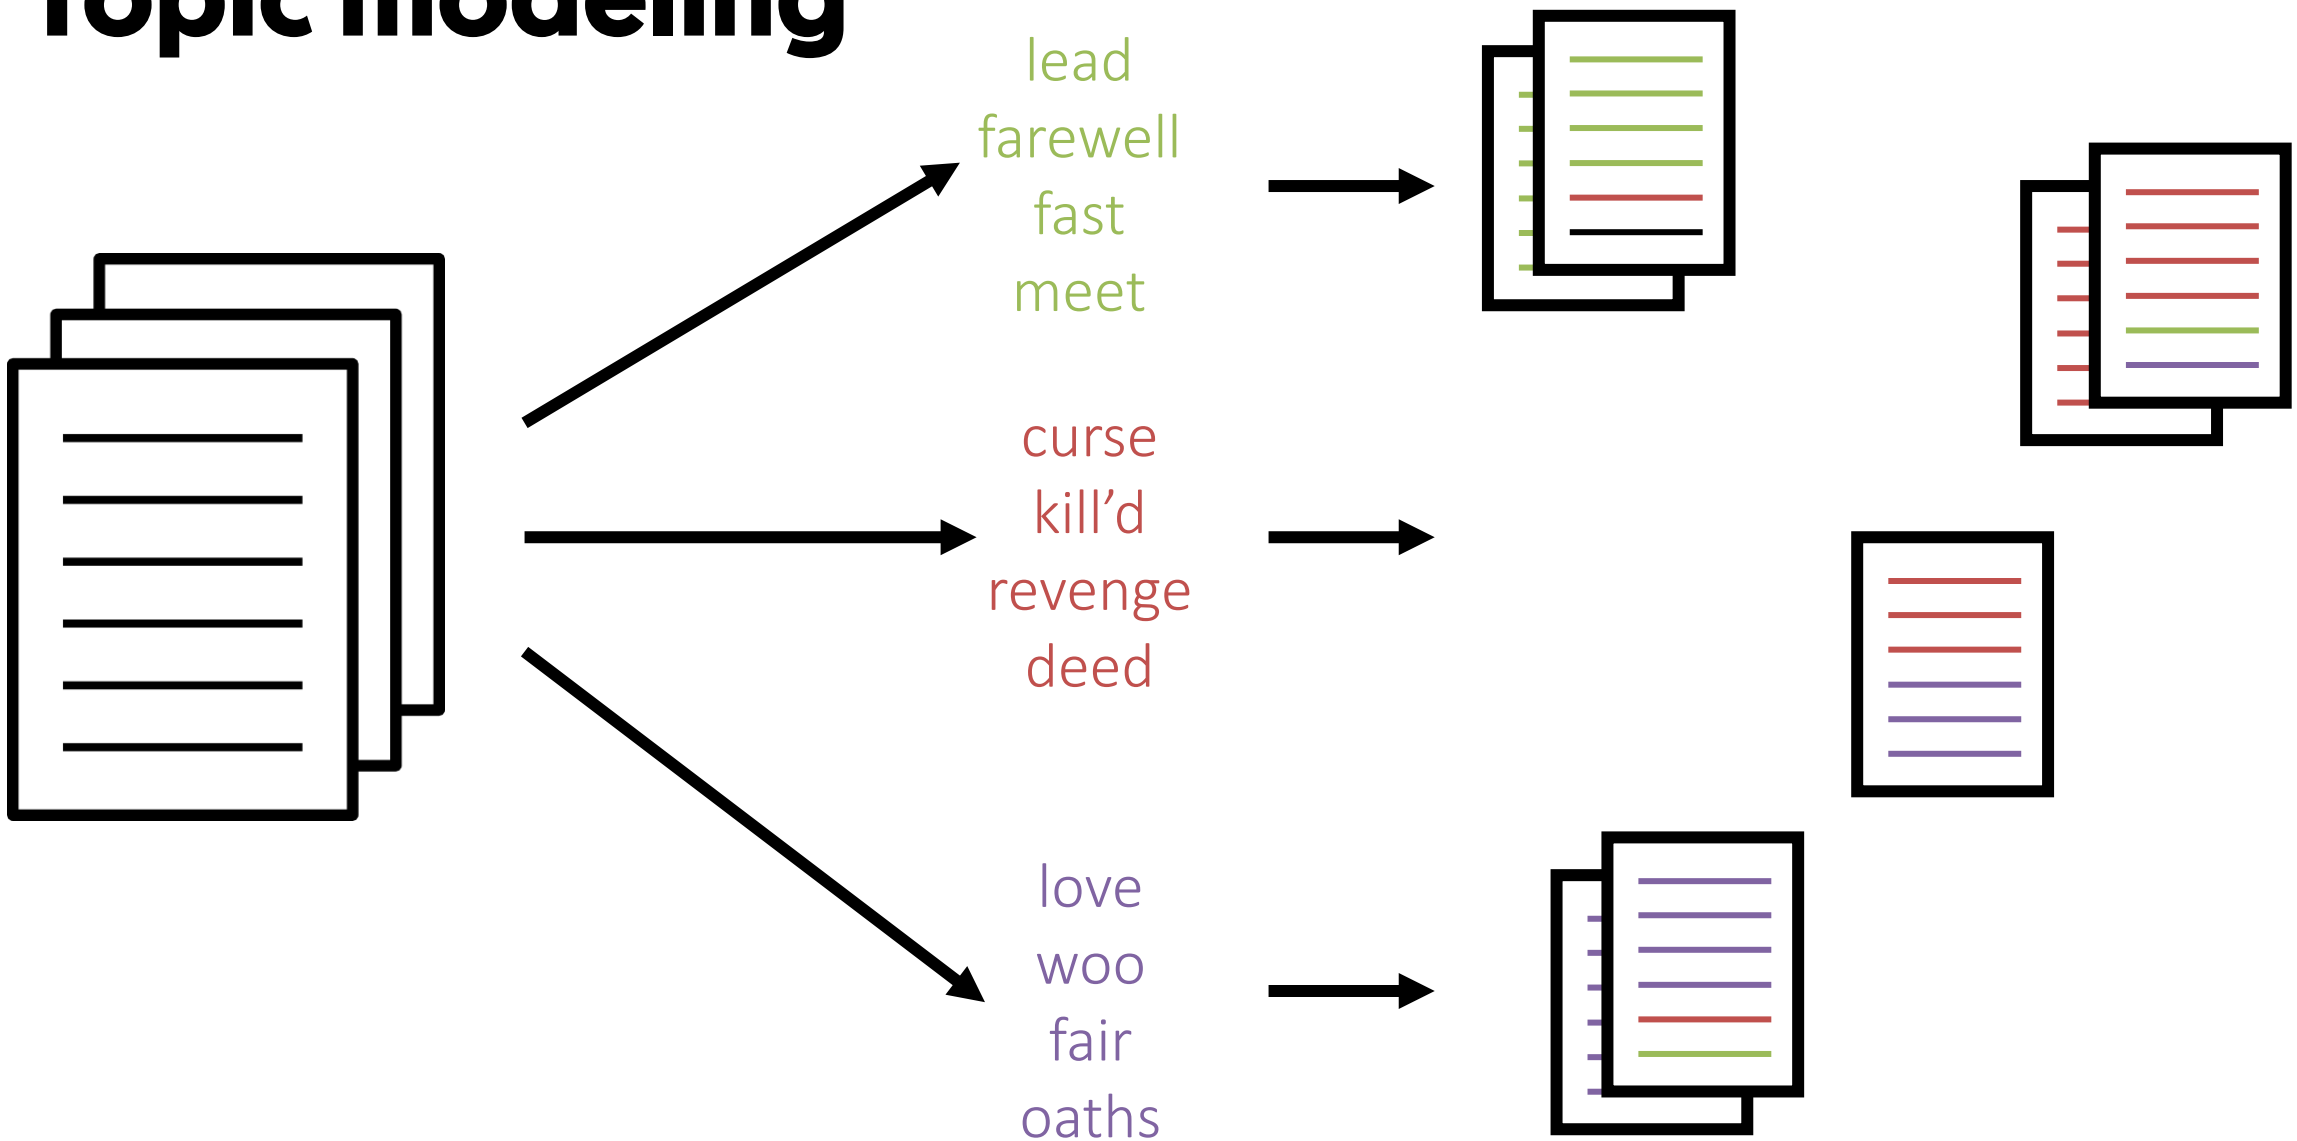

Word vector embeddings

Place **words** in a high-dimensional **vector** space

Words **similar in meaning** should be **close in space**

Infer similarity by **distributional semantics**:

similar context implies similar meaning

```
my pet cat is brown
my pet dog is brown
my big car is brown
```

Construct embeddings by processing a **corpus** of text

Several ways to build word embeddings

Word2Vec

Skip-gram model

Neural embedding

GLoVE

Co-occurrence model

Factor matrix by optimization

Several ways to build topic models (document embeddings)

LDA (Latent Dirichlet Allocation)

Standard algorithm

Iterative and probabalistic

NMF (Non-Negative Matrix Factors)

Fast Algorithm

Less widely adopted

Why use Word Vector Embeddings?

Learn about
Language or Corpora (Texts)

Find similar words/synonyms

Track changes of word usage

Exploring polysemy

Creating lexical resources

Evidence of bias

...

Natural Language Applications
Pre-Processing

Translation

Sentiment Analysis

Interpretation

...

Why Interpret Embeddings?

Gain insight into a model

Did you build a good model?

Gain insight into the modeling process

Which methods are better?

Gain insight about the underlying data

What does the model tell us?

Challenges of Word Vector Embeddings

Challenges of Document Embeddings

Large numbers of Words [or documents]

High-dimensional spaces

Complex relationships – meaningless positions [or questionable]

Complex processes for building embeddings

Complex downstream applications

No ground truth - subjective aspects

Interactive Analysis of Word Vector Embeddings

Florian Heimerl and Michael Gleicher
Department of Computer Sciences
University of Wisconsin – Madison
EuroVis 2018

Summary:

Word vector embeddings offer unique challenges

Task analysis of needs

characteristic	single target	multiple targets
similarity	(1) inspect local neighborhood	(2) compare local neighborhoods
average, offset	(3) inspect arithmetic structure	(4) compare arithmetic results
co-oc. probability	(5) analyze encoded probabilities	(6) compare probabilities
concept axis	(7) analyze vector relations	(8) compare vector relations
multiple	(9) discover interesting aspects	(10) compare interesting aspects

3 Designs for unmet needs

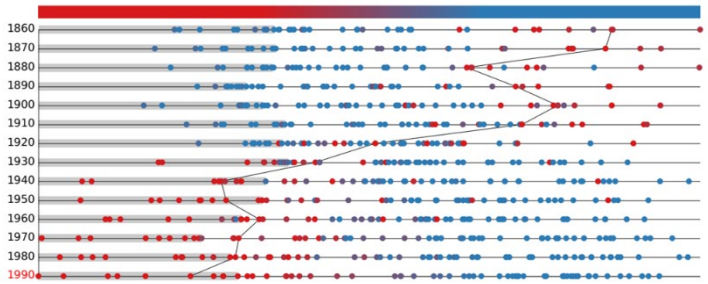

Buddy Plots

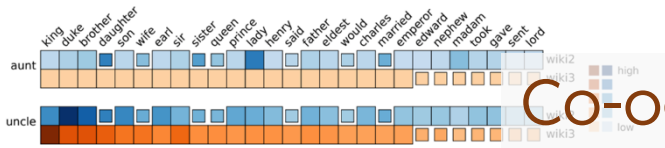

Co-occurance Matrices

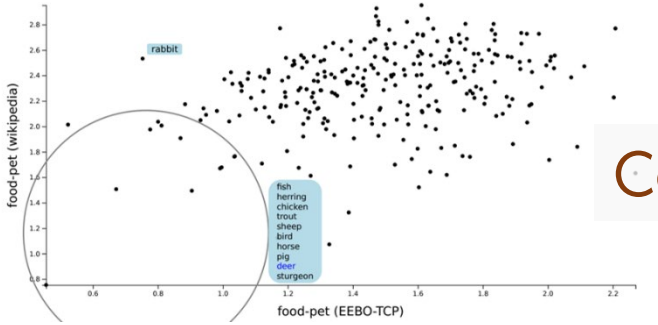

Concept Axis Plots

Why visualization?

Tasks are inherently human-centric

Variety of tasks involving interpretation
Linguistic and domain knowledge for applications

But what are those tasks?

Task Analysis:
What do people do with Word Embeddings?

Literature survey

111 papers from diverse communities

Consider use cases in Linguistics, HCI, Digitial Humanities, etc.

Augment this list by extrapolation

what tasks would users want – but aren't doing yet

Use cases suggest tasks

Learn about
Language or Corpora (Texts)

Find similar words/synonyms

Track changes of word usage

Exploring polysemy

Creating lexical resources

Natural Language Applications
Pre-Processing

Translation

Sentiment Analysis

Interpretation

Interpretation

Identify items of interest

Probe values of interest

Evaluation

Intrinsic (good embedding?)

Extrinsic (applications success?)

Linguistic Tasks and Characteristics

We identified 7 distinct *linguistic tasks* within the literature

4 characteristics pertinent to those tasks: similarity, arithmetic structures, concept axis, and co-occurrences

linguistic tasks	characteristics	examples
rank word pairs	similarity	[BDK14, PSM14]
compare concepts	average, similarity	[RBS17, SLMJ15]
find analogies	offset, similarity	[SLMJ15, LG14]
view neighbors	similarity	[HLJ16, YWL* 16]
select synonyms	similarity	[BDK14, FDJ* 14]
project based on concepts	concept axis	[BCZ* 16, FRMW17]
predict contexts	co-oc. probability	[SN16, LJW* 15]

Prior Work:
Word Vector Embedding Visualizations

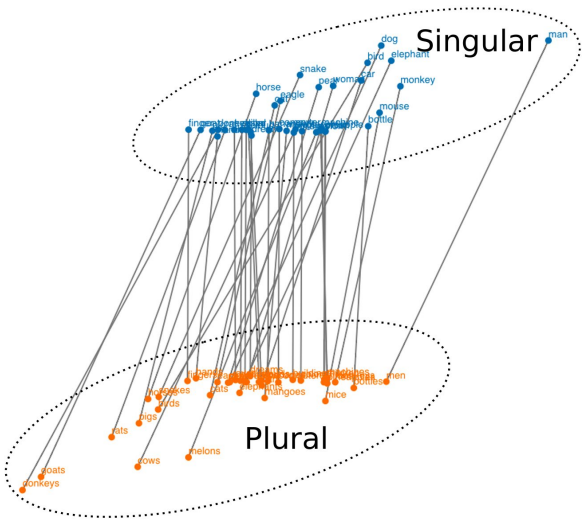

Liu et al., InfoVis 2017
Analogy Relationships

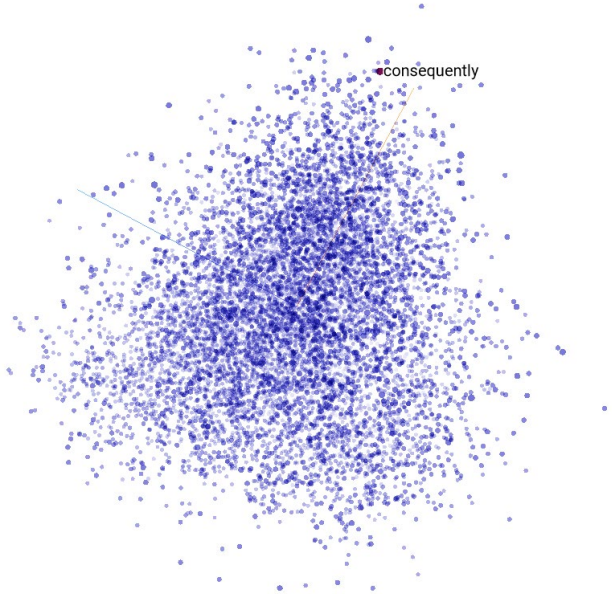

Smilkov et al., 2016
Distance Relationships

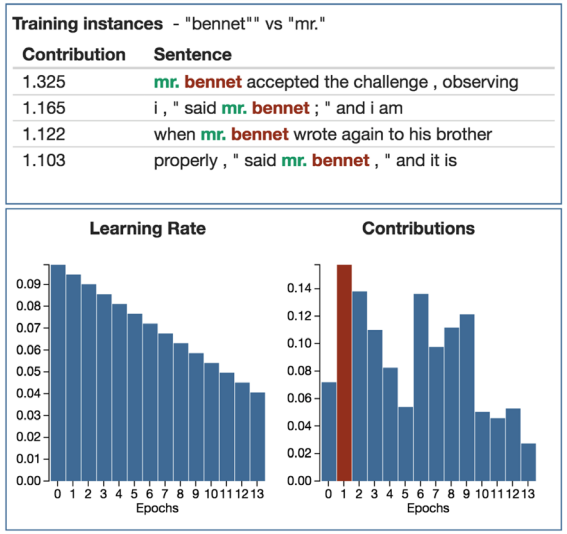

Rong and Adar, 2016
Training Process

Tasks and Characteristics vs. Visualizations

Rank word pairs	similarity
View neighbors	similarity ← Smilkov, et al 2016
Select synonyms	similarity
Compare concepts	average, similarity
Find analogies	offset, similarity ← Liu, et al 2017
Project on Concept	concept axis
Predict Contexts	co-oc. probability

Tasks and Characteristics vs. Visualizations

Rank word pairs	similarity
View neighbors	**similarity** ← Smilkov, et al 2016
Select synonyms	**similarity**
Compare concepts	**average, similarity**
Find analogies	offset, similarity ← Liu, et al 2017
Project on Concept	**concept axis**
Predict Contexts	**co-oc. probability**

Tasks we seek designs for in this paper

1. Similarities (local distances)

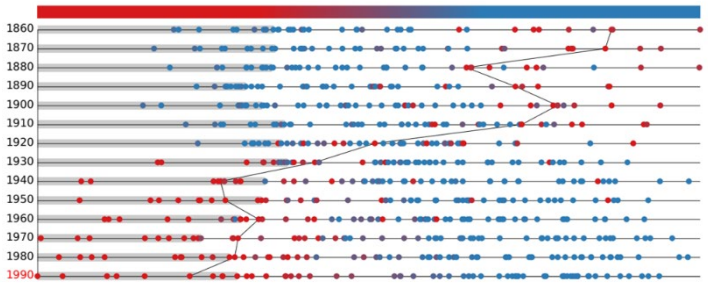

Buddy Plots

2. Co-Occurrances

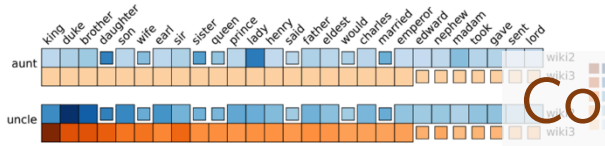

Co-occurance Matrices

3. Concept Axes

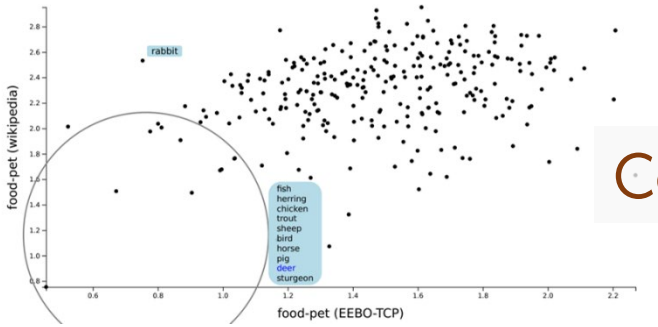

Concept Axis Plots

1. Similarities (local distances)

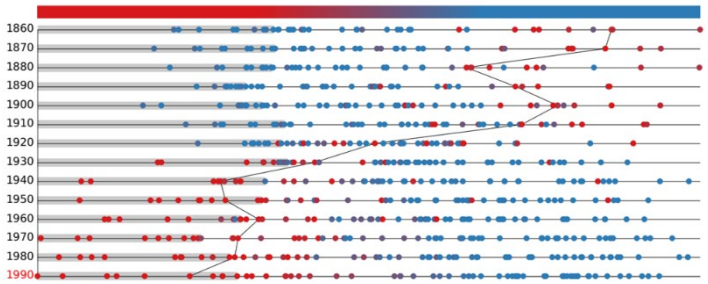

Buddy Plots

2. Co-Occurrances

Co-occurance Matrices

3. Concept Axes

Concept Axis Plots

Similarities: Understanding local distances

Distances are meaningful

- even if absolute values are not

What is close to a word?

Are there groups of words that are similar?

Ordered lists are useful

- Density (how many can you show)
- Sense of relative distances
- Comparison between words

Embedding Projector
Smilkov, et al. 2016

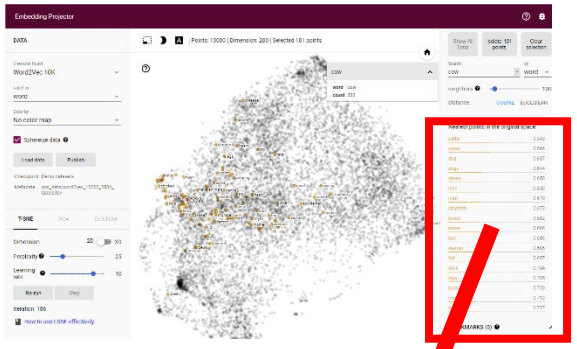

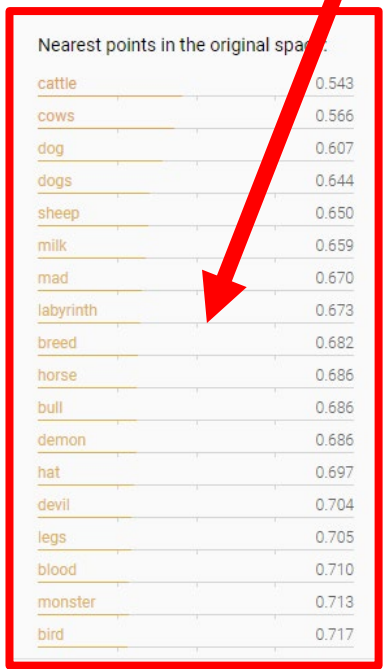

Buddy Plots (1D lists)

Map distance (to selected reference) to horizontal axis

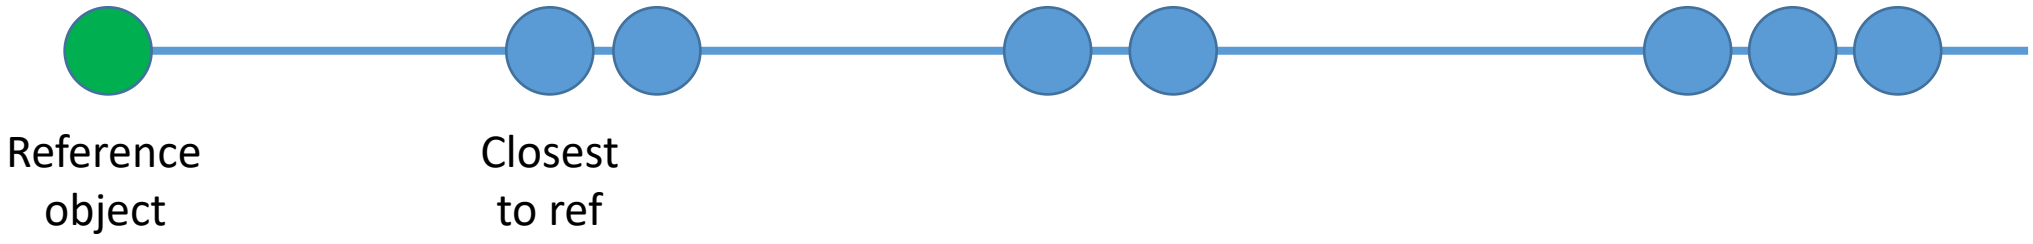

Alexander and Gleicher, 2016 – for Topic Models

Not dimensionality reduction?

Map distance to word to horizontal axis

Focus on a single point – other relations not preserved

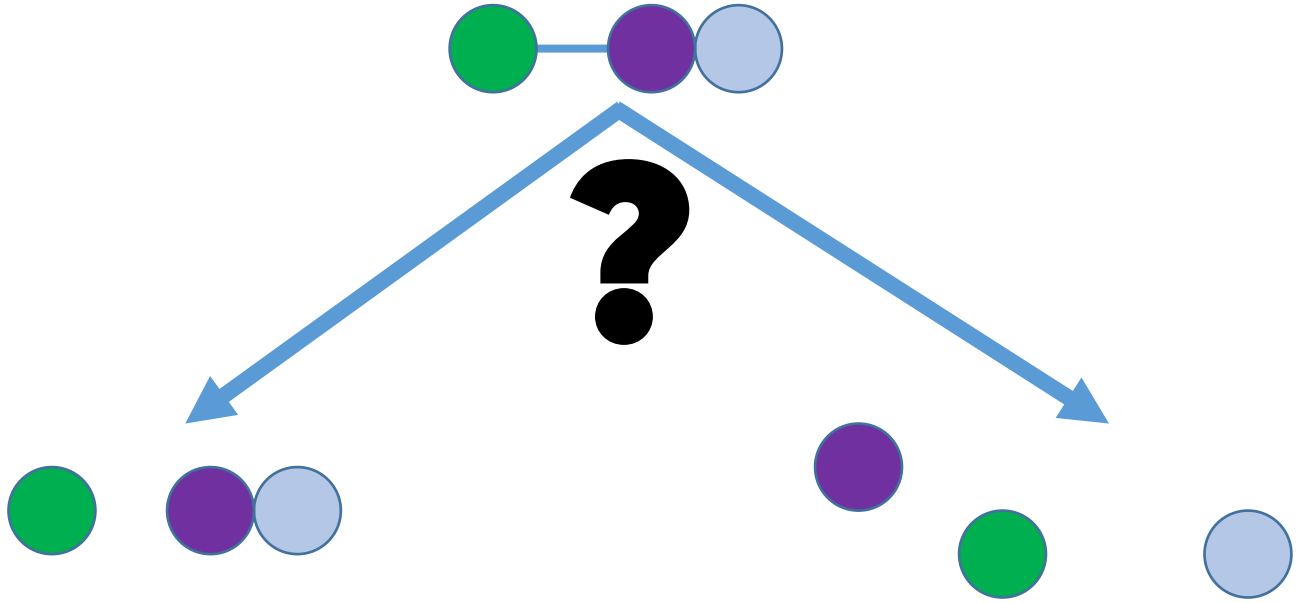

Embedding Projector
Smilkov, et al. 2016

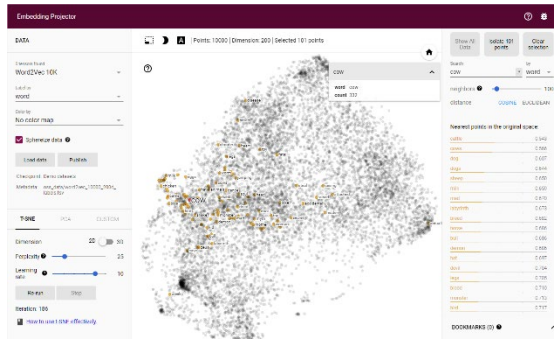

Nearest points in the original space:

cattle	0.543
cows	0.566
dog	0.607
dogs	0.644
sheep	0.650
milk	0.659
mad	0.670
labyrinth	0.673
breed	0.682
horse	0.686
bull	0.686

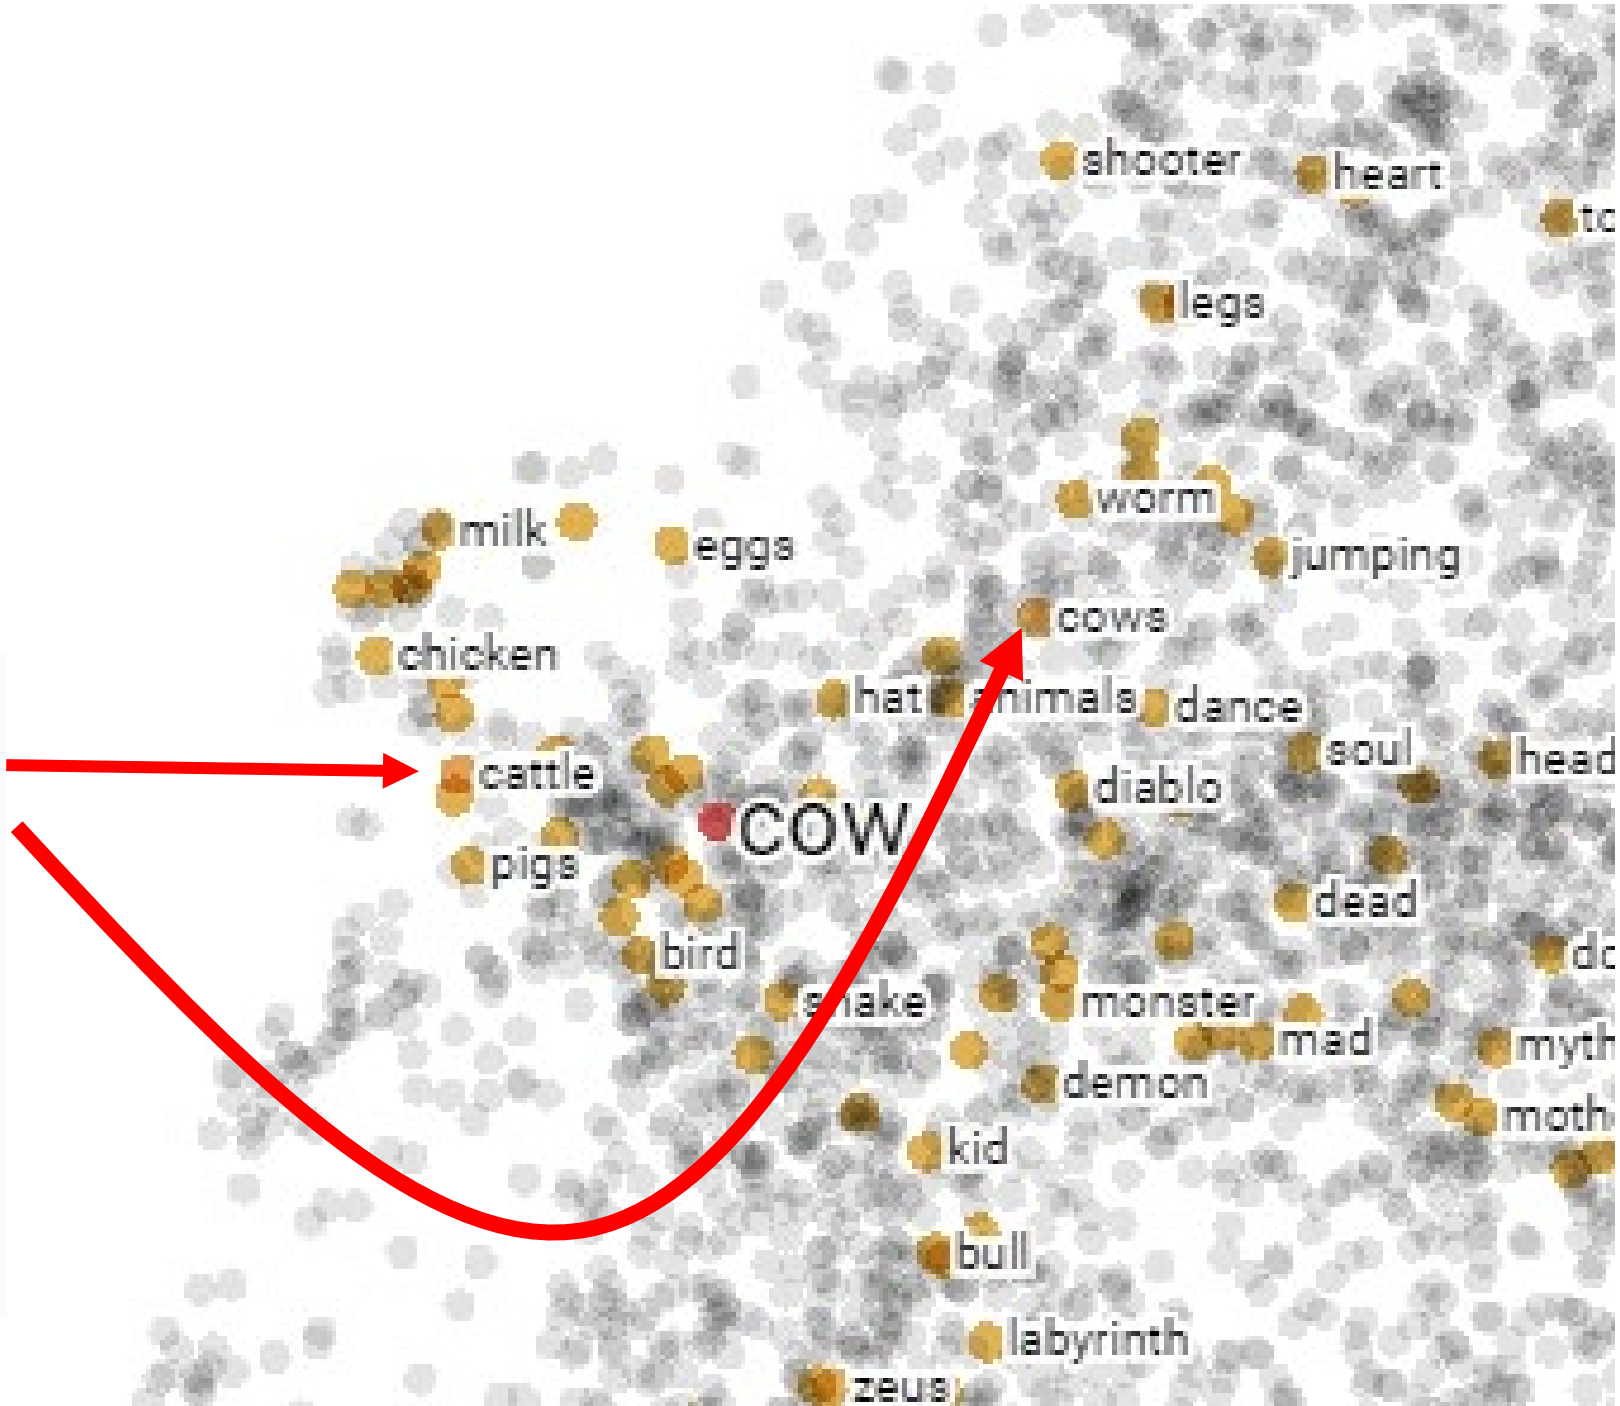

Stacked/Chained buddy plots

Use color to encode distance

Stacked/Chained buddy plots

Use color to encode distance **in the reference row**

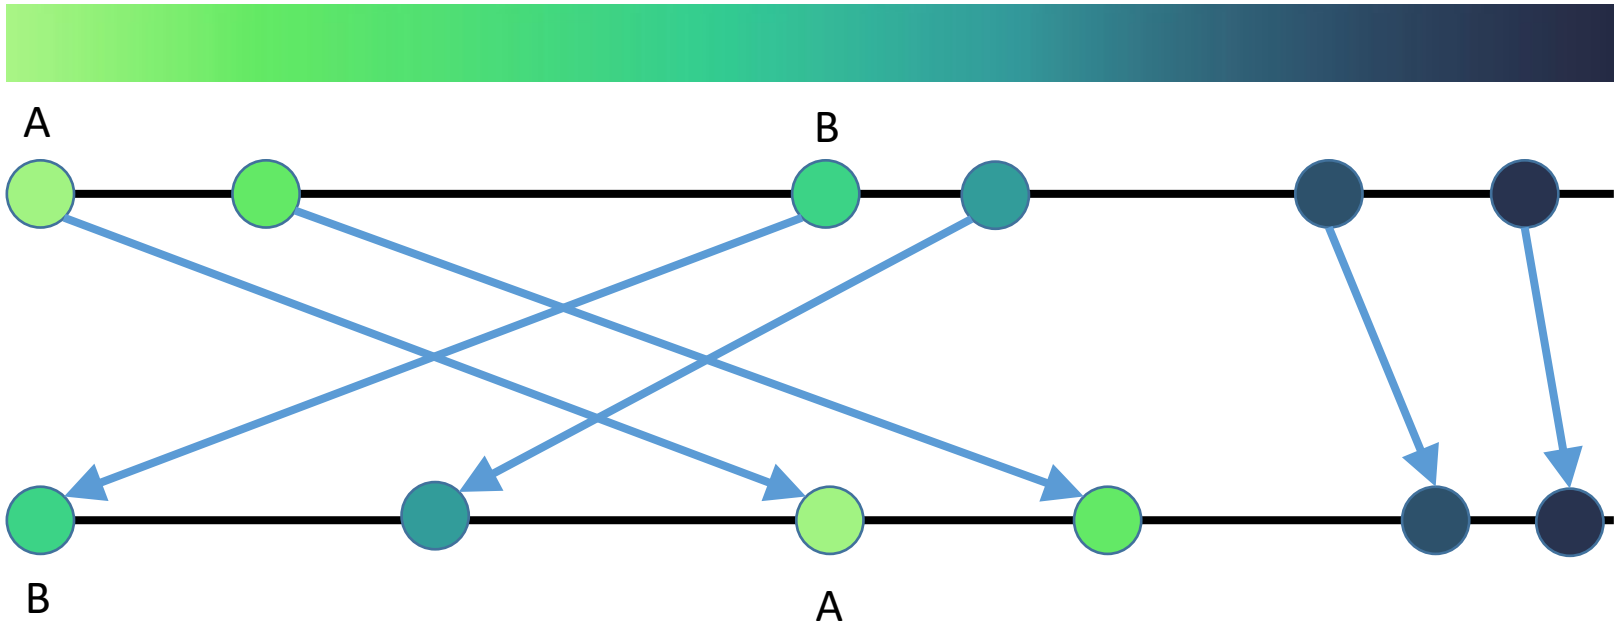

Stacked/Chained buddy plots

Use color to encode distance **in the reference**

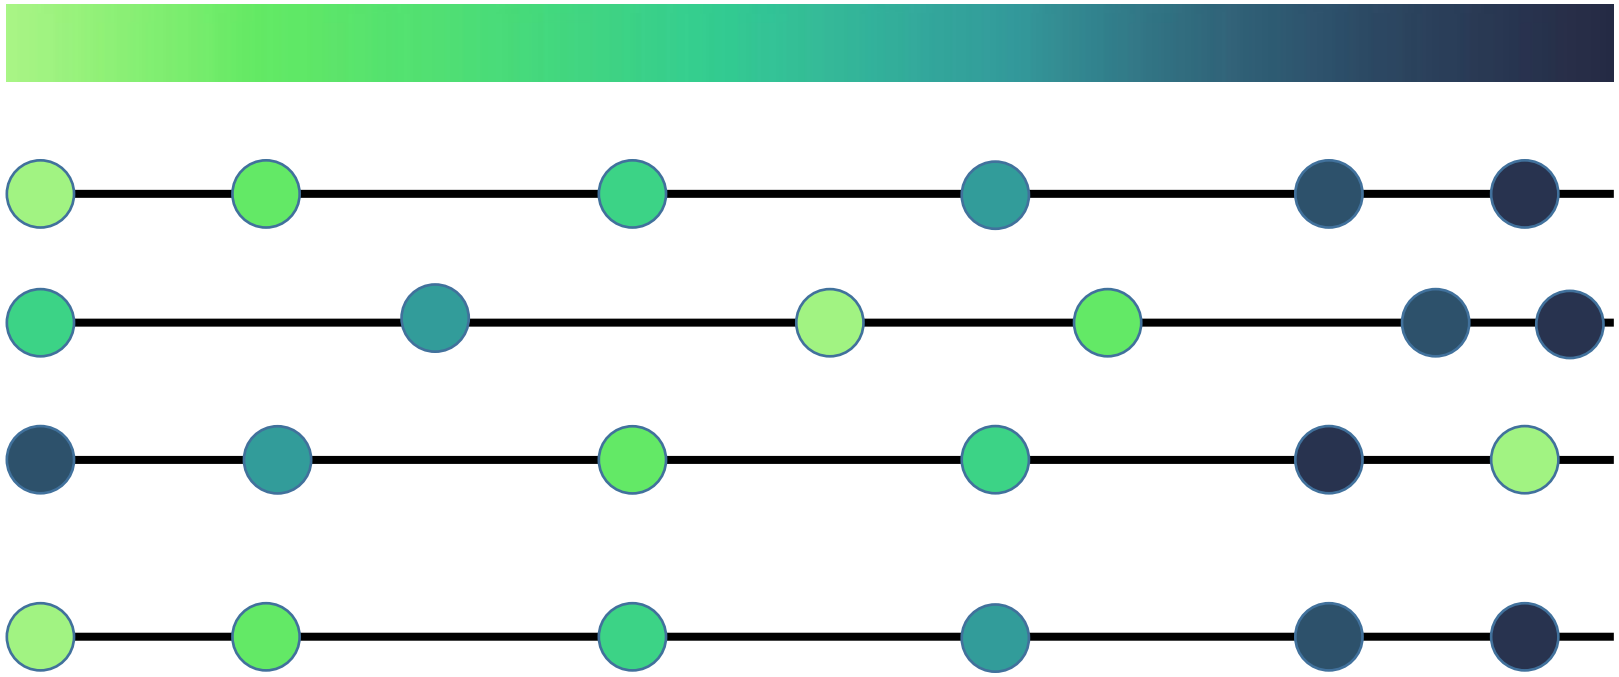

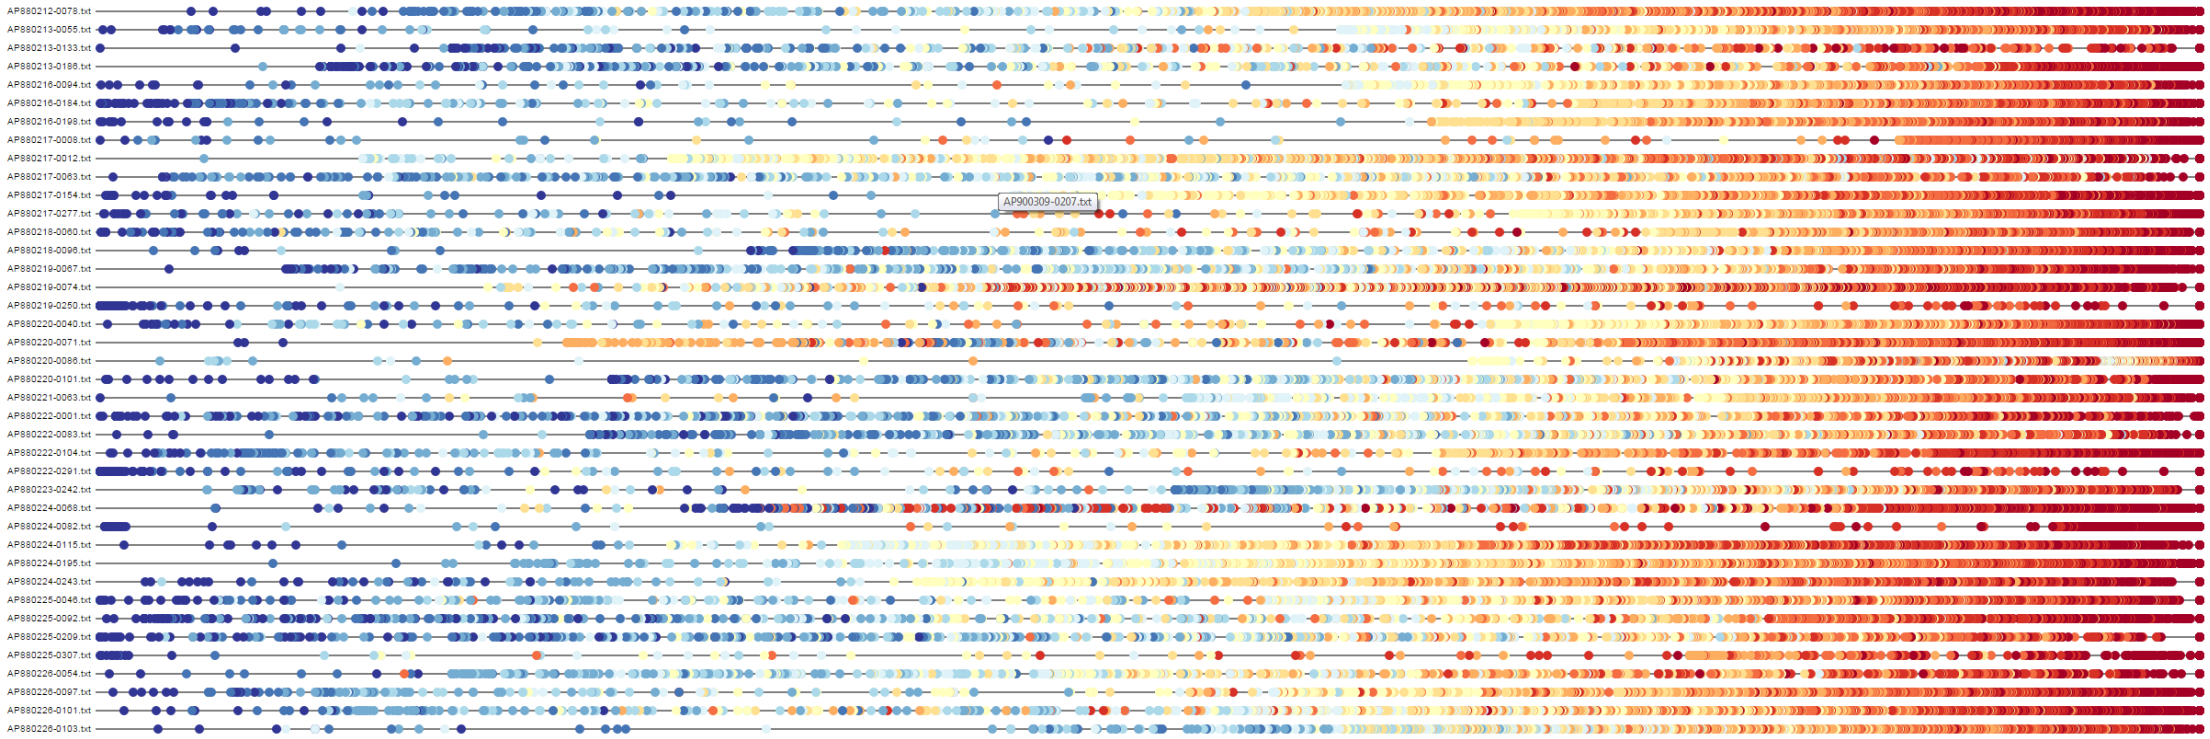
AP880212-0078.txt
AP880213-0055.txt
AP880213-0133.txt
AP880213-0186.txt
AP880216-0094.txt
AP880216-0184.txt
AP880216-0198.txt
AP880217-0008.txt
AP880217-0012.txt
AP880217-0063.txt
AP880217-0154.txt
AP880217-0277.txt
AP880218-0060.txt
AP880218-0096.txt
AP880219-0067.txt
AP880219-0074.txt
AP880219-0250.txt
AP880220-0040.txt
AP880220-0071.txt
AP880220-0086.txt
AP880220-0101.txt
AP880221-0063.txt
AP880222-0001.txt
AP880222-0083.txt
AP880222-0104.txt
AP880222-0291.txt
AP880223-0242.txt
AP880224-0068.txt
AP880224-0082.txt
AP880224-0115.txt
AP880224-0195.txt
AP880224-0243.txt
AP880225-0046.txt
AP880225-0092.txt
AP880225-0209.txt
AP880225-0307.txt
AP880226-0054.txt
AP880226-0097.txt
AP880226-0101.txt
AP880226-0103.txt
AP900309-0207.txt

Same word... different embedding

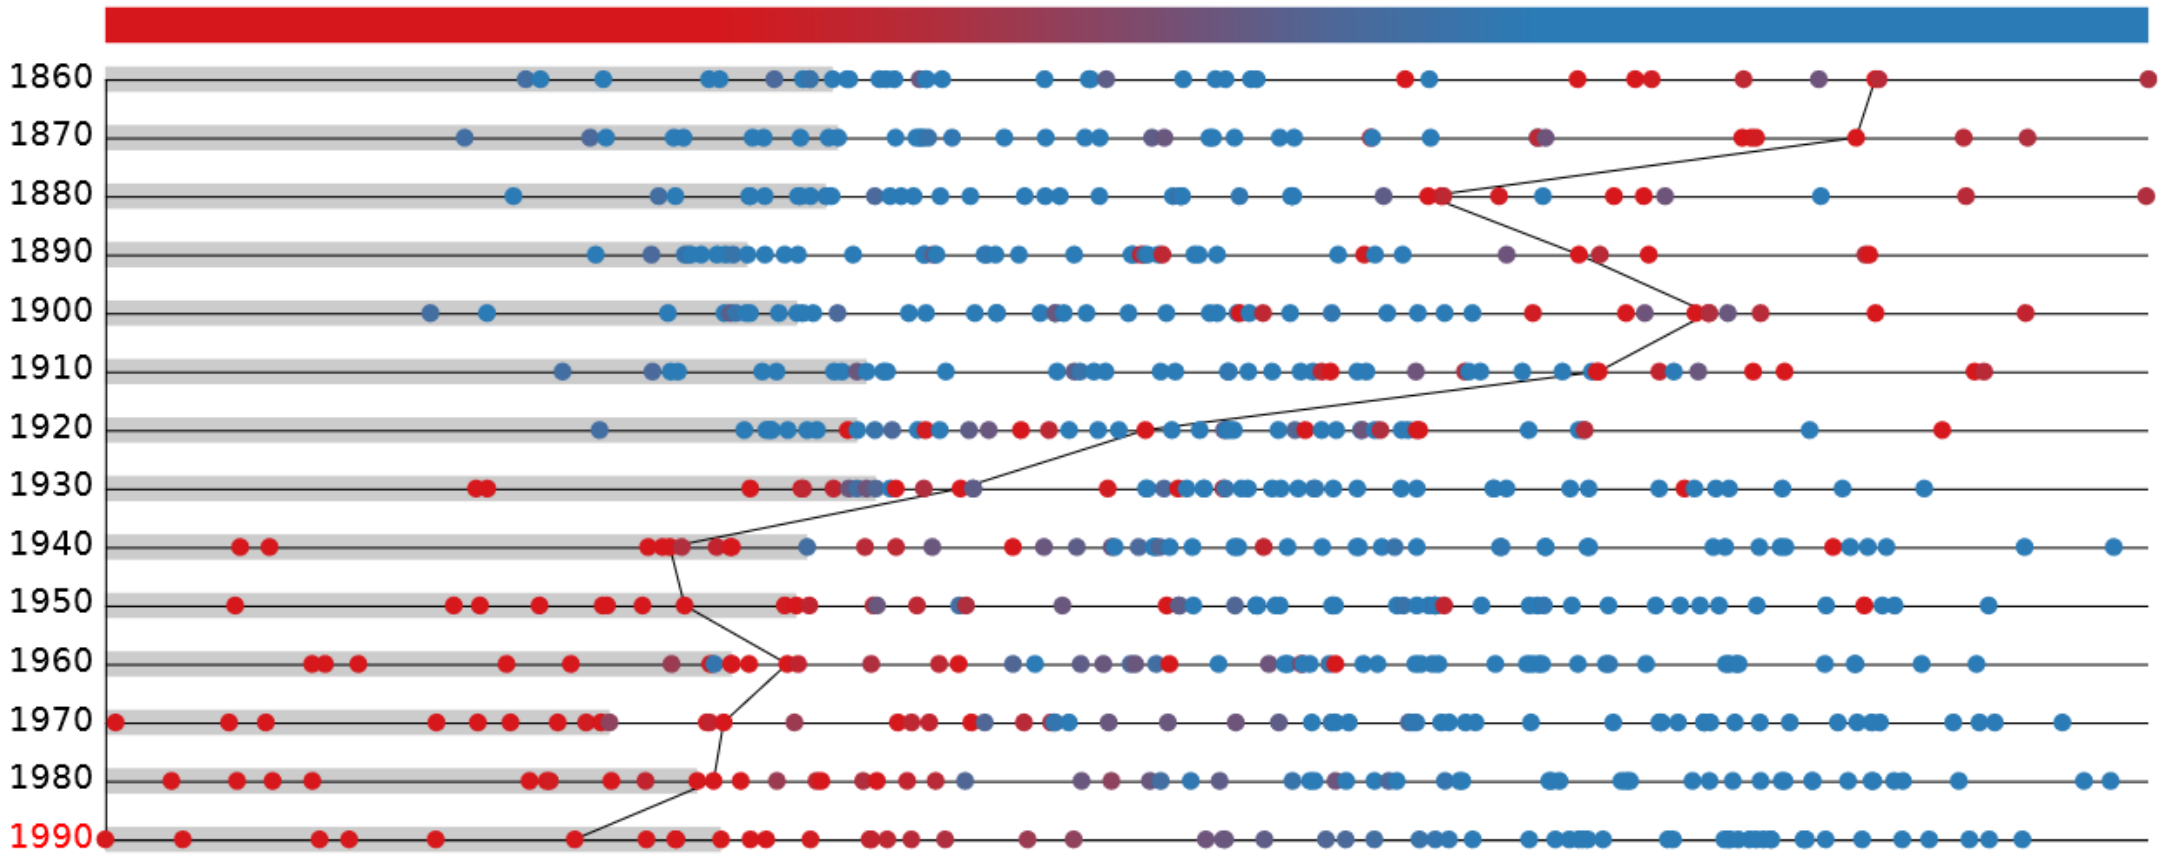

Broadcast

1. Similarities (local distances)

Buddy Plots

2. Co-Occurrances

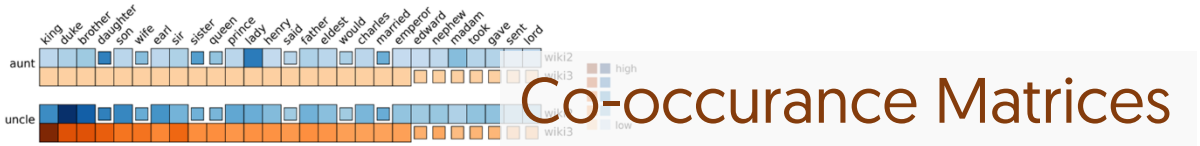

Co-occurance Matrices

3. Concept Axes

Concept Axis Plots

Why are words similar? Understandng word co-occurence

Similarity based on co-occurrence

- count how often one word occurs near another

Co-occurrence matrix

- main form of input data
- many models approximate the matrix (reconstruct)

Useful for understanding and diagnosing models

Co-occurrence view

How do we view the massive matrix?

- color encoding (heat map) – density, relative values
- select rows (specify words of interest)
- select columns (metrics of interestingness – given rows)
  - highest values
  - highest variance

Co-occurrence matrix view

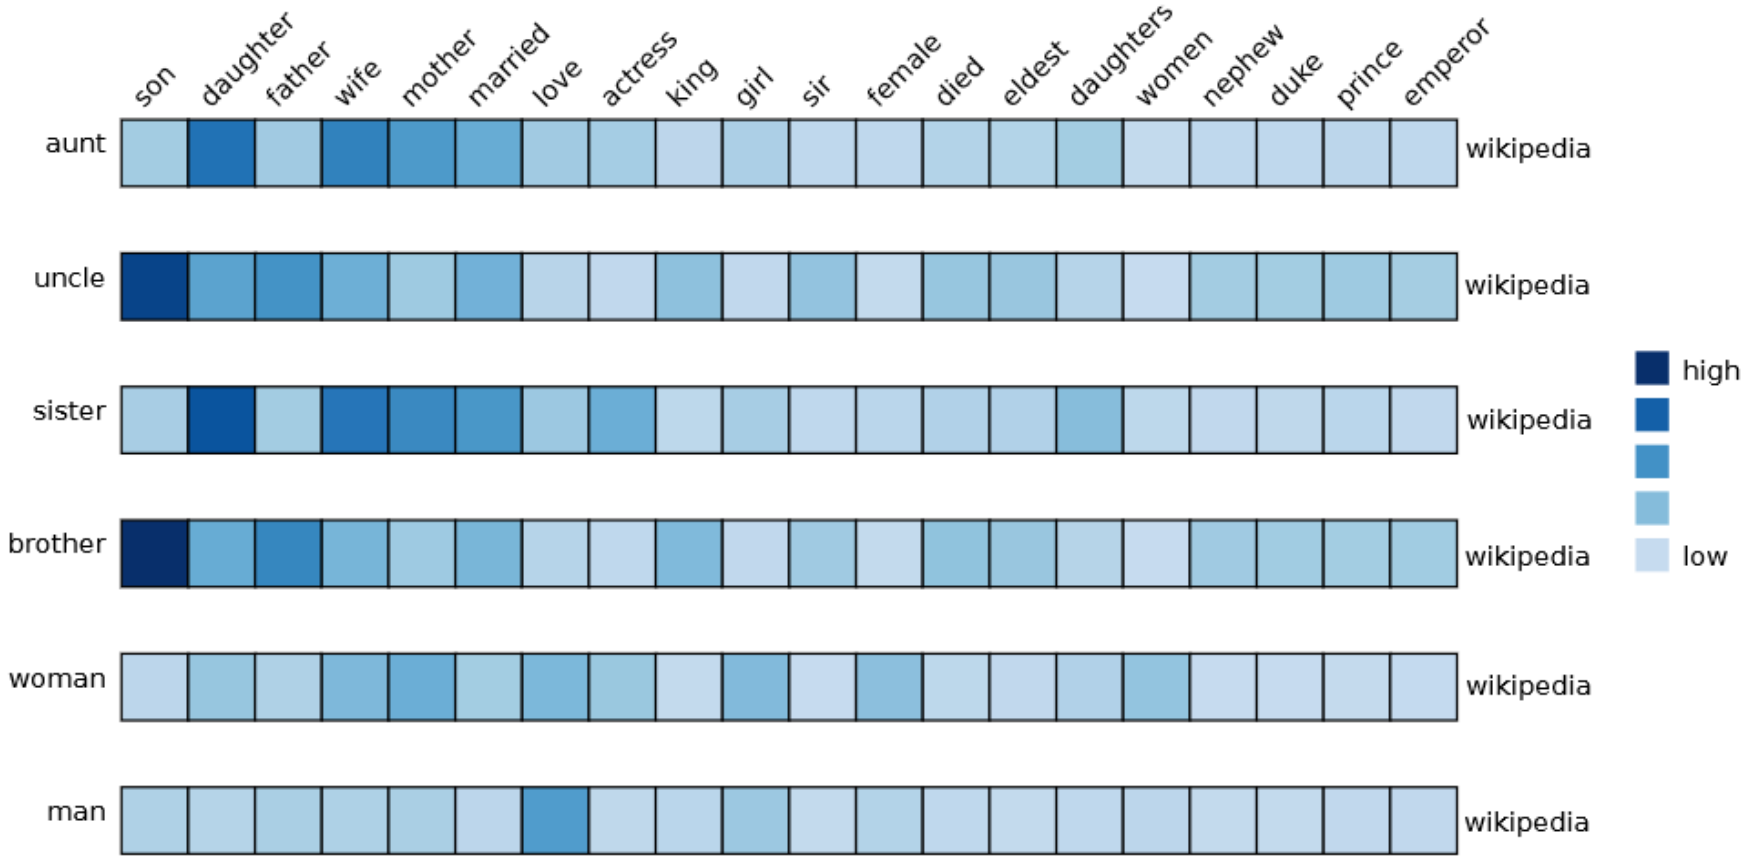

1. Similarities (local distances)

Buddy Plots

2. Co-Occurrances

Co-occurance Matrices

3. Concept Axes

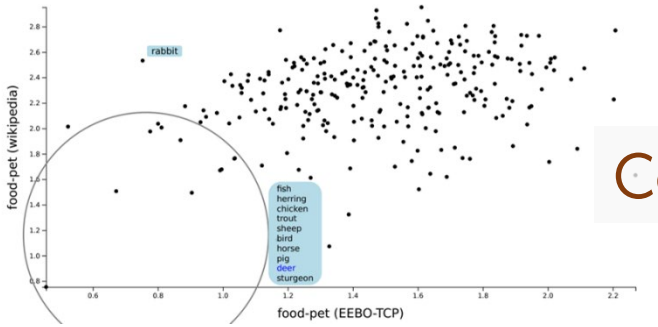

Concept Axis Plots

Concept Axes: Understanding Semantic Directions

Opposing concepts make an axis

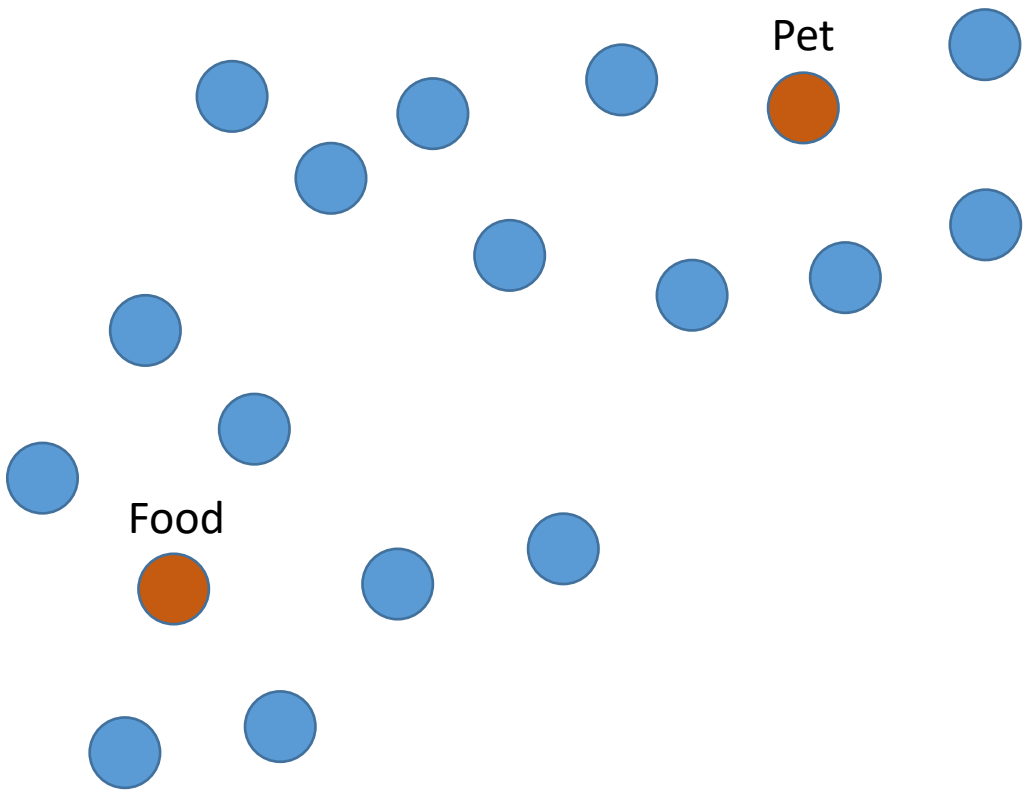

Concept Axes: Understanding Concept Axes

Define an axis from one concept to another

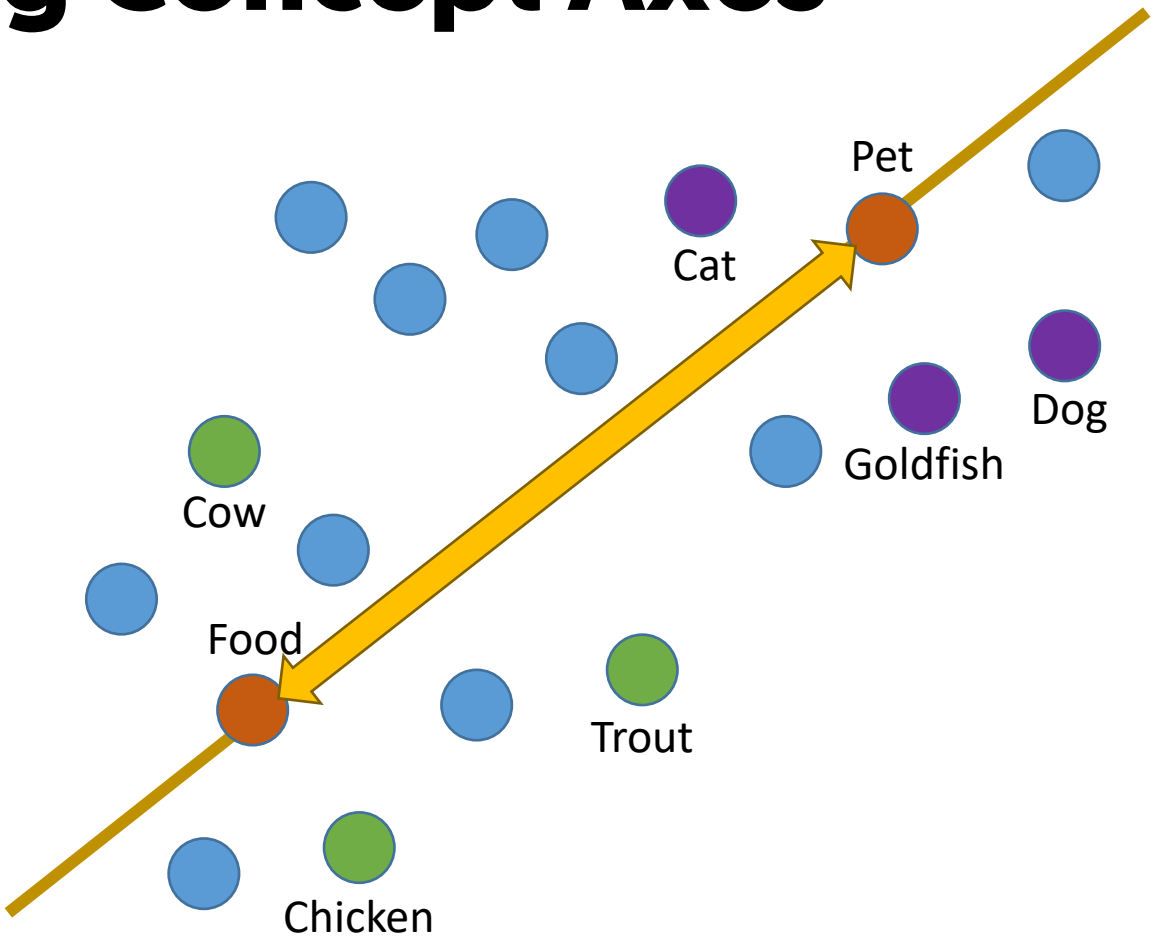

Ways to define axes

Vector between two concepts

Interaxis - Kim et al., 2015

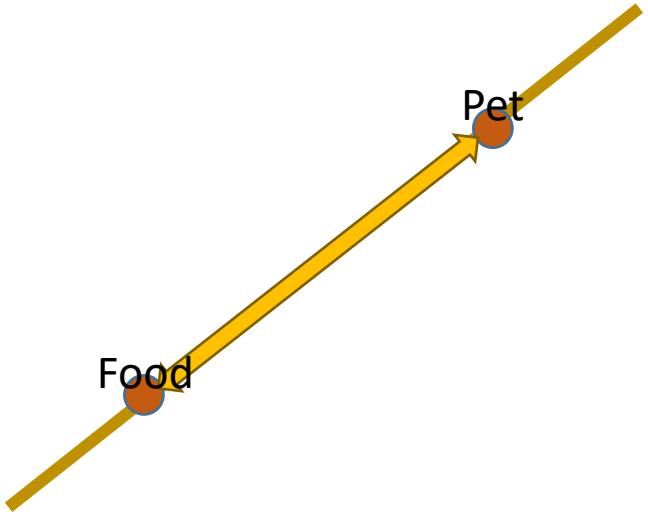

Classifier between two groups

Explainers – Gleicher, 2013

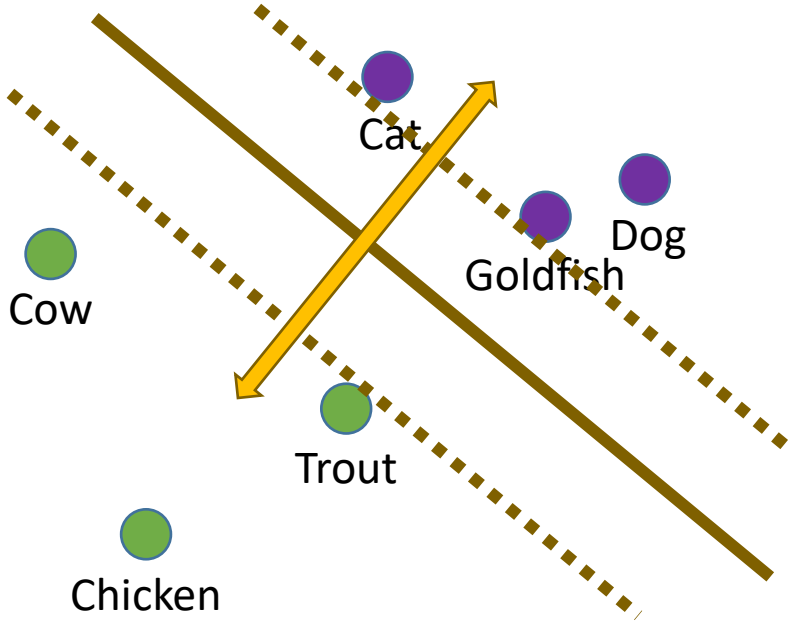

Multiple Concept Axes

Use multi-variate plots

2D = Scatterplot

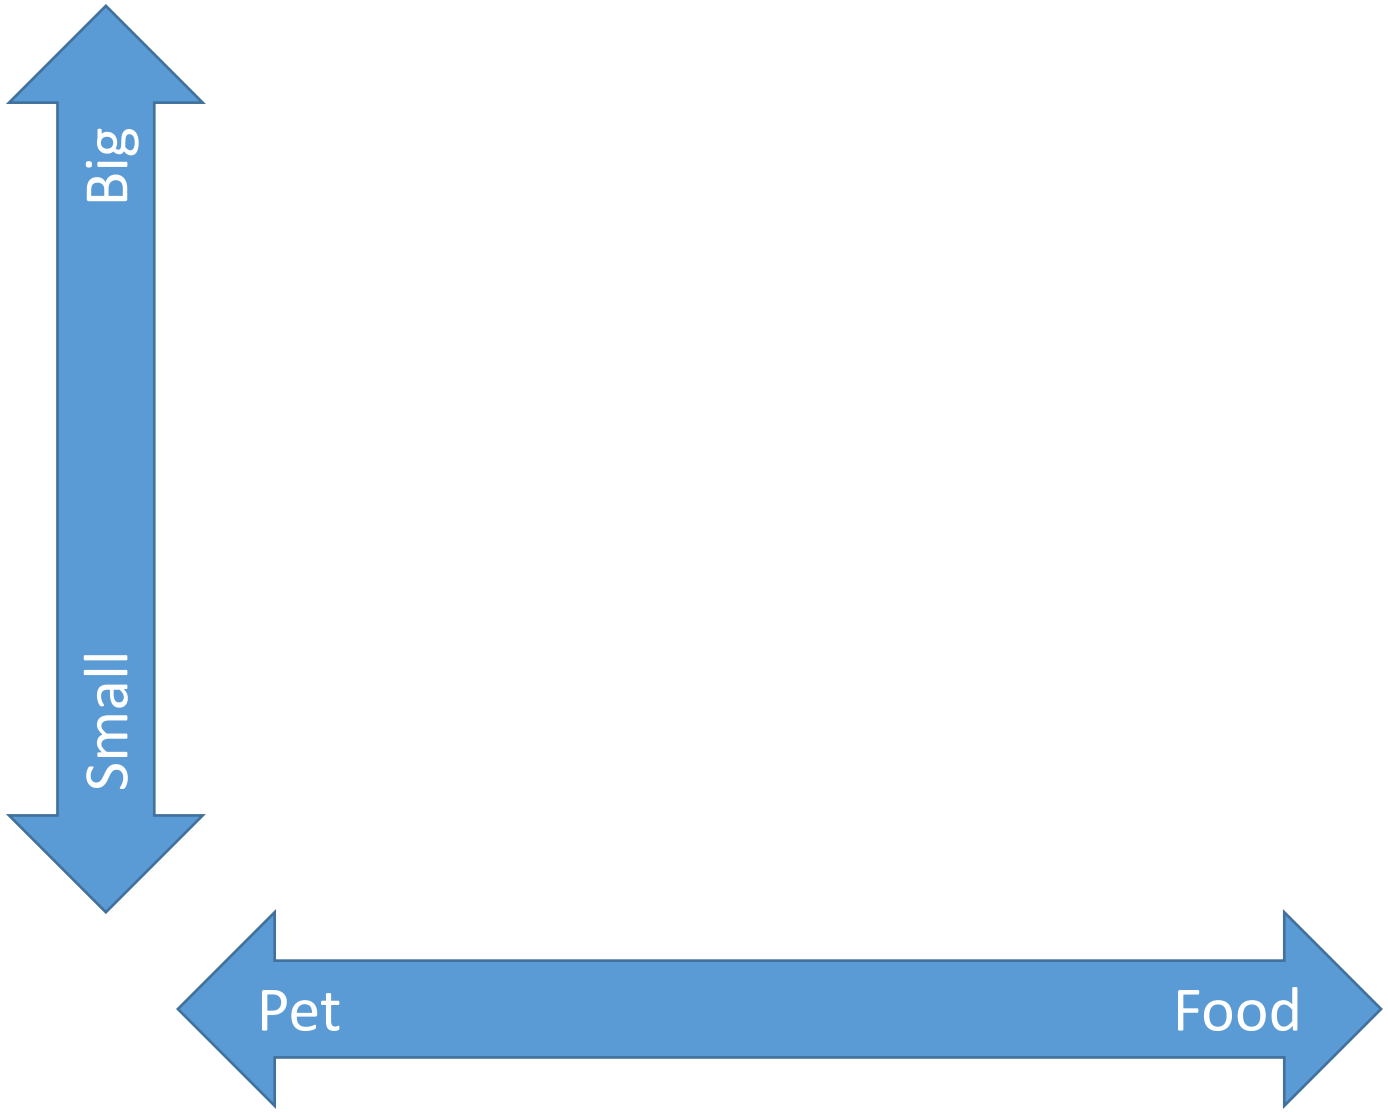

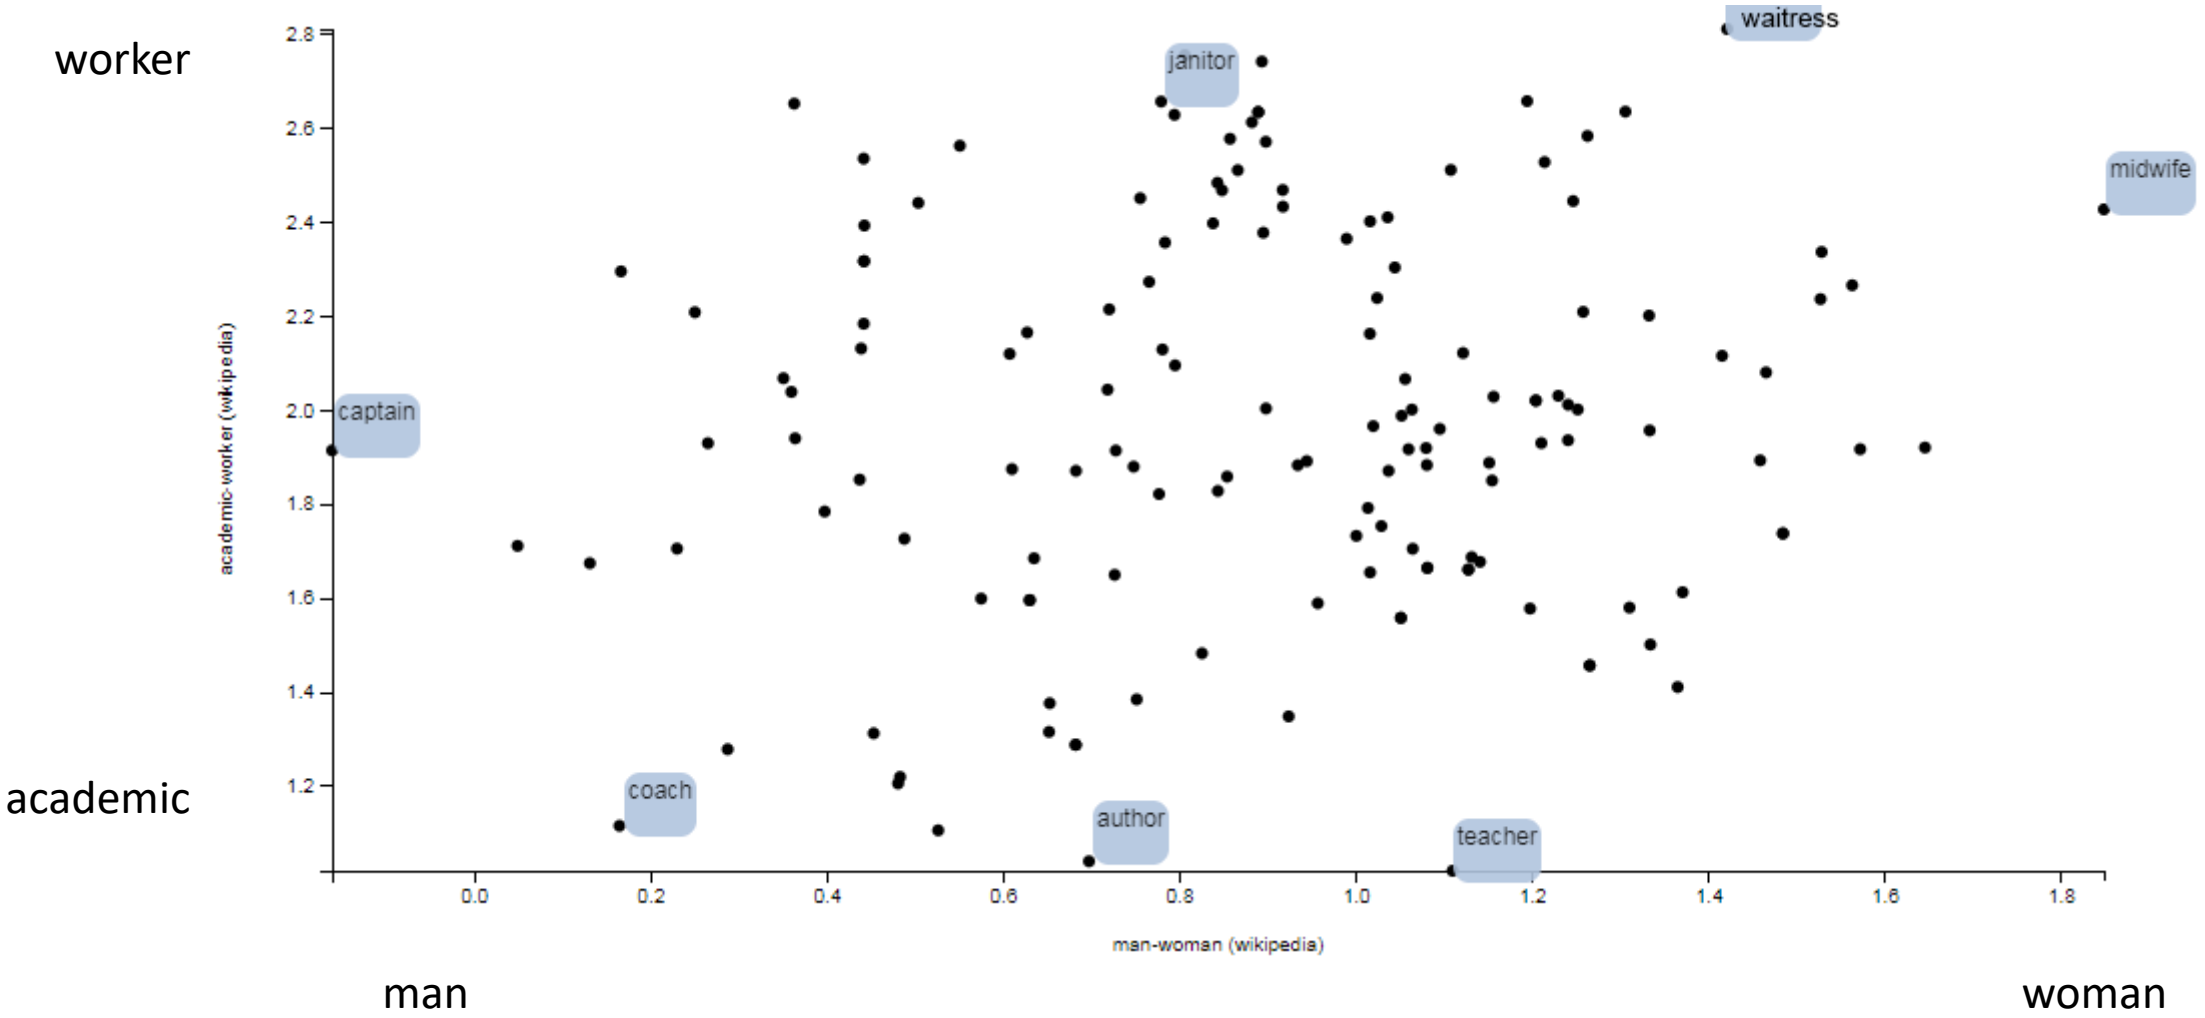
worker
academic
man
woman
waitress
janitor
midwife
captain
coach
author
teacher
2.8
2.6
2.4
2.2
2.0
1.8
1.6
1.4
1.2
academic-worker (wikipedia)
0.0
0.2
0.4
0.6
0.8
1.0
1.2
1.4
1.6
1.8
man-woman (wikipedia)

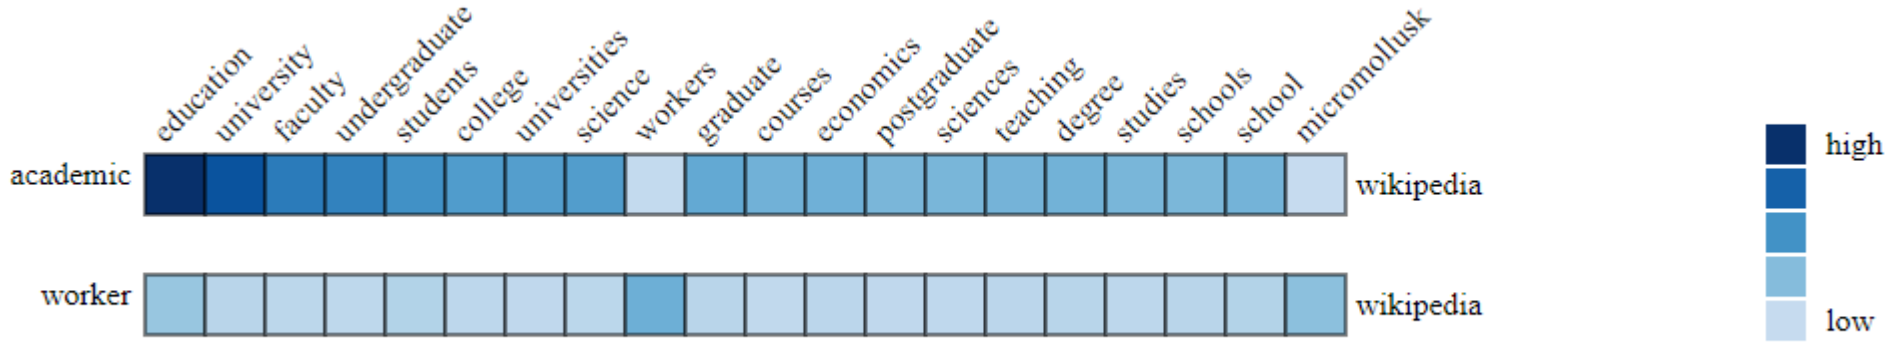
education
university
faculty
undergraduate
students
college
universities
science
workers
graduate
courses
economics
postgraduate
sciences
teaching
degree
studies
schools
school
micromollusk
academic
worker
wikipedia
wikipedia
high
low

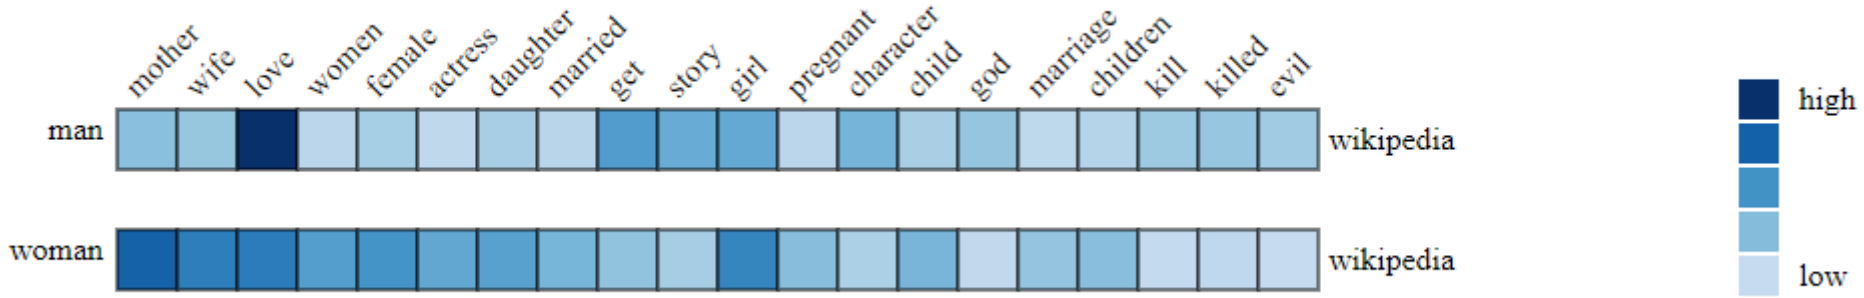
mother
wife
love
women
female
actress
daughter
married
get
story
girl
pregnant
character
child
god
marriage
children
kill
killed
evil
man
woman
wikipedia
wikipedia
high
low

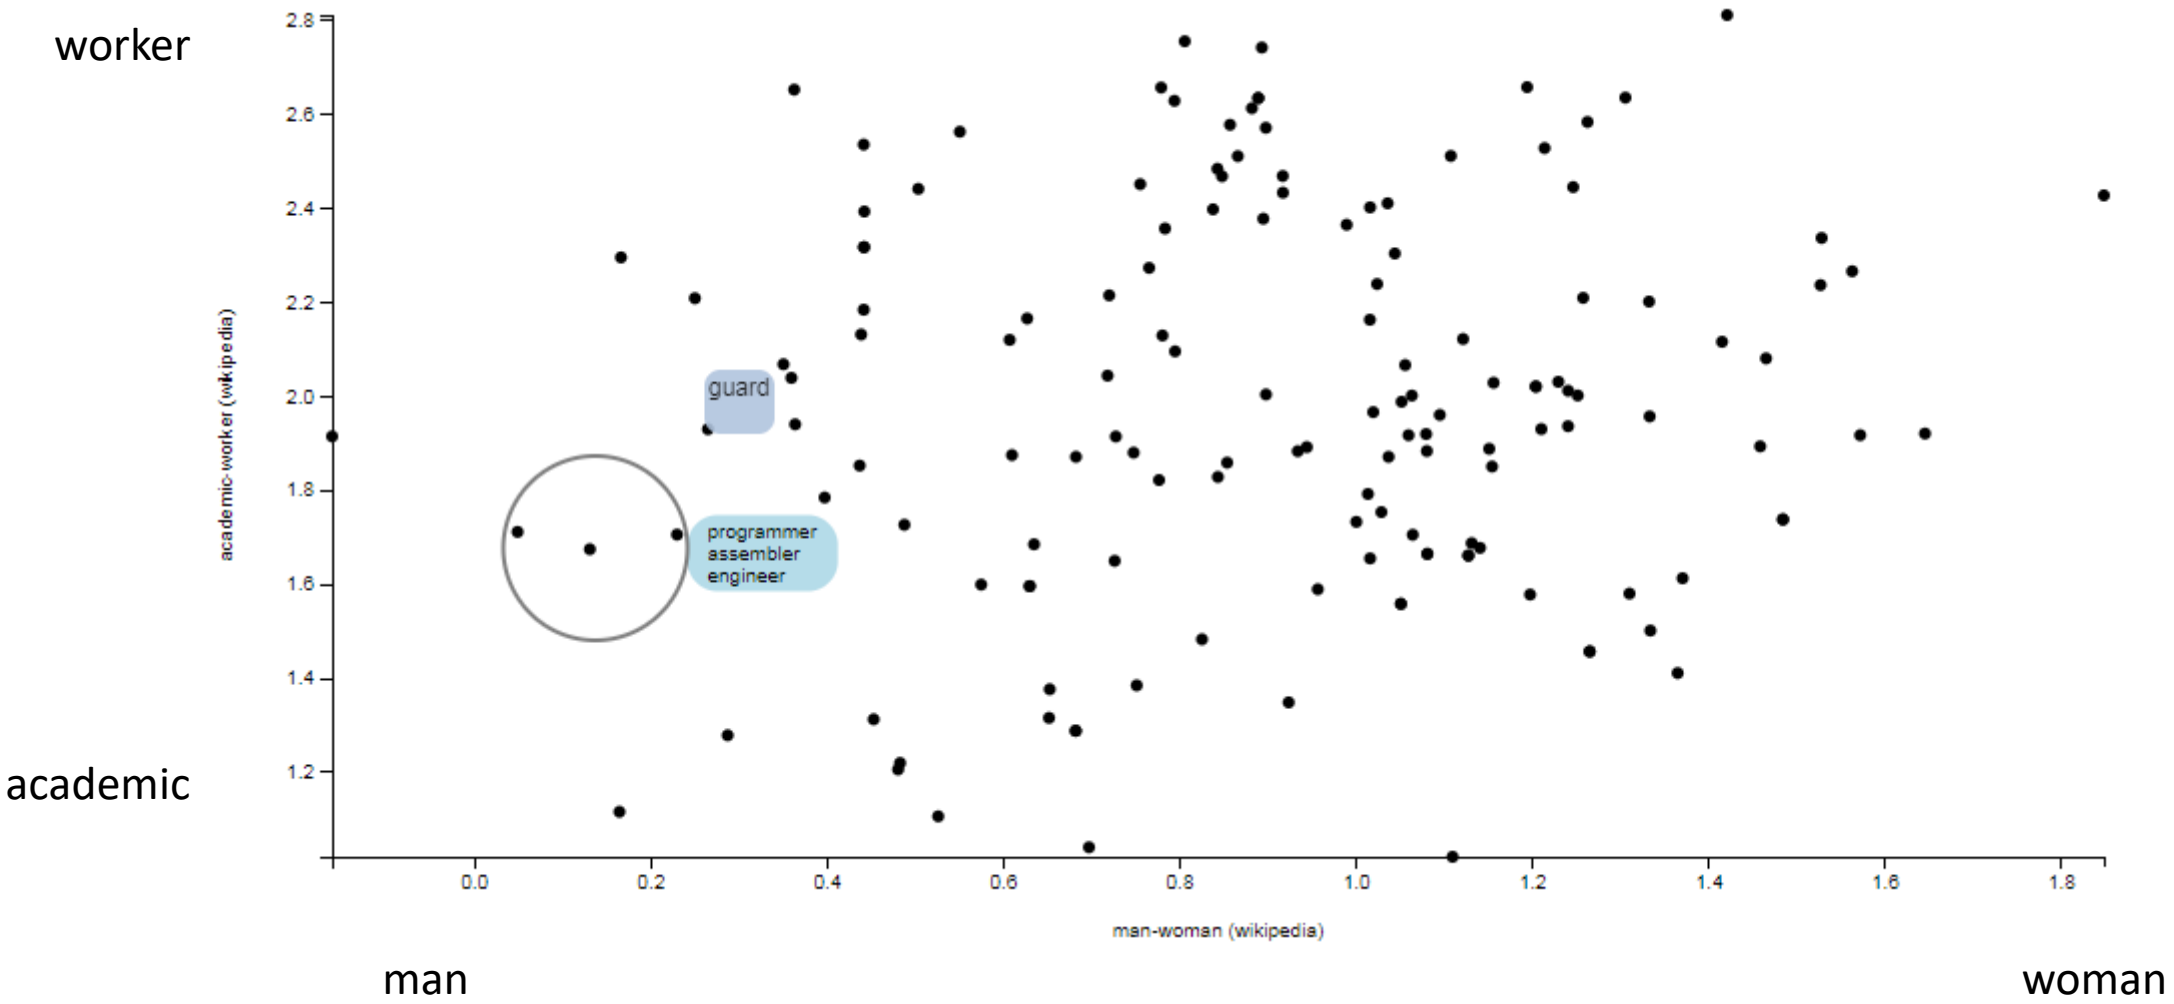
worker
academic
man
woman
guard
programmer
assembler
engineer
academic-worker (wikipedia)
man-woman (wikipedia)
2.8
2.6
2.4
2.2
2.0
1.8
1.6
1.4
1.2
0.0
0.2
0.4
0.6
0.8
1.0
1.2
1.4
1.6
1.8

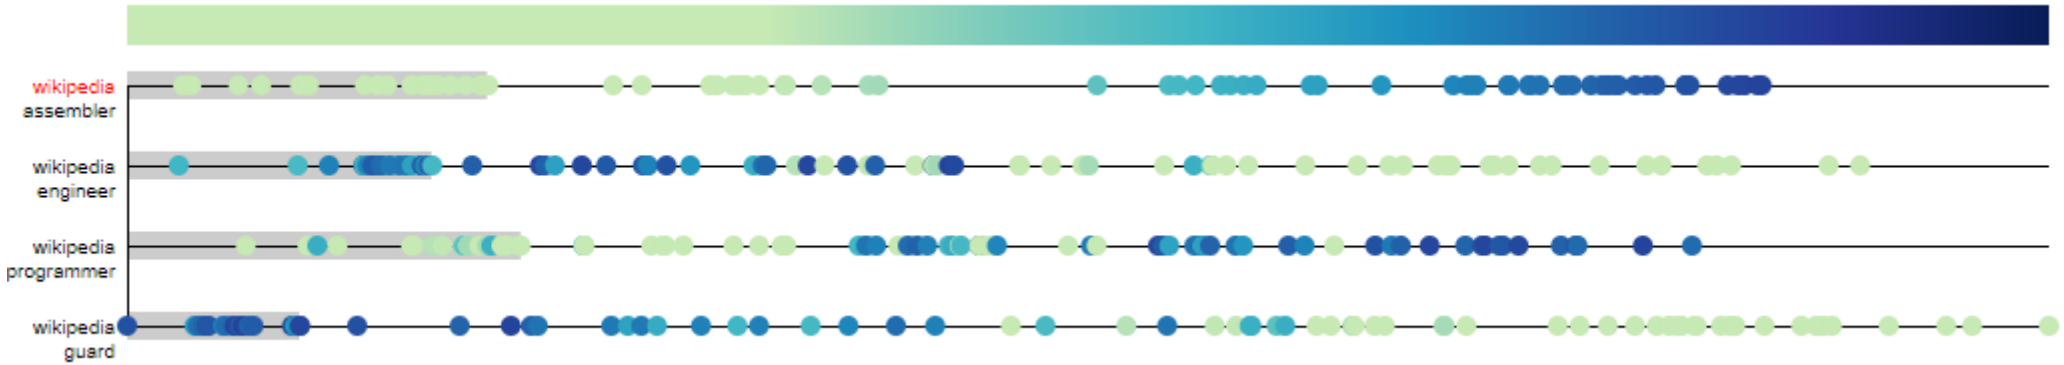
wikipedia
assembler
wikipedia
engineer
wikipedia
programmer
wikipedia
guard

wikipedia
assembler
fortran
compiler
lisp
cobol
compilers
gnu
macros
matlab
assemblers
javascript
debugger
templates
perl
bios
algol
emacs
implementations
runtime
executable
preprocessor
programmers
interpreter
pascal
interfaces
interface
software
input
functionality
computers
keyboard
hardware
computer
graphics
editing
hack
engineering
designed
engineers
inventor
hacker
developer
employed
aviation
command
assigned
worked
personnel
pioneer
military
force
guards
headed
army
marines
captain
duty
commanding
retired
soldiers

Limitations

Implementation	Usability and Scalability
Effectiveness	Evaluation (of designs)
Completeness	More Tasks Identifying Probes **Explicit Comparison** Connection to (model) evaluation Feedback to model building

Application: Word meaning change

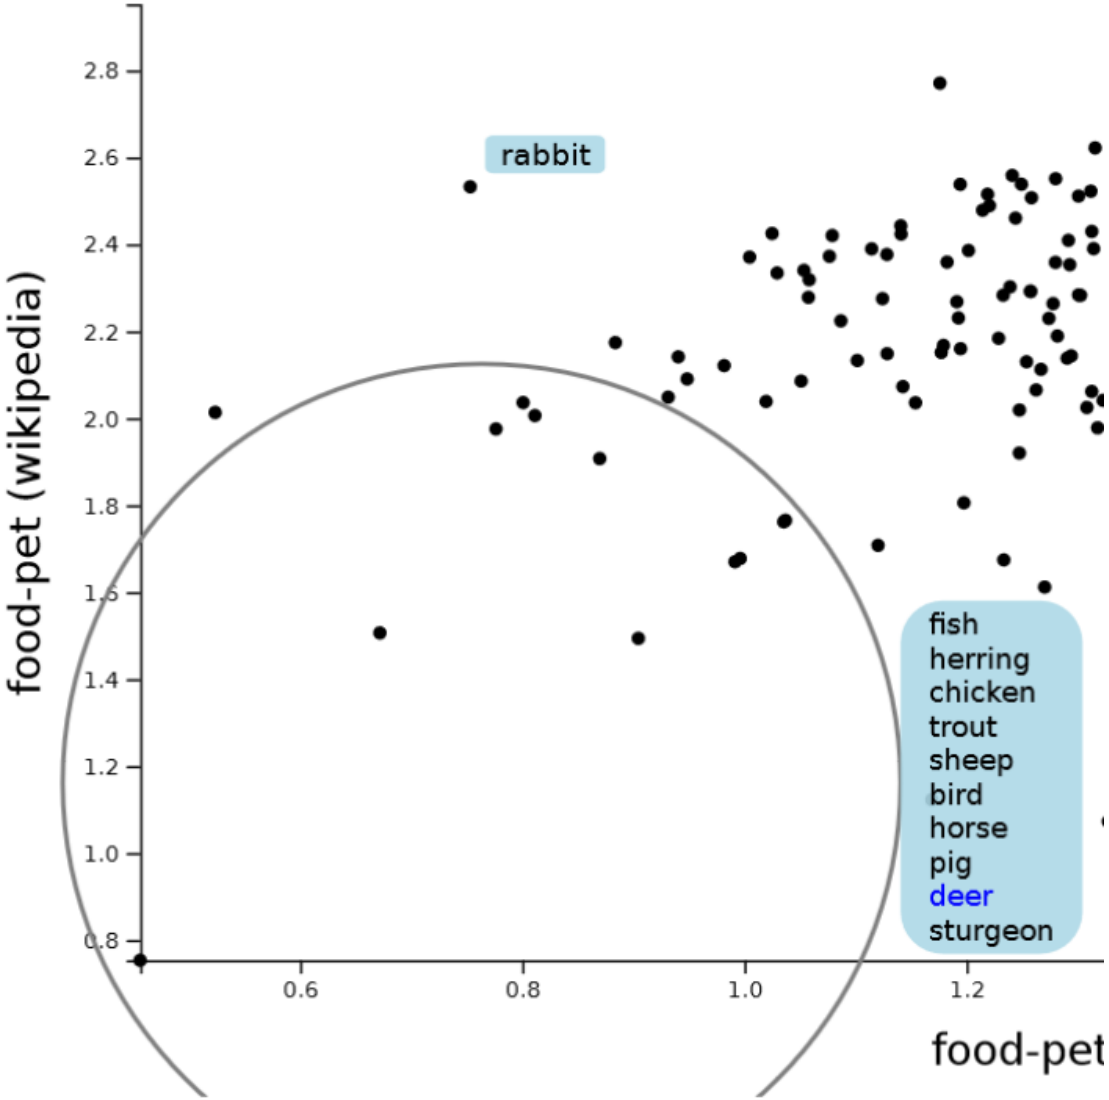

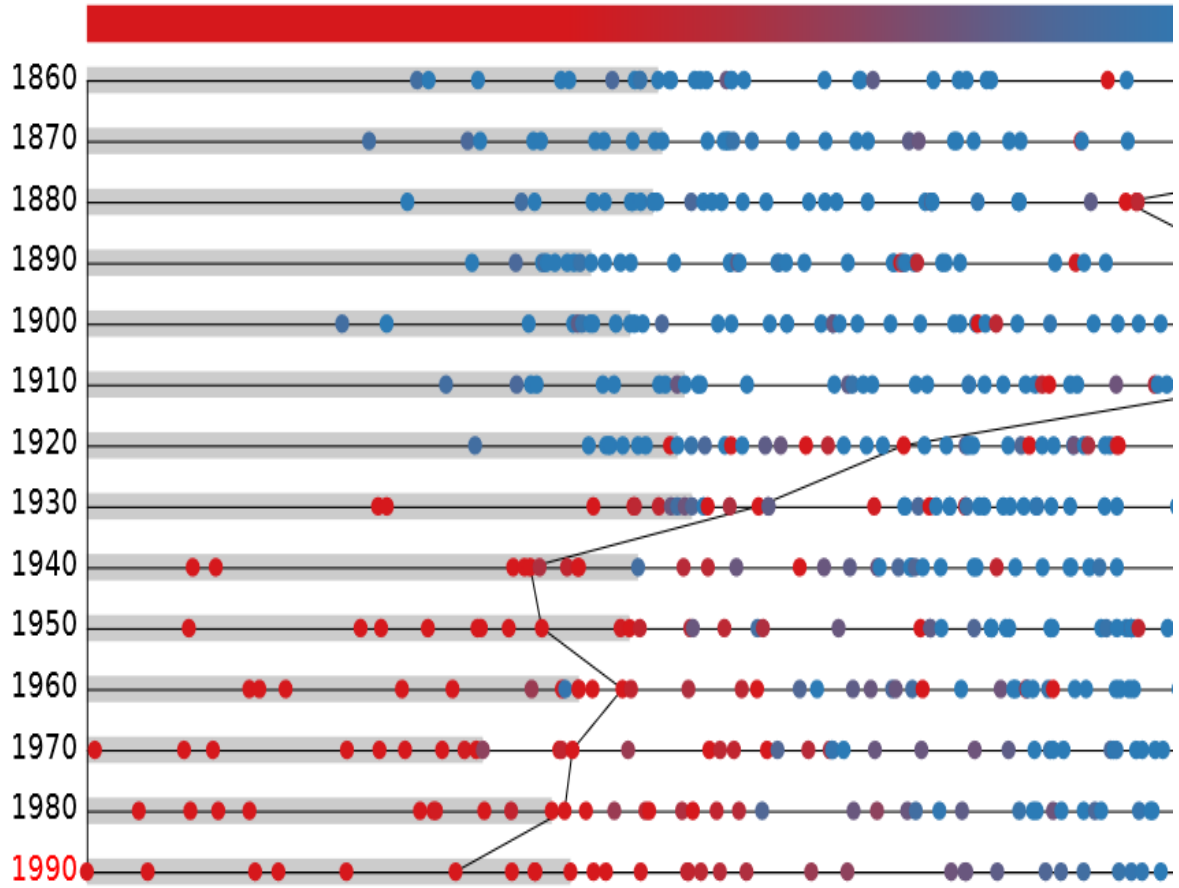

Application: Stability Assessment

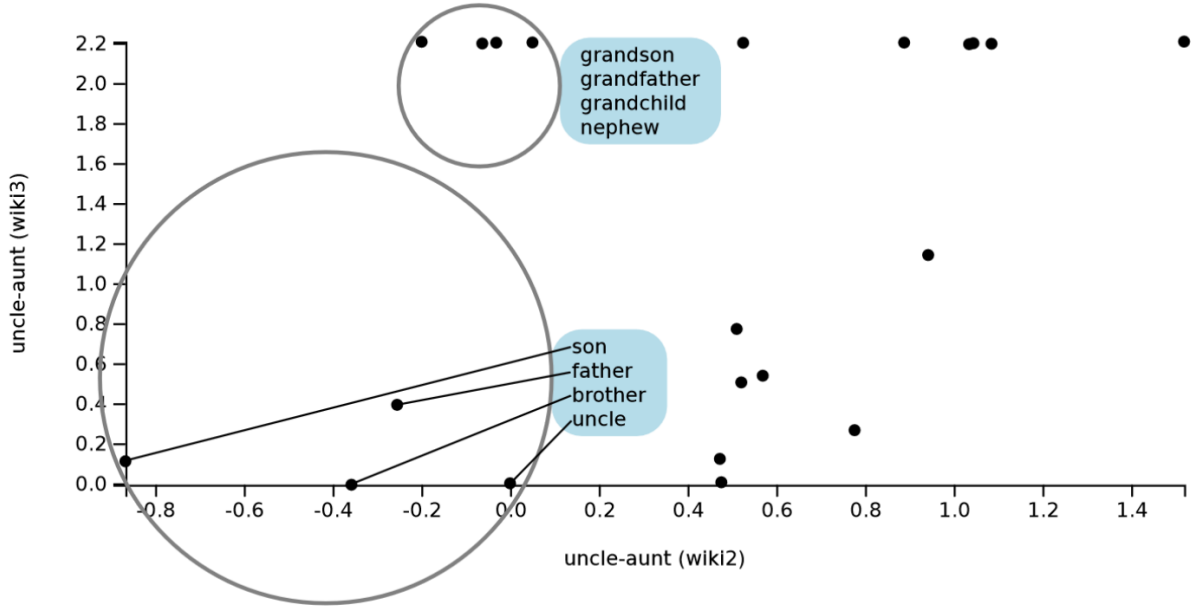

Comparison?

For interpretation

Do the differences show something?

Word meaning change

(between corpora)

Correlations / Biases

(within corpora)

For modeling (selection)

How are these models different?

Which is better?

Comparison?

Comparison is an important task in Data Analysis

Comparison is special!

(since it involves multiple things)

It's an important special case

It deserves special attention

Almost all Data Analysis can be viewed as comparison

Comparison is a lens

(to look at problems)

It's a generally useful tool

It deserves special attention

How do I think about comparison?

to help me develop tools to help people do it

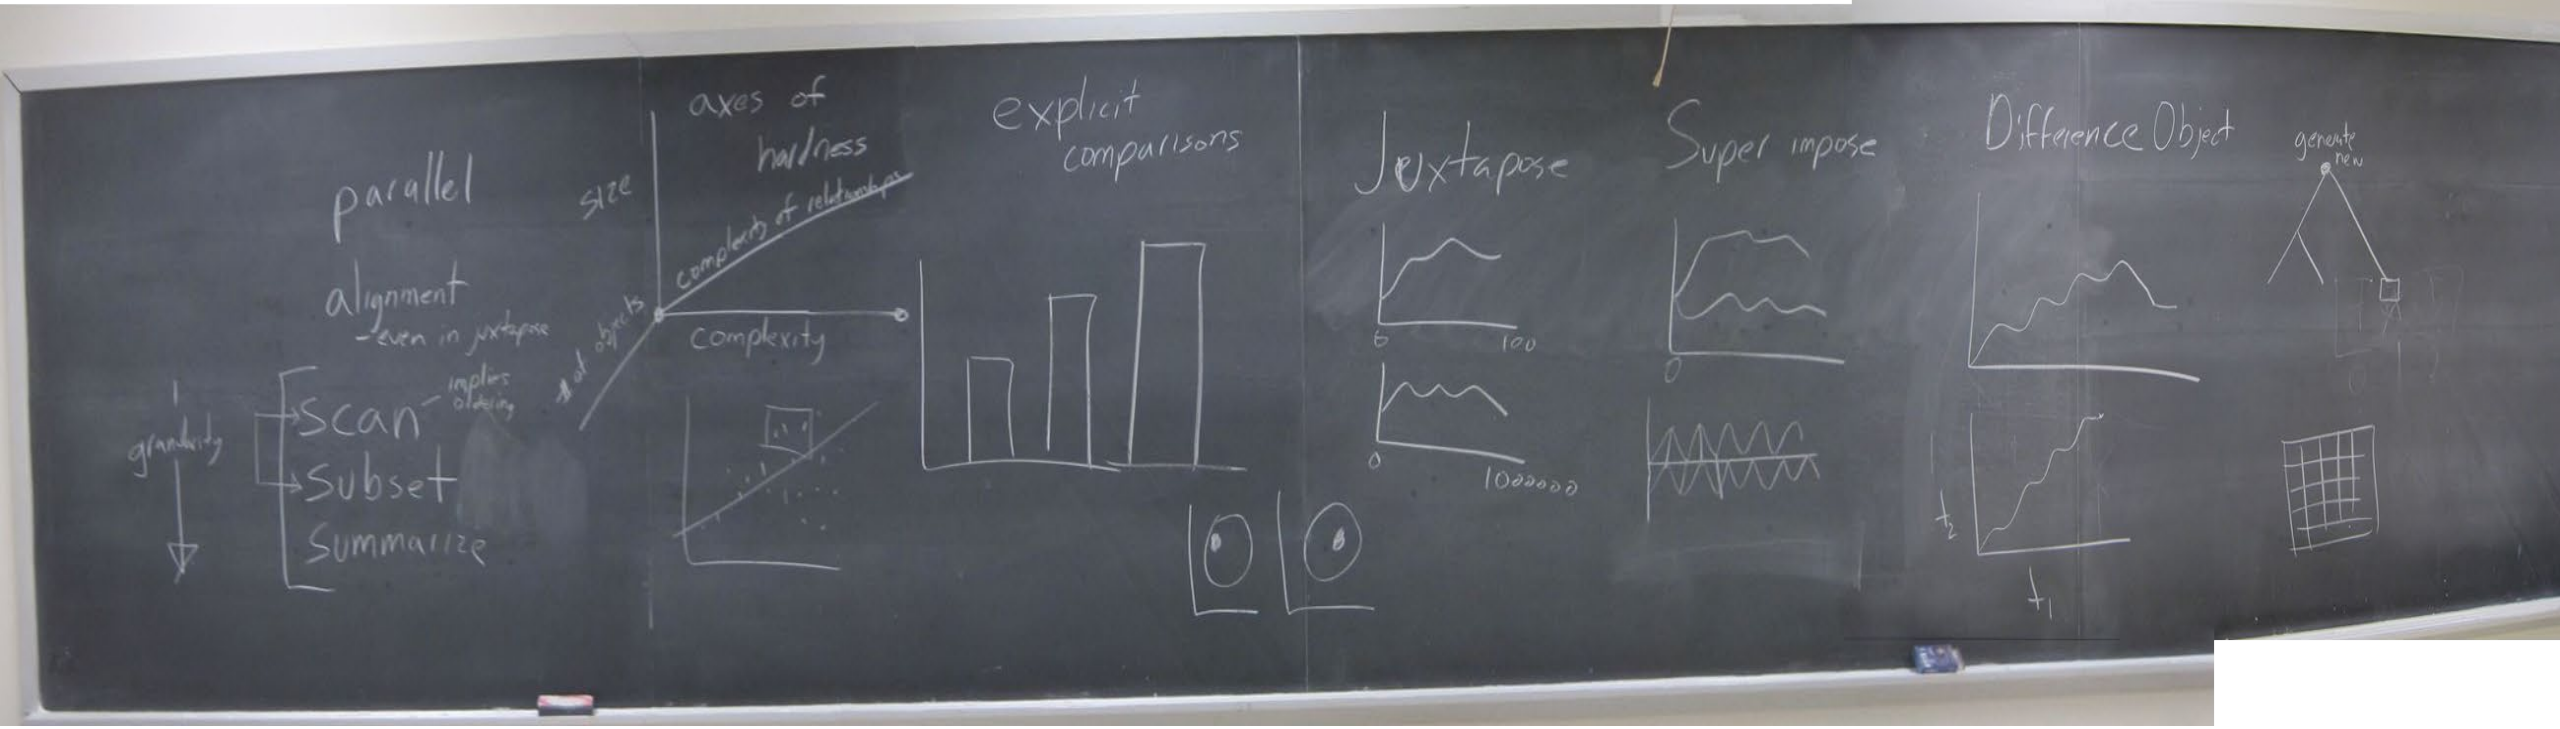

My CS838 (Data Visualization) class
March 11, 2010

Maybe I shouldn’t publish this…

It’s my secret weapon to devise new tools…

Considerations for Visualizing Comparison

Michael Gleicher

Department of Computer Sciences

University of Wisconsin – Madison

InfoVis 2017

What is this paper?

Considerations for Visualizing Comparison

4 questions to ask

when designing a visualization or tool

?

Dictionary.com definitions compare

compare

[kuh m-**pair**]

Spell Syllables

CITE

Examples Word Origin

See more synonyms on Thesaurus.com

verb (used with object), **compared, comparing.**

1. to examine (two or more objects, ideas, people, etc.) in order to note similarities and differences:
   to compare two pieces of cloth; to compare the governments of two nations.
2. to consider or describe as similar; liken: *"Shall I compare thee to a summer's day?"*.
3. *Grammar.* to form or display the degrees of comparison of (an adjective or adverb).

To examine (two or more objects, ideas, people, etc.) in order to note similarities and differences

OED Oxford English Dictionary
The definitive record of the English language

Quick search:

« Previous | Back to Results | Next » Help on Dictionary Entry | Print | Save | Email | Cite

compare, *v.*[1]

Text size: A A

View as: Outline | Full entry Quotations: Show all | Hide all Keywords: On | Off

Pronunciation: Brit. /kəmˈpɛː/, U.S. /kəmˈpɛ(ə)r/
Forms: Also ME *Sc.* **comper.**
Frequency (in current use): •••••••
Etymology: < Old French *comperer* (from 14th cent. *comparer*) = Provençal *comparar*, Spanish ... (Show More)

1.

a. *trans.* To speak of or represent as similar; to liken. Const. *to*. (With negative, in such phrases as *not to be compared to*, usually implying great inferiority in some respect.) Thesaurus »

1447 O. Bokenham *Lyvys Seyntys* (1835) 9 Seynt Margrete On to that gemme [may] weel comparyd be.
1489 (▸ a1380) J. Barbour *Bruce* (Adv.) I. 403 Off manheid and mekill mycht Till Ector dar I nane comper.
a1538 T. Starkey *Dial. Pole & Lupset* (1989) 31 The one may..be comparyd to the body & the other to the soule.
1611 *Bible* (King James) Prov. iii. 15 All the things thou canst desire, are not to be compared vnto her.
1699 W. Dampier *Voy. & Descr.* I. vii. 125 He compares it to a Sloe, in shape and taste.
1855 W. H. Prescott *Hist. Reign Philip II of Spain* I. I. iv. 113 He greatly offended the Flemings by comparing their ships to muscle-shells.
(Hide quotations)

†**b.** ***to compare***: (a thing) for one to compare, (a thing) to be compared, comparable (*to*, *with*). Thesaurus »

1484 Caxton tr. G. de la Tour-Landry *Bk. Knight of Tower* (1971) lv. 80 Suche men or wymmen be to compare to the wyf of Lothe.
1711 J. Addison *Spectator* No. 161. ¶9 An Imitation of the best Authors, is not to compare with a good Original.
(Hide quotations)

c. *intr.* To draw a com

1597 Shakespeare *Ri*

2.

a. *trans.* To mark or
or place together (act
Const. *with* (or *to*) ar

1509 A. Barclay *Bra*
?1531 J. Frith *Disput*
a1640 R. Burton *An*
1667 Milton *Paradis*
1710 R. Steele *Tatler*
1850 R. W. Emerson *Mo*
in any other [country].
1860 J. Tyndall *Glaciers of Alps* II. x. 283 To compare the motion of the eastern and western halves of the glacier.
1879 G. C. Harlan *Eyesight* viii. 106 This cramping tendency of town as compared to country.

This entry has not yet been fully updated (first published 1891).

Entry history
Entry profile

Previous version: OED2 (1989)

In this entry:
compare notes, to
compare, to

In other dictionaries:
Oxford Dictionaries
compare: view definition in Oxford Dictionaries
comparen, v. in Middle English Dictionary

To mark or point out the similarities and differences of (two or more things)

What is the comparison?

Why is it hard?

How to address the challenges?

Which visual design to use?

What is the comparison?

Comparative Elements
Targets
Actions

Why is it hard?

Comparative Challenges
Number of Targets
Large or Complex Targets
Complex Relationships

How to address the challenges?

Scalability Strategies
Scan Sequentially
Select Subset
Summarize Somehow

Which visual design to use?

Comparative Designs
Juxtapose
Superpose
Explicit Encoding

Wrong Question:
Is my problem Comparison?

Just about anything can be viewed as comparison

Not everything benefits from being viewed this way

Serendip: Topic Model-Driven Visual Exploration of Text Corpora

Eric Alexander, Joe Kohlmann, Robin Valenza, Michael Witmore, and Michael Gleicher, *Member, IEEE*

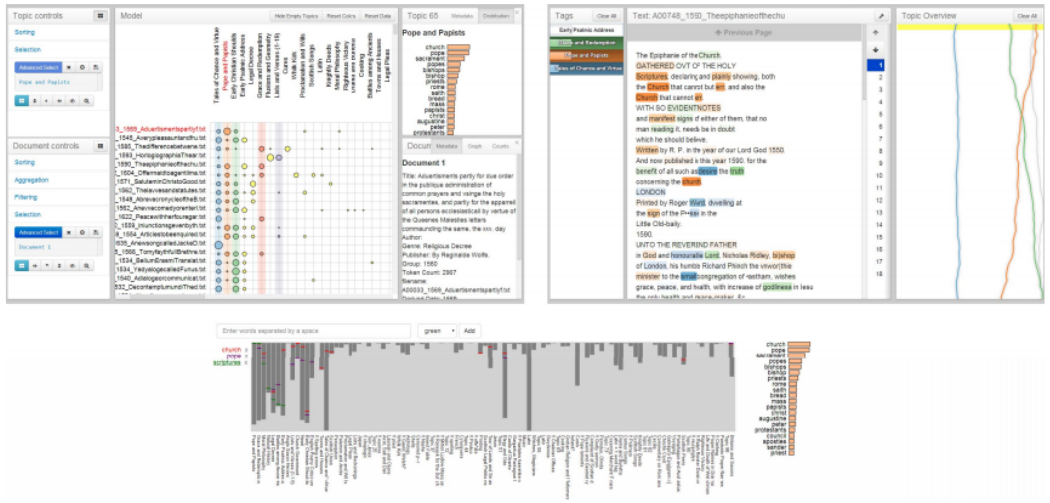

Fig. 1. The three main views of Serendip: CorpusViewer, TextViewer, and RankViewer.

Abstract— Exploration and discovery in a large text corpus requires investigation at multiple levels of abstraction, from a zoomed-out view of the entire corpus down to close-ups of individual passages and words. At each of these levels, there is a wealth of information that can inform inquiry—from statistical models, to metadata, to the researcher's own knowledge and expertise. Joining all this information together can be a challenge, and there are issues of scale to be combatted along the way. In this paper, we describe an approach to text analysis that addresses these challenges of scale and multiple information sources, using probabilistic topic models to structure exploration through multiple levels of inquiry in a way that fosters serendipitous discovery. In implementing this approach into a tool called Serendip, we incorporate topic model data and metadata into a highly reorderable matrix to expose corpus level trends; extend encodings of tagged text to illustrate probabilistic information at a passage level; and introduce a technique for visualizing individual word rankings, along with interaction techniques and new statistical methods to create links between different levels and information types. We describe example uses from both the humanities and visualization research that illustrate the benefits of our approach.

Index Terms—Text visualization, topic modeling.

1 Introduction

Exploration and discovery in large text corpora can be a daunting task. Corpora can easily grow to thousands or more texts, ranging in length from short snippets to long books. The task is further complicated by the range of questions that can be asked of such corpora, broad both in subject (making comparisons across time, genre, author, etc.) and in level of detail (corpus, document, passage, even word). Discoveries must often connect multiple subjects and levels of inquiry. Fortunately, there is considerable information to aid these inquiries. Beyond the texts themselves, there are statistical summaries of content, document metadata, and analysts' explicit and implicit knowledge of the documents and their context. However, mixing these different types of information across scales of inquiry is challenging. The information types, and the existing tools that support their use, generally focus solely on a particular scale.

In this paper, we introduce a topic modeling tool for text exploration that is designed to address the issues of inter-mixing scales of inquiry and information types. Our core idea is that to enable fluent fusion, a system must provide not only a set of views for looking at the data from multiple viewpoints, but also connections between the different types of information allowing a reader to move smoothly across scales, data types, and research questions. To achieve this, we have had to adapt existing views to work with different types of text corpora data, develop new views that address some unmet needs, and introduce statistical methods that help connect between different object types. The resulting system enables users to explore questions about collections

- *E. Alexander, J. Kohlmann, and M. Gleicher are with the Department of Computer Sciences at the University of Wisconsin-Madison. E-mail: ealexand@cs.wisc.edu, kohlmannj@gmail.com, gleicher@cs.wisc.edu.*
- *R. Valenza is with the Department of English at the University of Wisconsin-Madison. Email: valenza@wisc.edu.*
- *M. Witmore is with the Folger Shakespeare Library in Washington, D.C. Email: mwitmore@folger.edu.*

IEEE Symposium on Visual Analytics Science and Technology 2014
November 9-14, Paris, France
978-1-4799-6227-3/14/$31.00 ©2014 IEEE

Serendip, VAST '14

Alexander, E., Kohlmann, J., Valenza, R., Witmore, M., & Gleicher, M. (2014). Serendip: Topic model-driven visual exploration of text corpora. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*

Serendip: Topic model-driven visual exploration of text corpora

Use a **topic model** to guide exploration of a **text corpus**

Find patterns and connect back to specifics

Documents

Model

Documents/Passages

Goals

Support inquiry across levels of abstraction

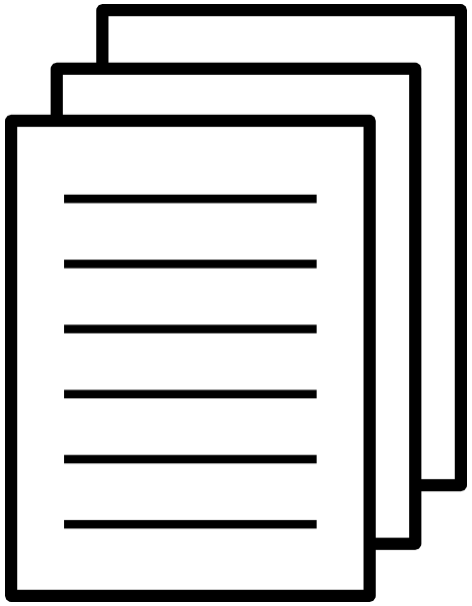

Corpus

Document

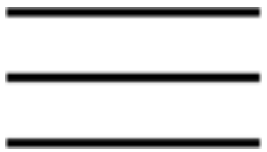

Passage

Word

Goals

Support inquiry across levels of abstraction

Combat issues of scale in the data

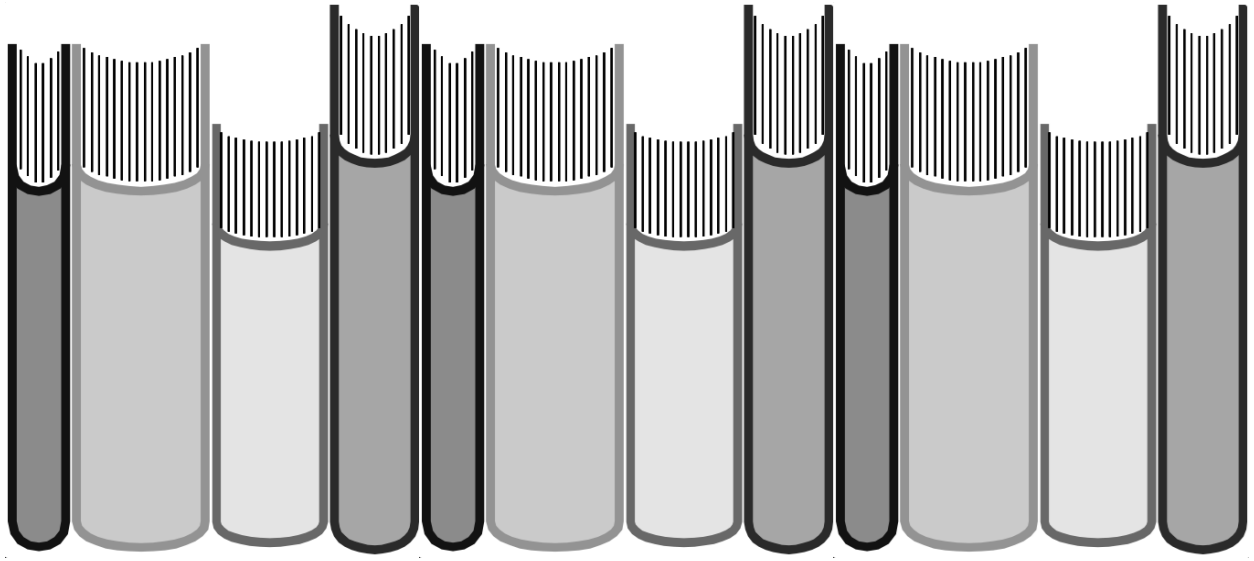

Many documents

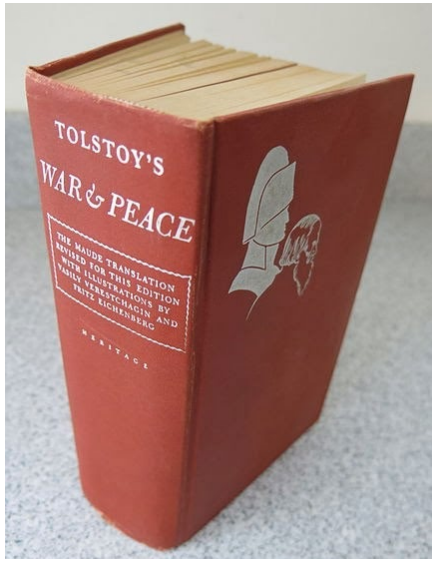

Long documents

Goals

Support inquiry across levels of abstraction

Combat issues of scale

Promote serendipitous discovery:

- Multiple entry points
- Highlight adjacency
- Flexible exploration

A. Thudt, U. Hinrichs, S. Carpendale 2012

Serendip

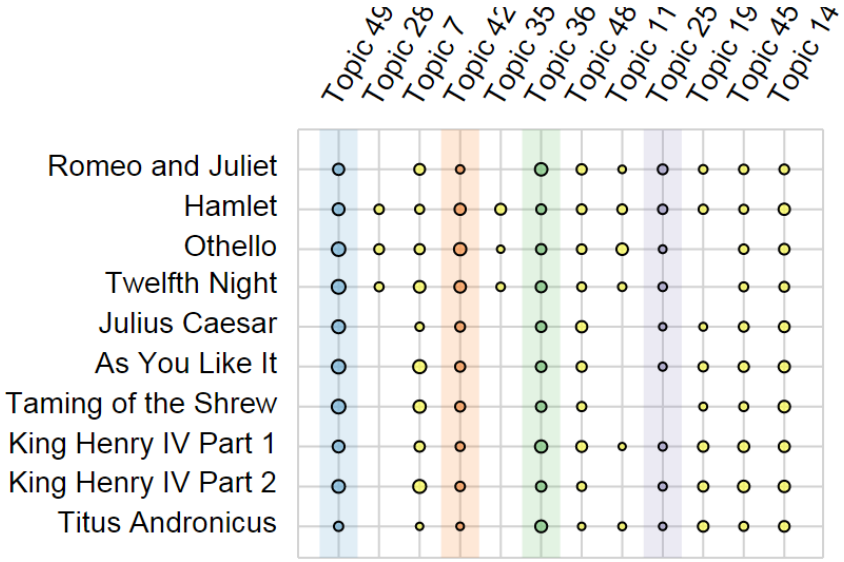

CorpusViewer

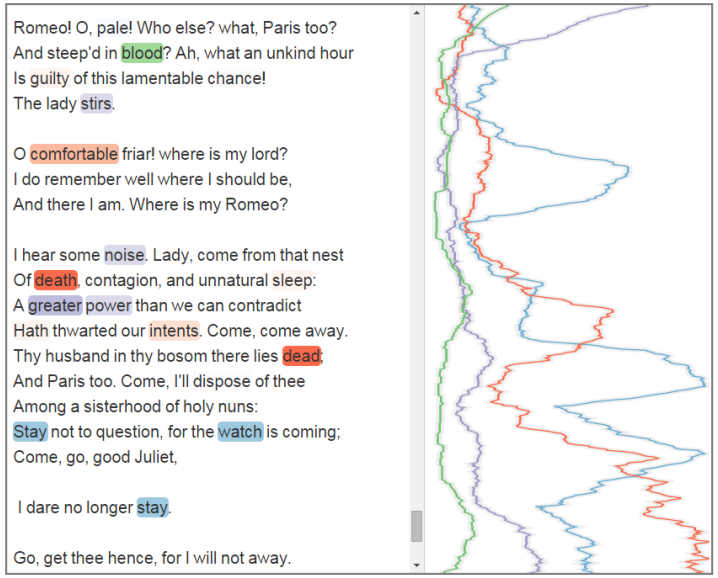

TextViewer

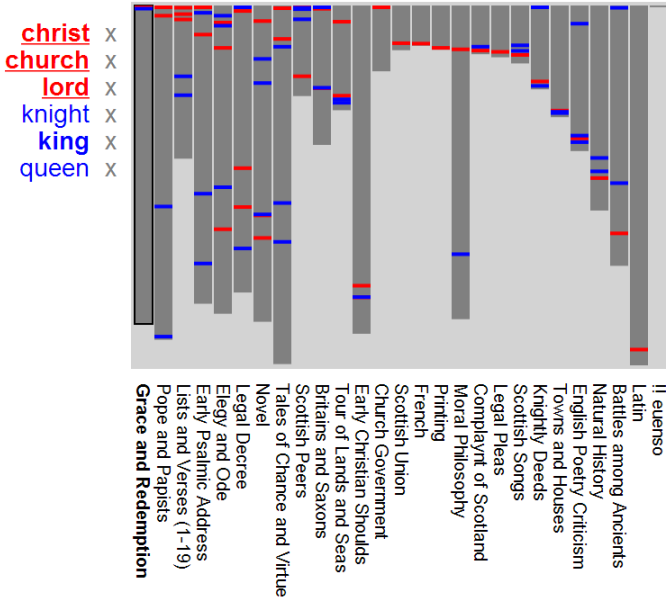

RankViewer

Basics of Serendip

Three different (interlinked views)

CorpusViewer

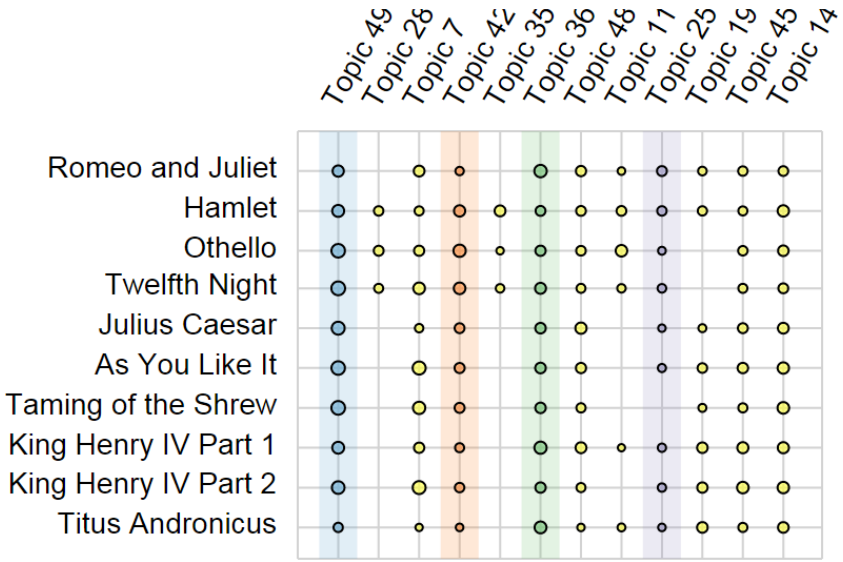

TextViewer

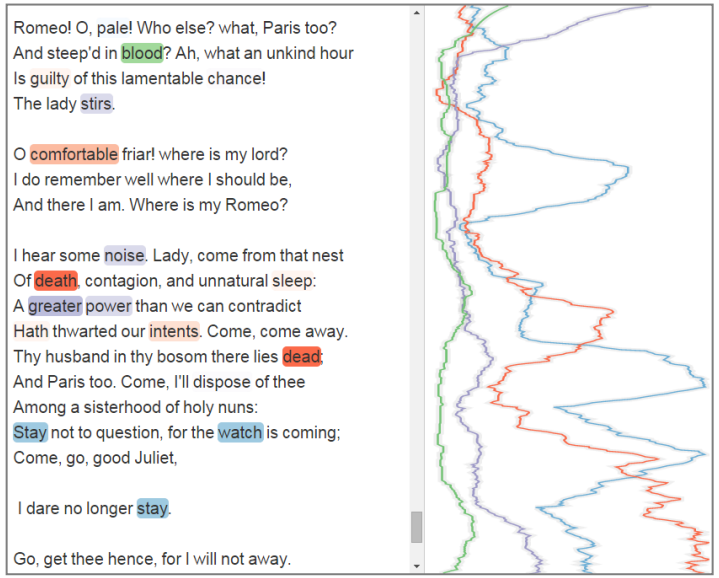

RankViewer

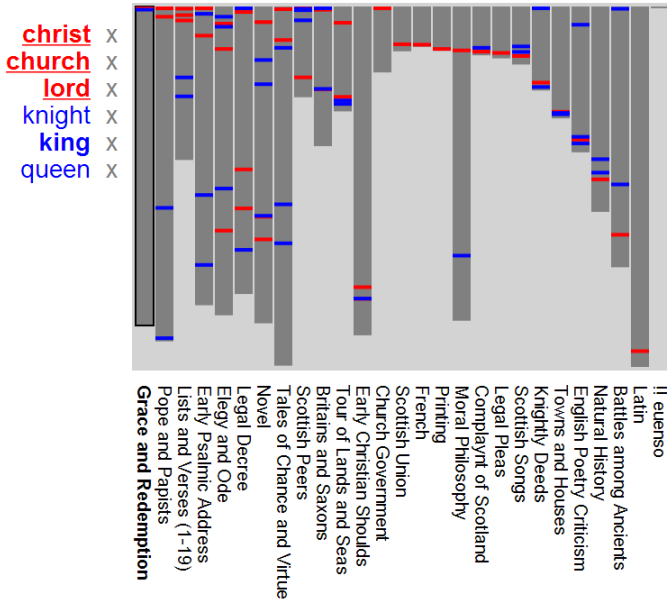

Corpus **Document** **Passage** **Word**


Topic controls
Sorting
Selection
Advanced Select
Pope and Papists
Document controls
Sorting
Aggregation
Filtering
Selection
Advanced Select
Document 1
Model
Hide Empty Topics
Reset Colors
Reset Data
Topic 65
Pope and Papists
church
pope
sacrament
popes
bishops
bishop
priests
rome
faith
bread
mass
papists
christ
augustine
peter
protestants
Document 1
Author:
Genre: Religious Decree
Publisher: By Reginalde Wolfe
Group: 1560
Token Count: 2367

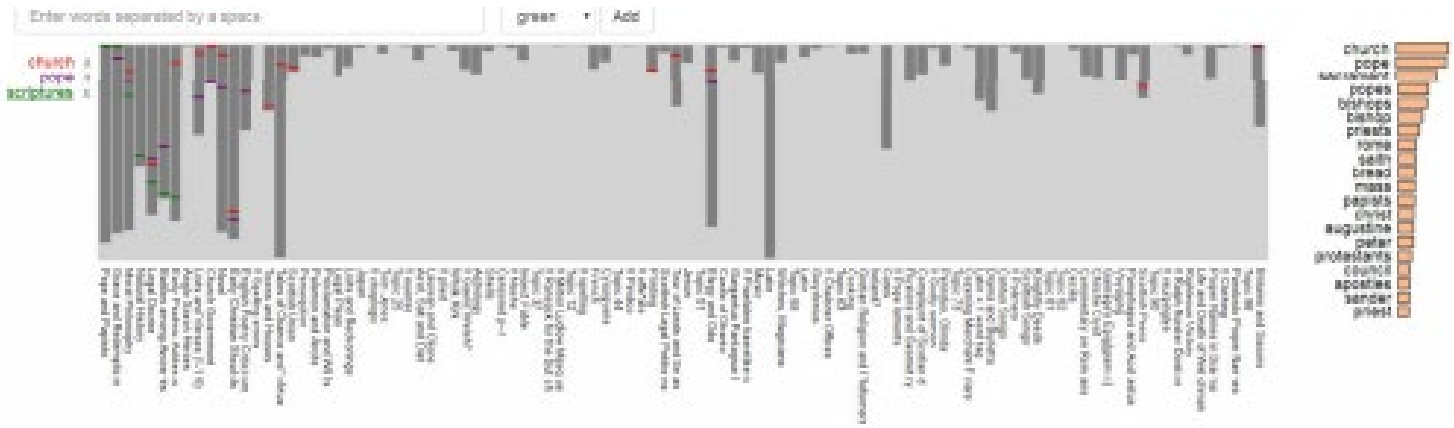
Enter words separated by a space
green
Add
church
pope
scriptures

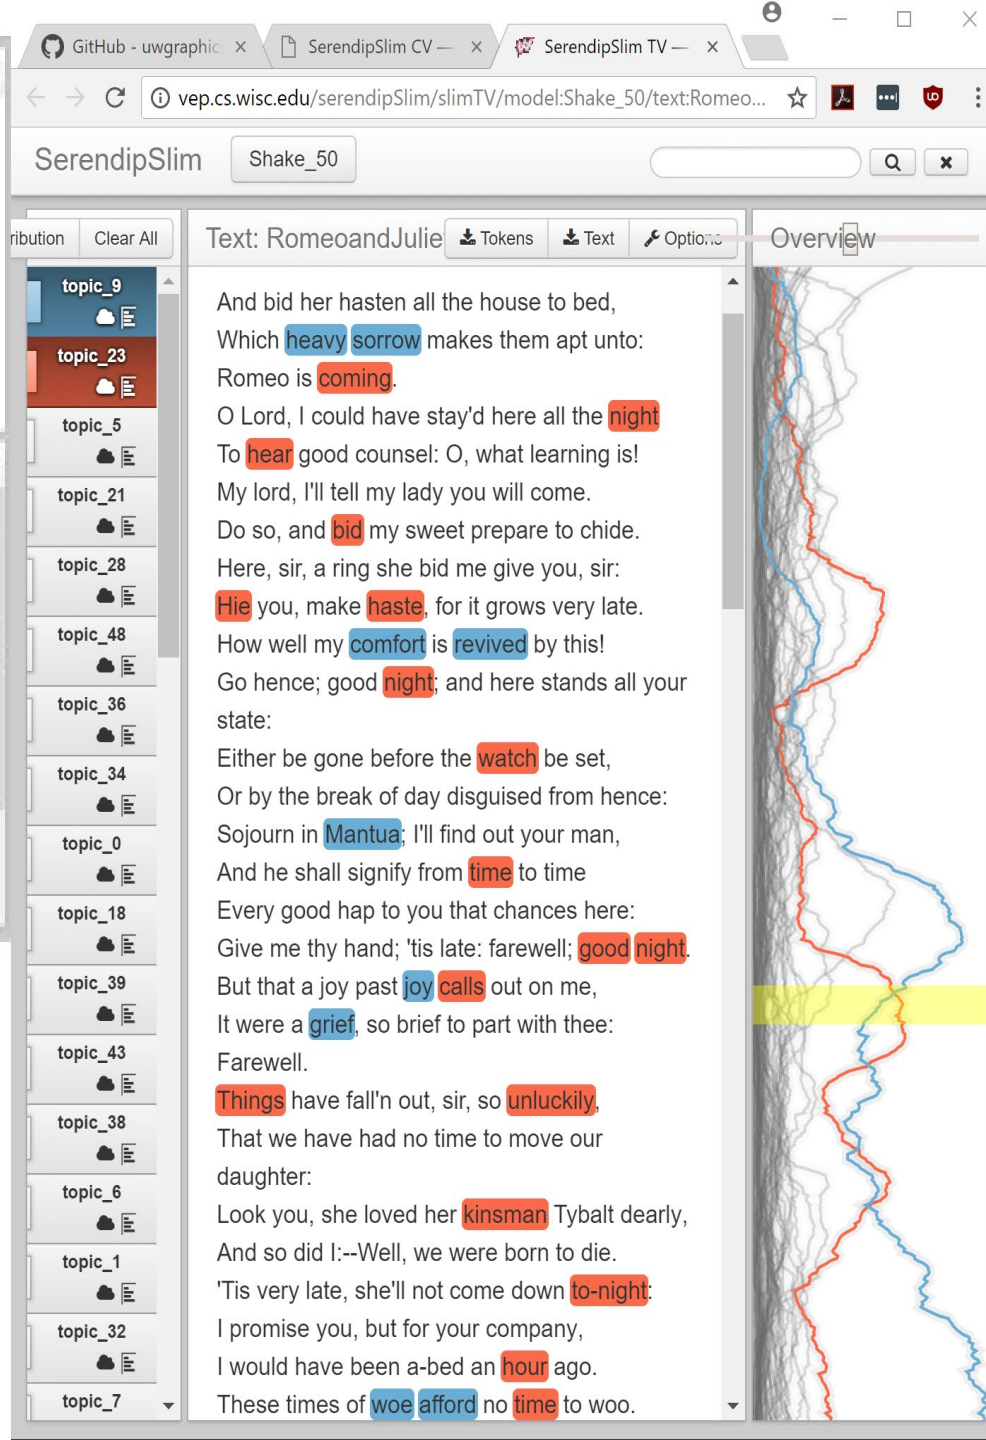
GitHub - uwgraphic
SerendipSlim CV —
SerendipSlim TV —
vep.cs.wisc.edu/serendipSlim/slimTV/model:Shake_50/text:Romeo...
SerendipSlim
Shake_50
ribution
Clear All
topic_9
topic_23
topic_5
topic_21
topic_28
topic_48
topic_36
topic_34
topic_0
topic_18
topic_39
topic_43
topic_38
topic_6
topic_1
topic_32
topic_7
Text: RomeoandJulie
Tokens
Text
Options
Overview
And bid her hasten all the house to bed,
Which heavy sorrow makes them apt unto:
Romeo is coming.
O Lord, I could have stay'd here all the night
To hear good counsel: O, what learning is!
My lord, I'll tell my lady you will come.
Do so, and bid my sweet prepare to chide.
Here, sir, a ring she bid me give you, sir:
Hie you, make haste, for it grows very late.
How well my comfort is revived by this!
Go hence; good night; and here stands all your
state:
Either be gone before the watch be set,
Or by the break of day disguised from hence:
Sojourn in Mantua; I'll find out your man,
And he shall signify from time to time
Every good hap to you that chances here:
Give me thy hand; 'tis late: farewell; good night.
But that a joy past joy calls out on me,
It were a grief, so brief to part with thee:
Farewell.
Things have fall'n out, sir, so unluckily,
That we have had no time to move our
daughter:
Look you, she loved her kinsman Tybalt dearly,
And so did I:--Well, we were born to die.
'Tis very late, she'll not come down to-night:
I promise you, but for your company,
I would have been a-bed an hour ago.
These times of woe afford no time to woo.

Is this comparison?

No!

A tool for exploring a topic model!

We didn't describe it as comparison

Tool for looking at one topic model

Unclear how users think about it

Yes!

Comparison thinking really helped!

We did think about comparison

Tool for using topic models

Our users had comparison tasks

~~Is this comparison?~~ I don't care!

No!

A tool for exploring a topic model!

A survey of comparison would have missed this.

parison

model

Unclear how users think about it

Yes!

Comparison thinking really helped!

We did

Tool for

It's a great example of comparison ideas

Our users had comparison tasks

What is the comparison?

Comparative Elements
Targets
Actions

Why is it hard?

Comparative Challenges
Number of Targets
Large or Complex Targets
Complex Relationships

How to address the challenges?

Scalability Strategies
Scan Sequentially
Select Subset
Summarize Somehow

Which visual design to use?

Comparative Designs
Juxtapose
Superpose
Explicit Encoding

Question 1:

~~What is the comparison?~~

What are the elements of the comparison?

The Elements of a Comparison . . .

To **examine**
(**two or more objects, ideas, etc.**)
in order to note
similarities and differences

To **mark** or **point** out the
similarities and differences of
(**two or more things**)

Targets — Set of things being compared

Action — What to do with the **relationship** among them

Question 1A:
The Elements: Targets

Do you know what you are comparing?

Explicit Comparisons – the system has the set of targets

Implicit Comparisons – the system may not know all the targets

compare against an implicit baseline

compare against the user's knowledge

compare with targets only the user knows

Question 1A:
The Elements: Targets

What is being compared? – Comparison Targets

Does the **model** match my expectations?

What **documents** are similar?

How do **groups of documents** differ?

What **words** indicate these differences?

How are **words** used differently?

Where in **texts** are these differences?

Do the **patterns** match other things I know?

Question 1B:
The Elements: Actions

Verbs on relationships

Try to be more specific than “examine” or “compare”

Truth in Advertising: I didn’t have this worked out in 2010

Question 1B:
The Elements: Actions

What to do with the relationship? Comparison Actions (Verbs)

Does the **model** match my expectations?	**Measure/Quantify** relationship
What **documents** are similar?	**Identify** similar things
How do **groups of documents** differ?	**Measure/Quantify** relationship
What **words** indicate these differences?	**Dissect** a difference
How are **words** used differently?	**Identify** meaningful differences
Where in **texts** are these differences?	**Contextualize** the relationships
Do the **patterns** match other things I know?	**Identify** similar things

What is the comparison?

Comparative Elements
Targets
Actions

Why is it hard?

Comparative Challenges
Number of Targets
Large or Complex Targets
Complex Relationships

How to address the challenges?

Scalability Strategies
Scan Sequentially
Select Subset
Summarize Somehow

Which visual design to use?

Comparative Designs
Juxtapose
Superpose
Explicit Encoding

Question 2:
Why is this comparison hard?

If it isn't hard, you probably don't need to think about it (much)

Abstractly

Too many targets to compare

Large or Complex Targets

Complex Relationships

Serendip

Lots of documents

Long / complex documents

Complicated models

Challenges of Scale!

Only Scalability Challenges?

Many different comparisons

Challenges from the kind (not scale)

- Hard target types (implicit)
- Hard action types (dissection)
- Hard combinations (dissect implicit)

Tasks influence scalability challenges

Solutions must respond to both!

Serendip Comparison Example: Where does this happen? (contextualize)

Task Challenge:

Contextualize – fit user knowledge

Strategy: show in context

Design: use text as scaffold

Scalability Challenge:

Long Documents

Strategy: summarize

Design: overview + detail

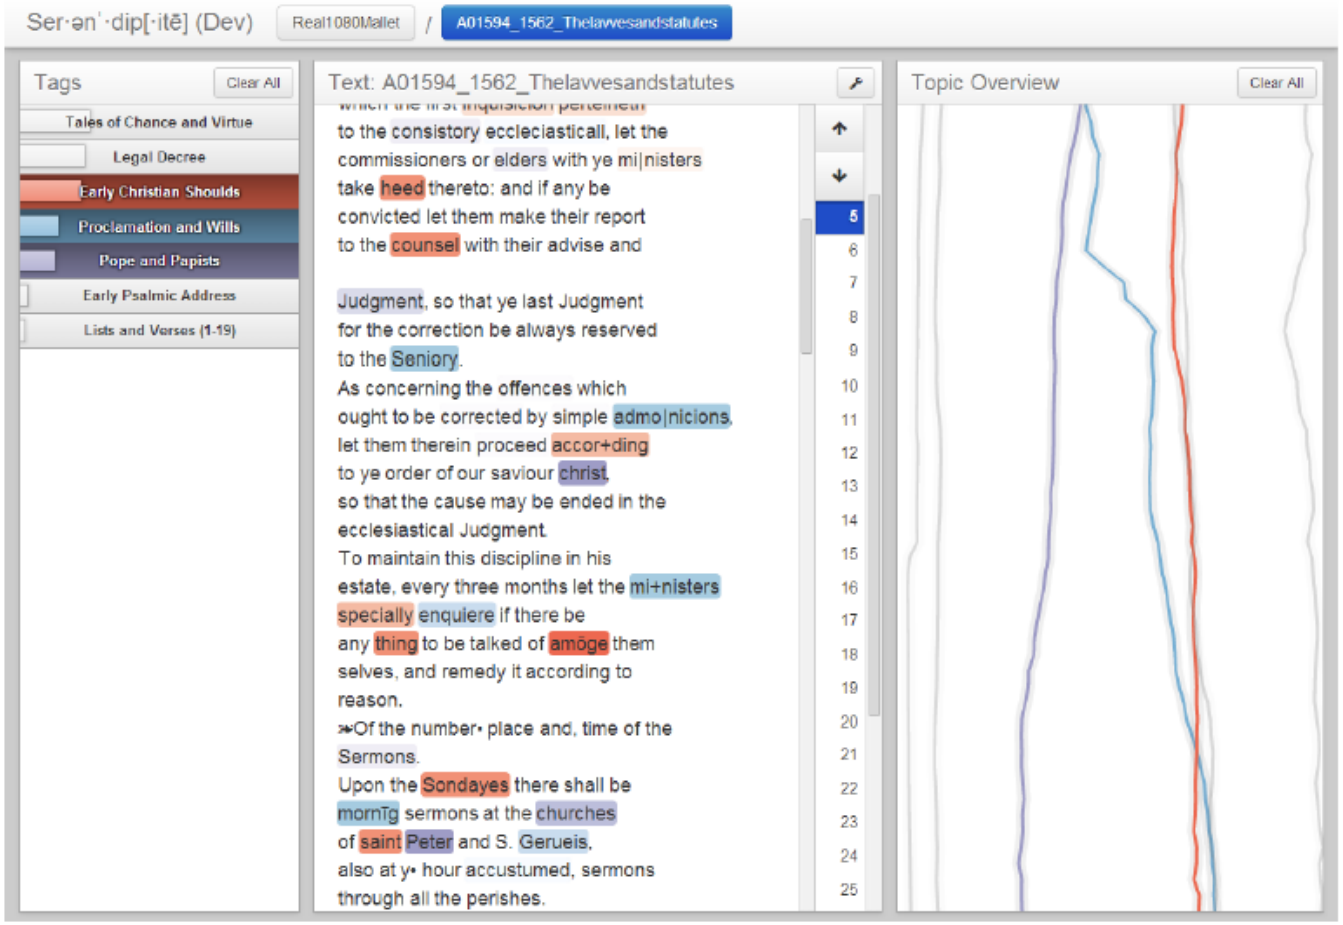

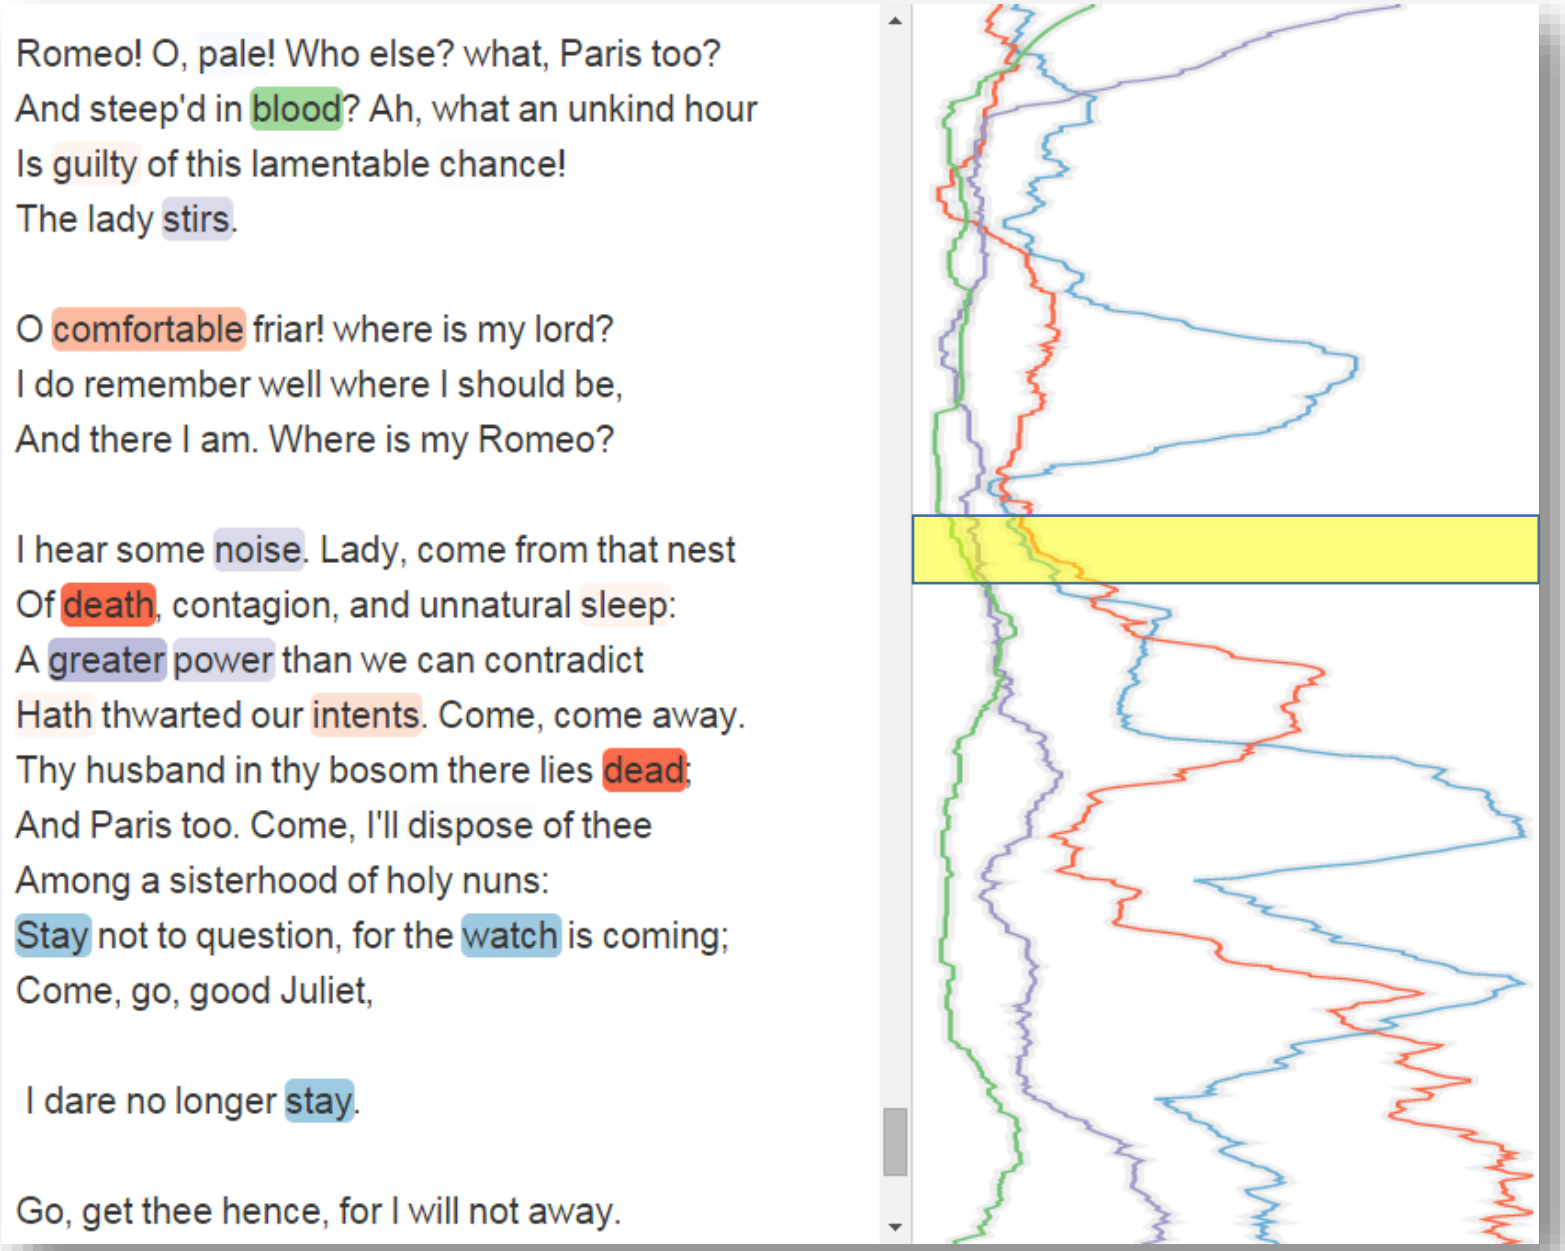

Document

Line-graph Overview

Combating scale

Showing trends

Affording navigation

What is the comparison?

Comparative Elements
Targets
Actions

Why is it hard?

Comparative Challenges
Number of Targets
Large or Complex Targets
Complex Relationships

How to address the challenges?

Scalability Strategies
Scan Sequentially
Select Subset
Summarize Somehow

Which visual design to use?

Comparative Designs
Juxtapose
Superpose
Explicit Encoding

Question 3:
What is your strategy for those challenges?

Abstractly

Scan Sequentially

Select Subset

Summarize Somehow

Scalability Strategies!

Sarikaya, Gleicher & Szafir. (2018). Design Factors for Summary Visualization in Visual Analytics. *Computer Graphics Forum*, *37*(3), 145–156.

EuroVis 2018.

Serendip Comparison Example: Compare Groups

Task Challenge:

Implicit targets – what groups?

Strategy: make explicit

Design: user specifies groups

Scalability Challenge:

Lots of documents

Strategy: summarize

Design: how to present statistics?

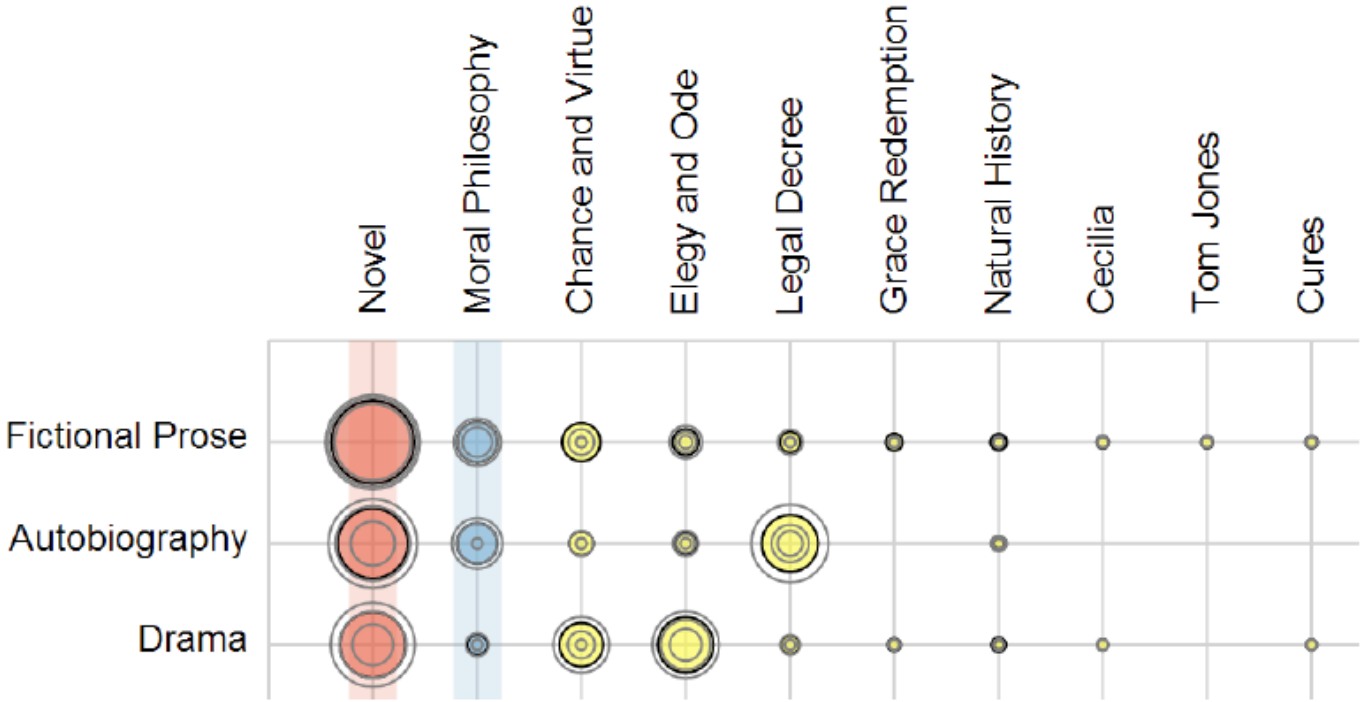

What is the comparison?

Comparative Elements
Targets
Actions

Why is it hard?

Comparative Challenges
Number of Targets
Large or Complex Targets
Complex Relationships

How to address the challenges?

Scalability Strategies
Scan Sequentially
Select Subset
Summarize Somehow

Which visual design to use?

Comparative Designs
Juxtapose
Superpose
Explicit Encoding

Question 4:
What Visual Design for Comparison?

Abstractly

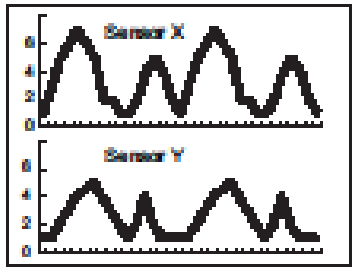

Juxtaposition

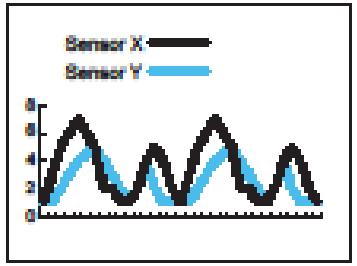

Superposition

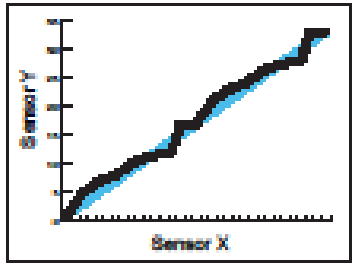

Explicit Encoding

Gleicher, M., Albers, D., Walker, R., Jusufi, I., Hansen, C. D., & Roberts, J. C. (2011). Visual comparison for information visualization. *Information Visualization*, *10*(4), 289–309.

What is the comparison?

Comparative Elements
Targets
Actions

Why is it hard?

Comparative Challenges
Number of Targets
Large or Complex Targets
Complex Relationships

How to address the challenges?

Scalability Strategies
Scan Sequentially
Select Subset
Summarize Somehow

Which visual design to use?

Comparative Designs
Juxtapose
Superpose
Explicit Encoding

Should I think about comparison?

Maybe… If it helps

If you're so good at comparison...

We should be able to compare complex things

Task-Driven Comparison of Topic Models

Eric Alexander and Michael Gleicher

Department of Computer Sciences

University of Wisconsin – Madison

VAST 2015

Why compare models?

Algorithm builders: “How does my technique compare?”

Many models to compare: LDA, NNMF, LSA, etc.

Model builders: “Which model is the best?”

Model selection in huge parameter space w/ no ground truth

Model Users: “Does this mean what I think it means?”

Validate findings across multiple runs
Understand and build trust in the algorithm(s)

Model Users: Compare corpora

Are there interesting similarities and differences?
See through algorithmic differences

Our approach

Single model tasks → Model comparison tasks → **Visual techniques**

Understanding topics: Topic alignment

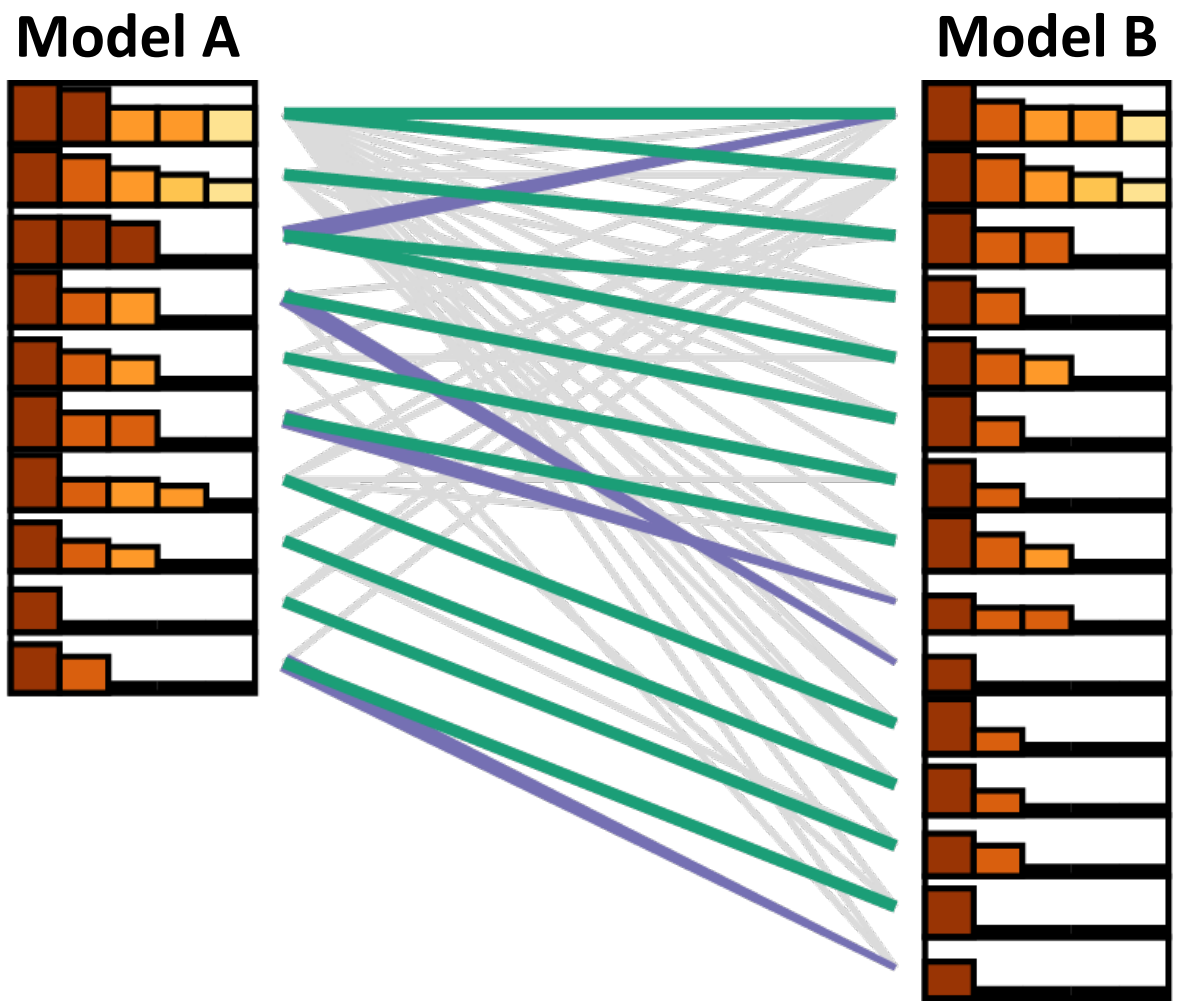

Understanding topics: Topic similarity

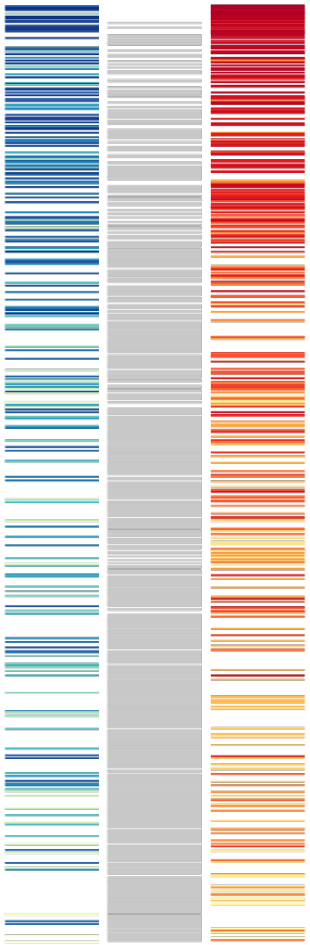

Understanding document similarity

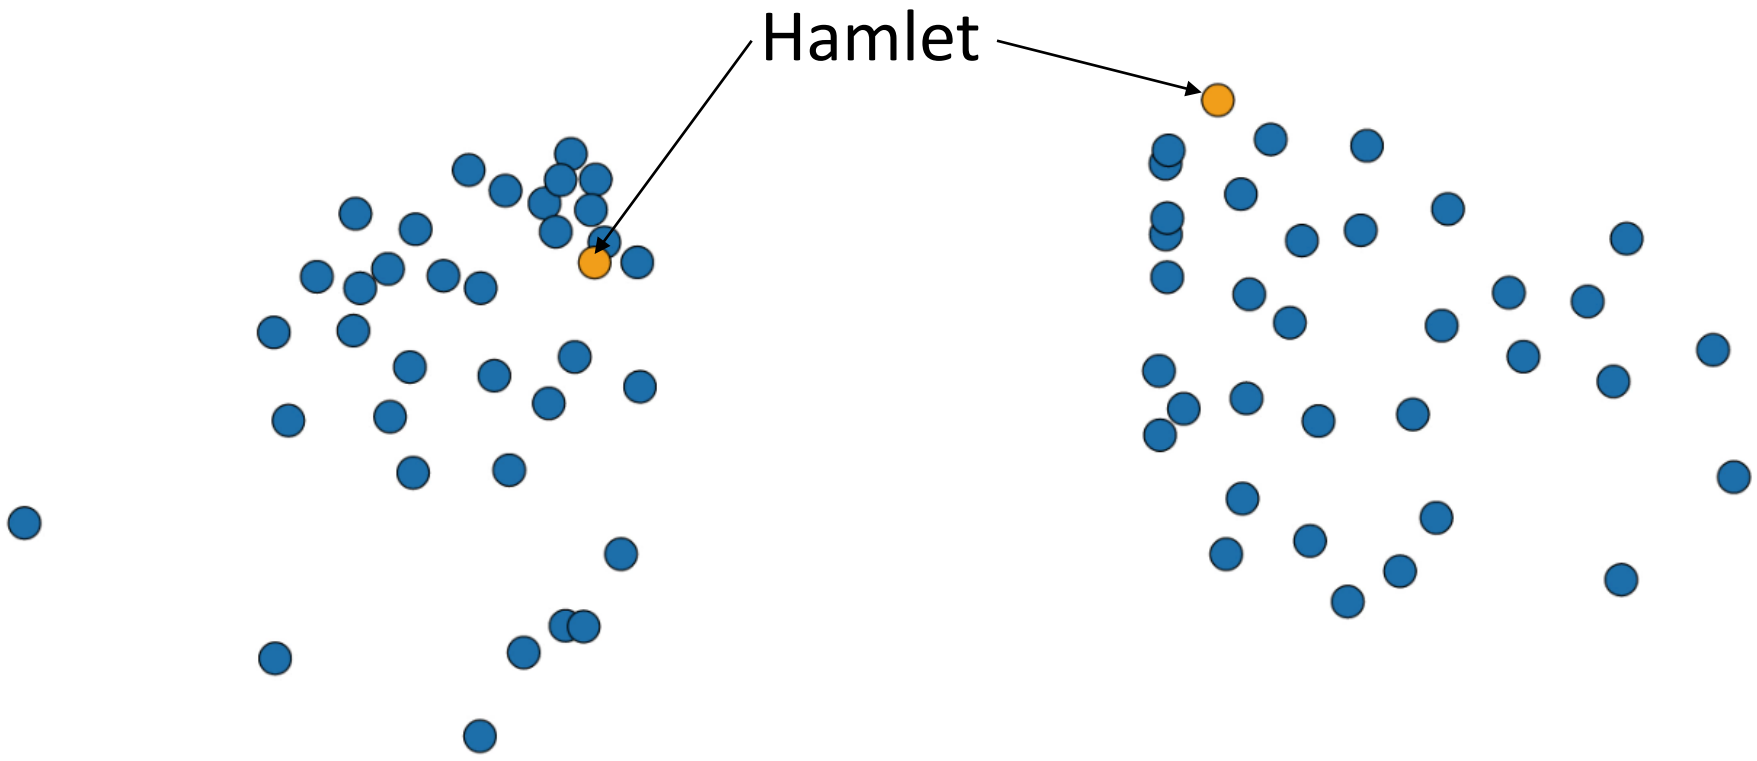

MDS on distances
from Model A

MDS on distances
from Model B

Understanding document similarity:
Buddy Plots

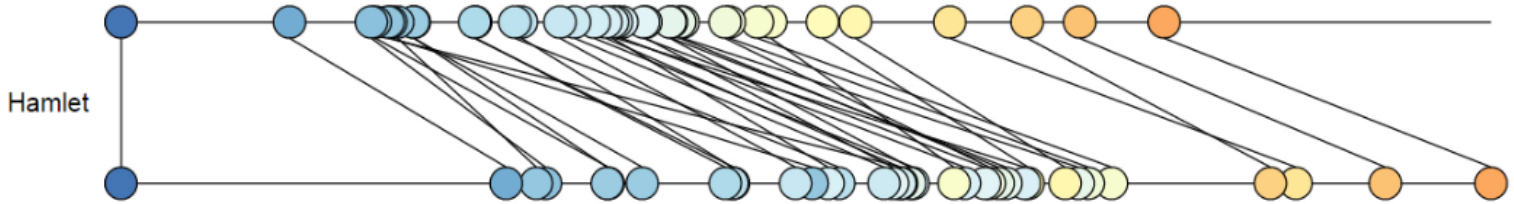

Understanding document similarity:

Cluster comparison (Visualizing 20 runs of k-means, k=5)

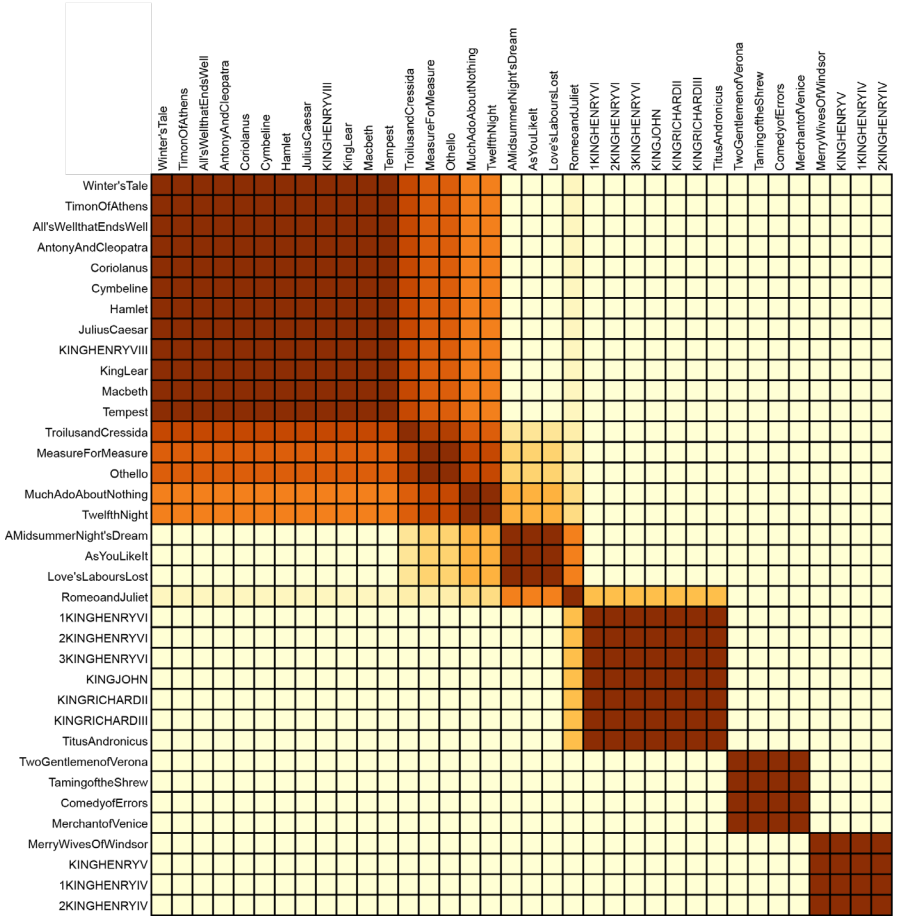

10 topics

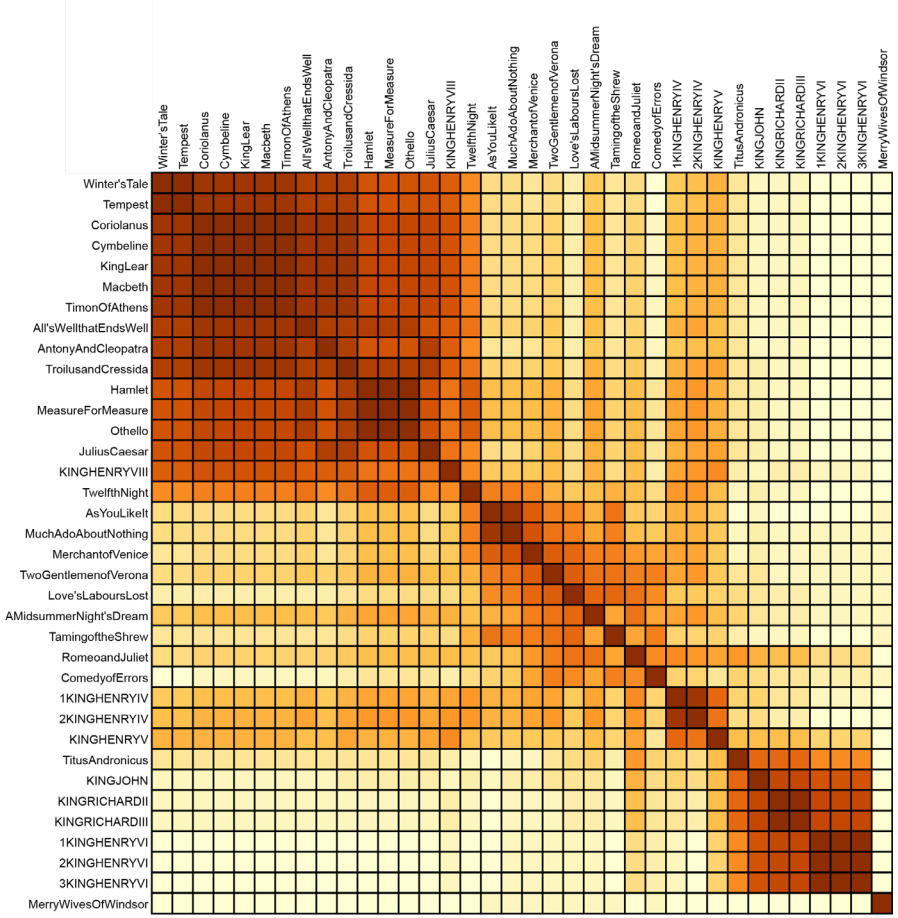

50 topics

Understanding change:
Asymmetrical Flow Diagrams

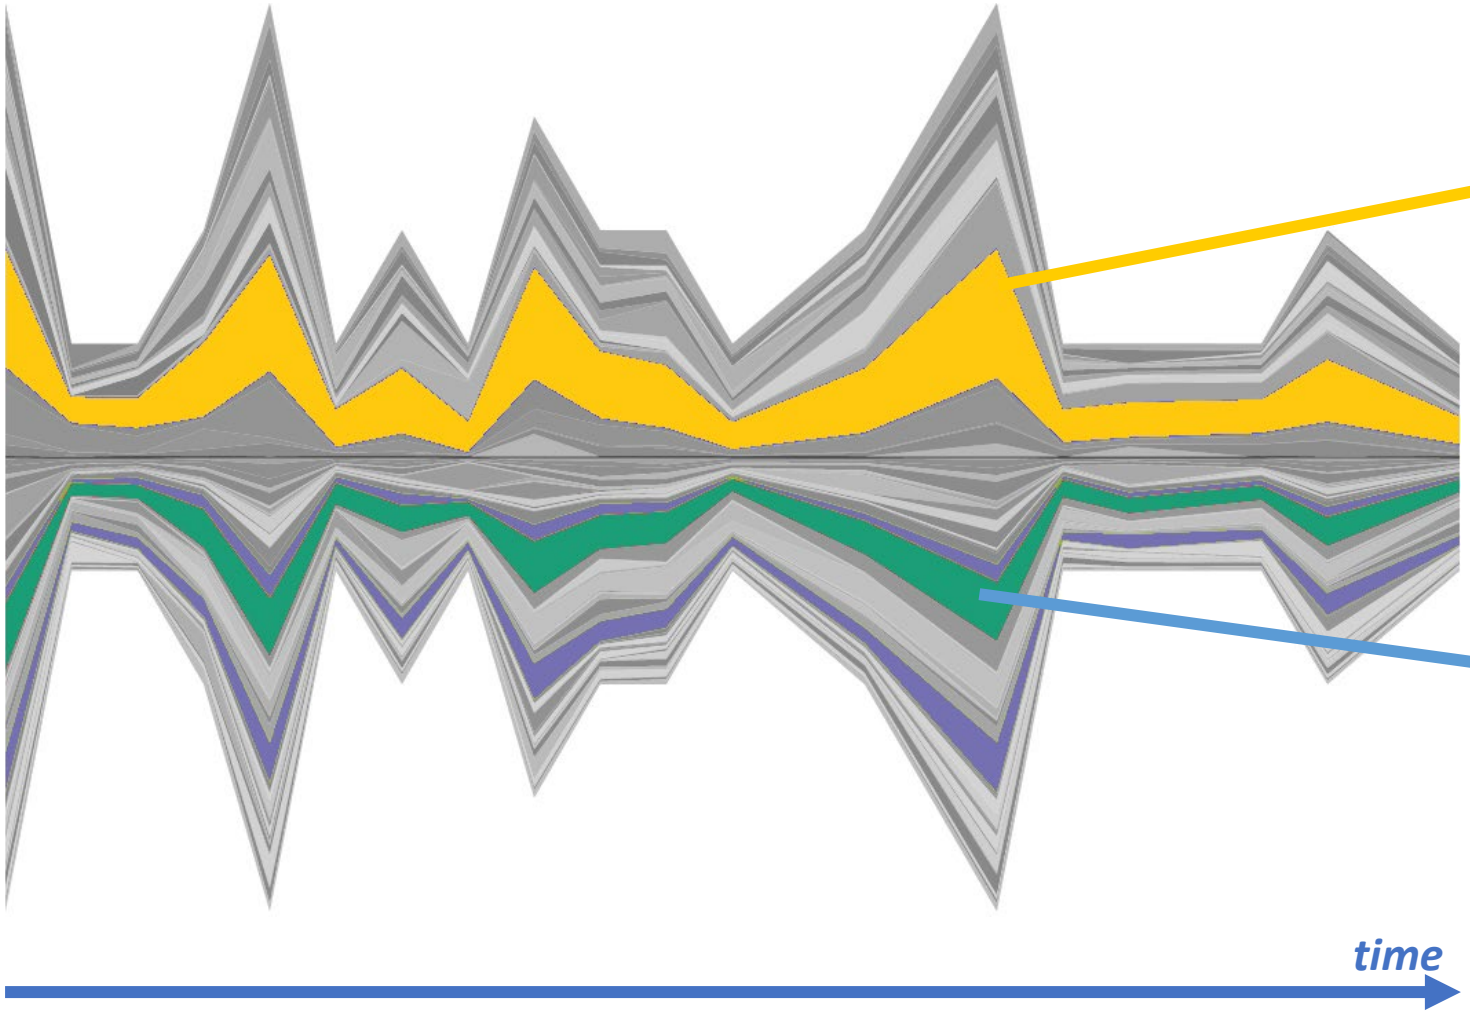

Hovering over a topic from one model…

…highlights its most closely aligned topics in the other model.

Interactive Visual Comparison of Object Embeddings

Florian Heimerl, Christoph Kralj, Torsten Möller and Michael Gleicher

University of Wisconsin – Madison, University of Vienna

Submitted for Publication

The Problem:
Compare two embeddings

We are interested in the relationships between objects

local structure

Not their positions in space

global structure

back to the introductory example...

10 dimensional data mapped to 2 dimensions

867 objects

2 runs of TSNE (different random seeds)

Perplexity 30, Learning Rate 200

Local structure (near neighbors) should be preserved

(that's what the algorithm does)

Very different shapes?

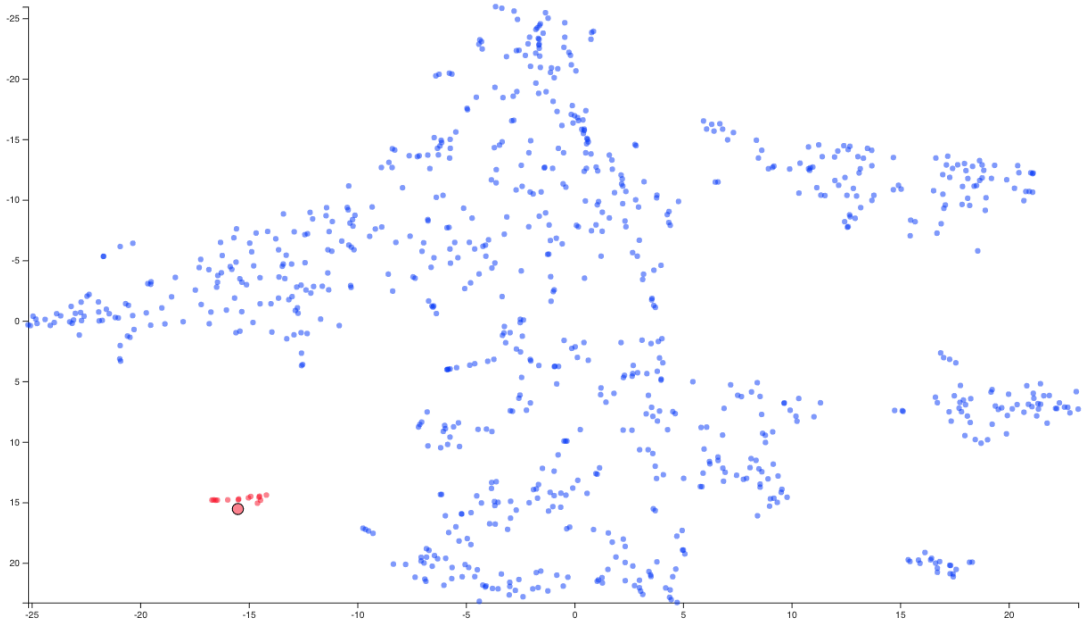

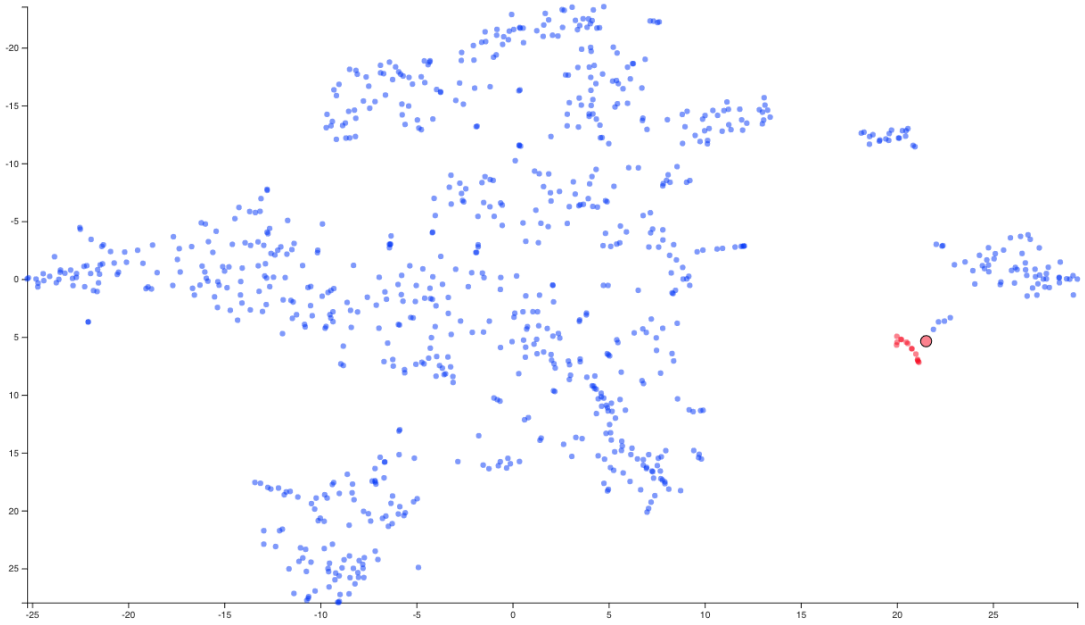

Similar neighbors

(points close in one are close in the other)

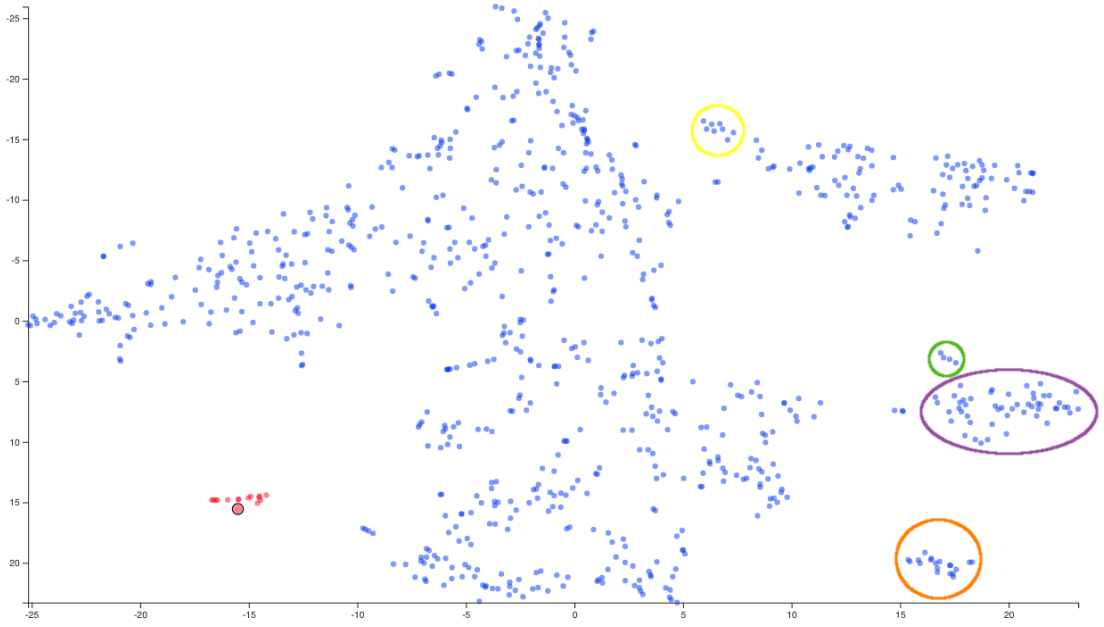

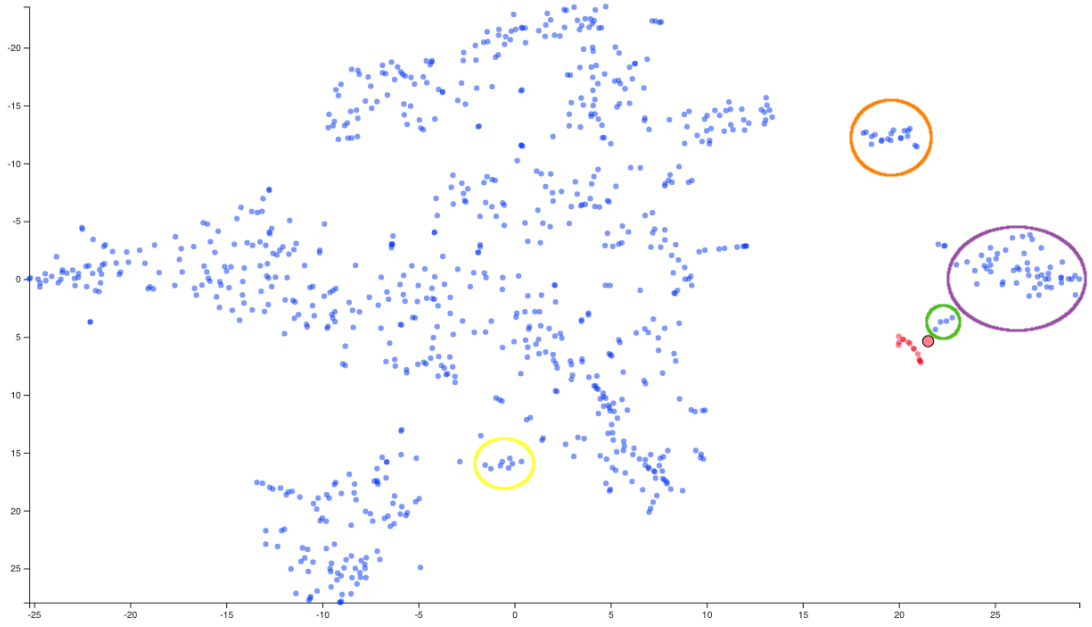

Is the local structure really similar?

measure difference/similarity?	Metrics (multiple)
assess and localize it?	Summary views
interpret and diagnose it?	Link between views
identify exemplars?	Connect to detail views
understand context?	Connect back to global views

Similar neighbors

(points close in one are close in the other)

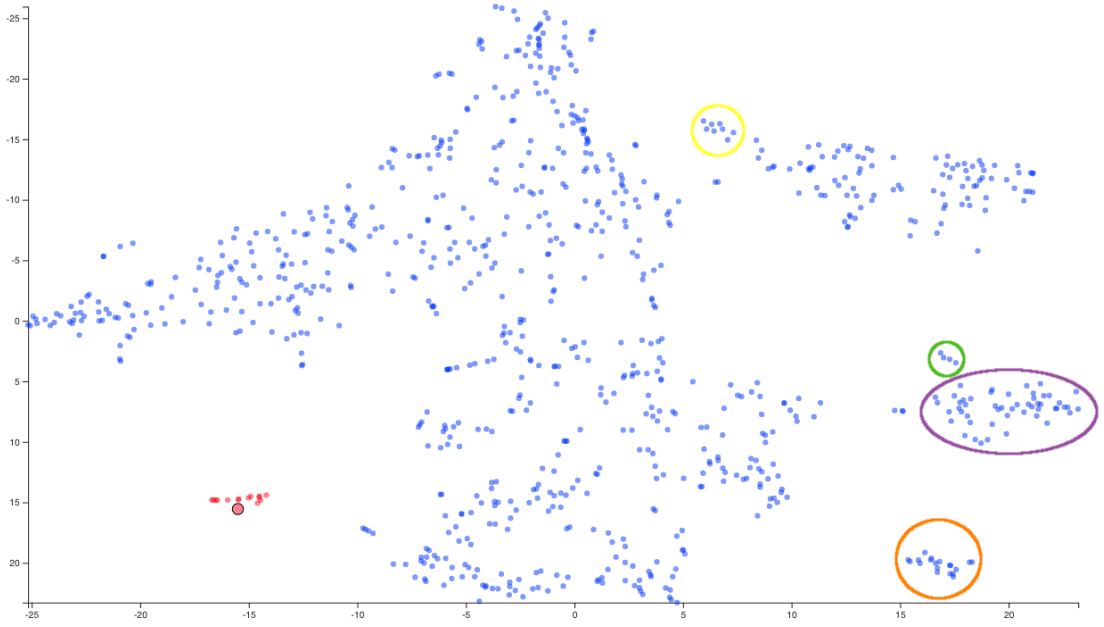

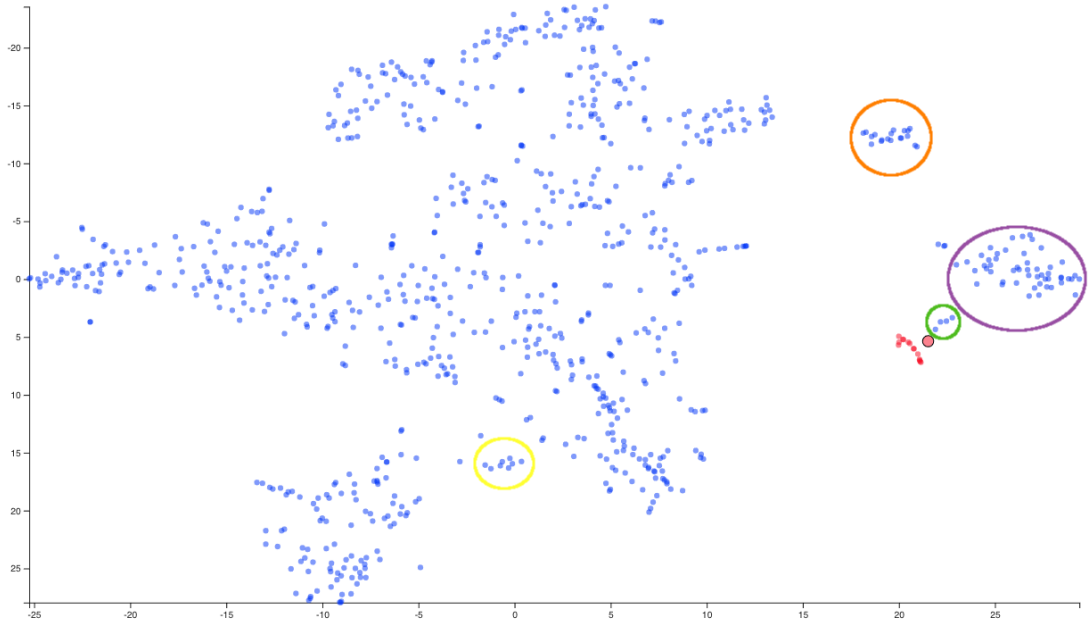

How many objects have the same nearest neighbor?

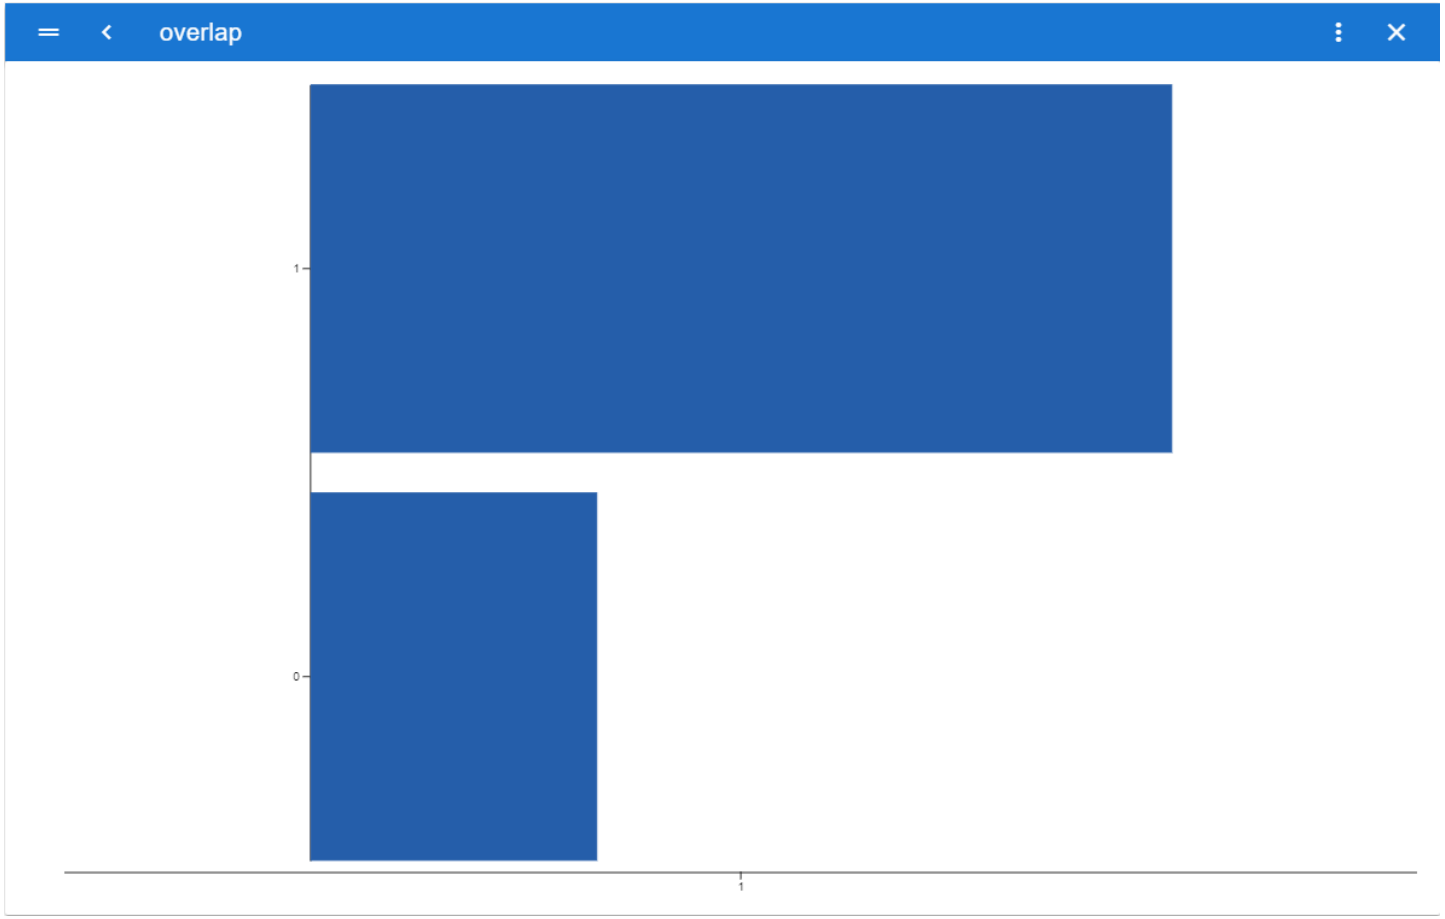

651

216

Color Encoding

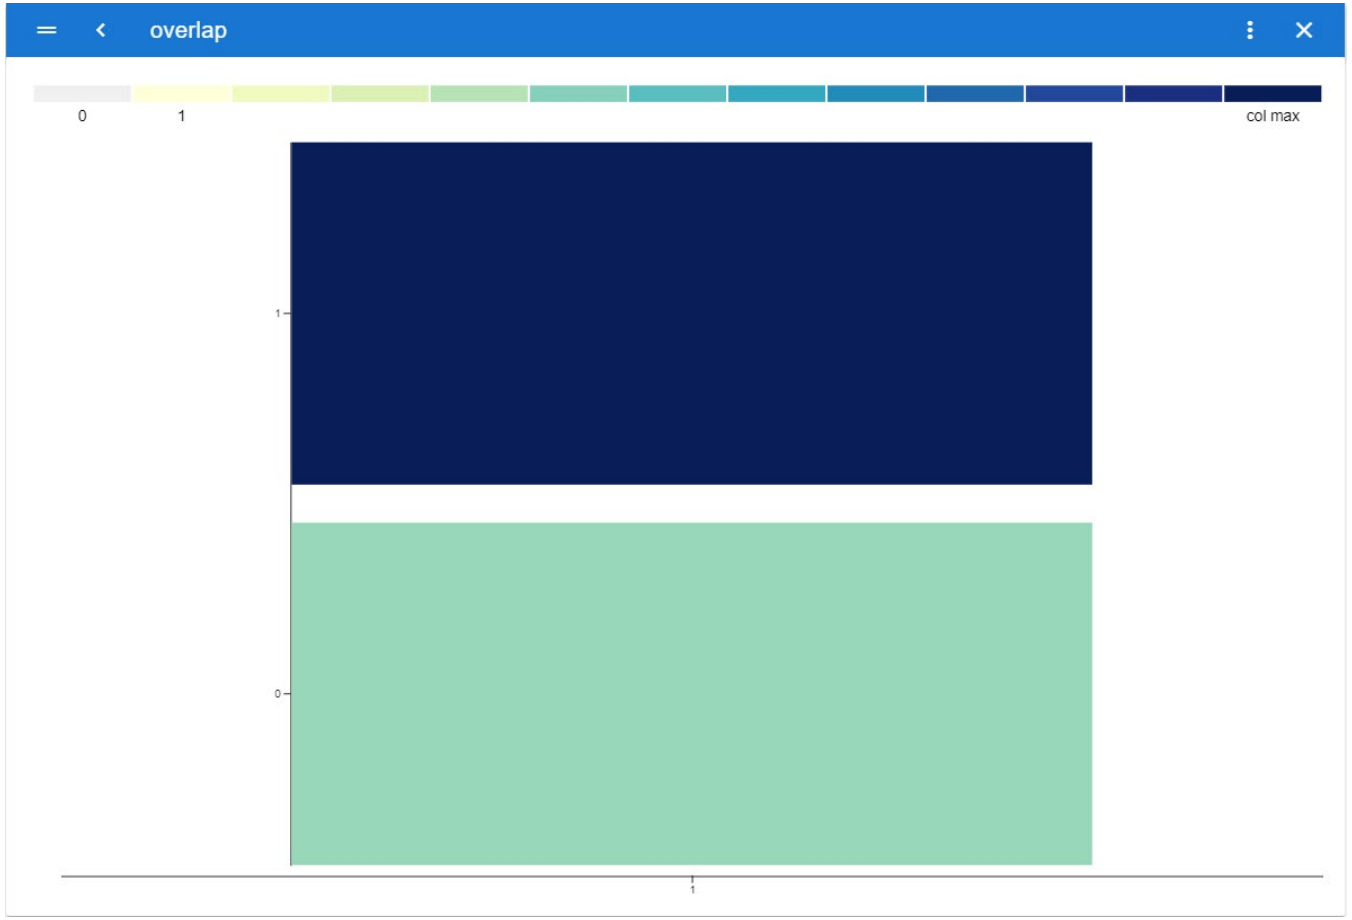

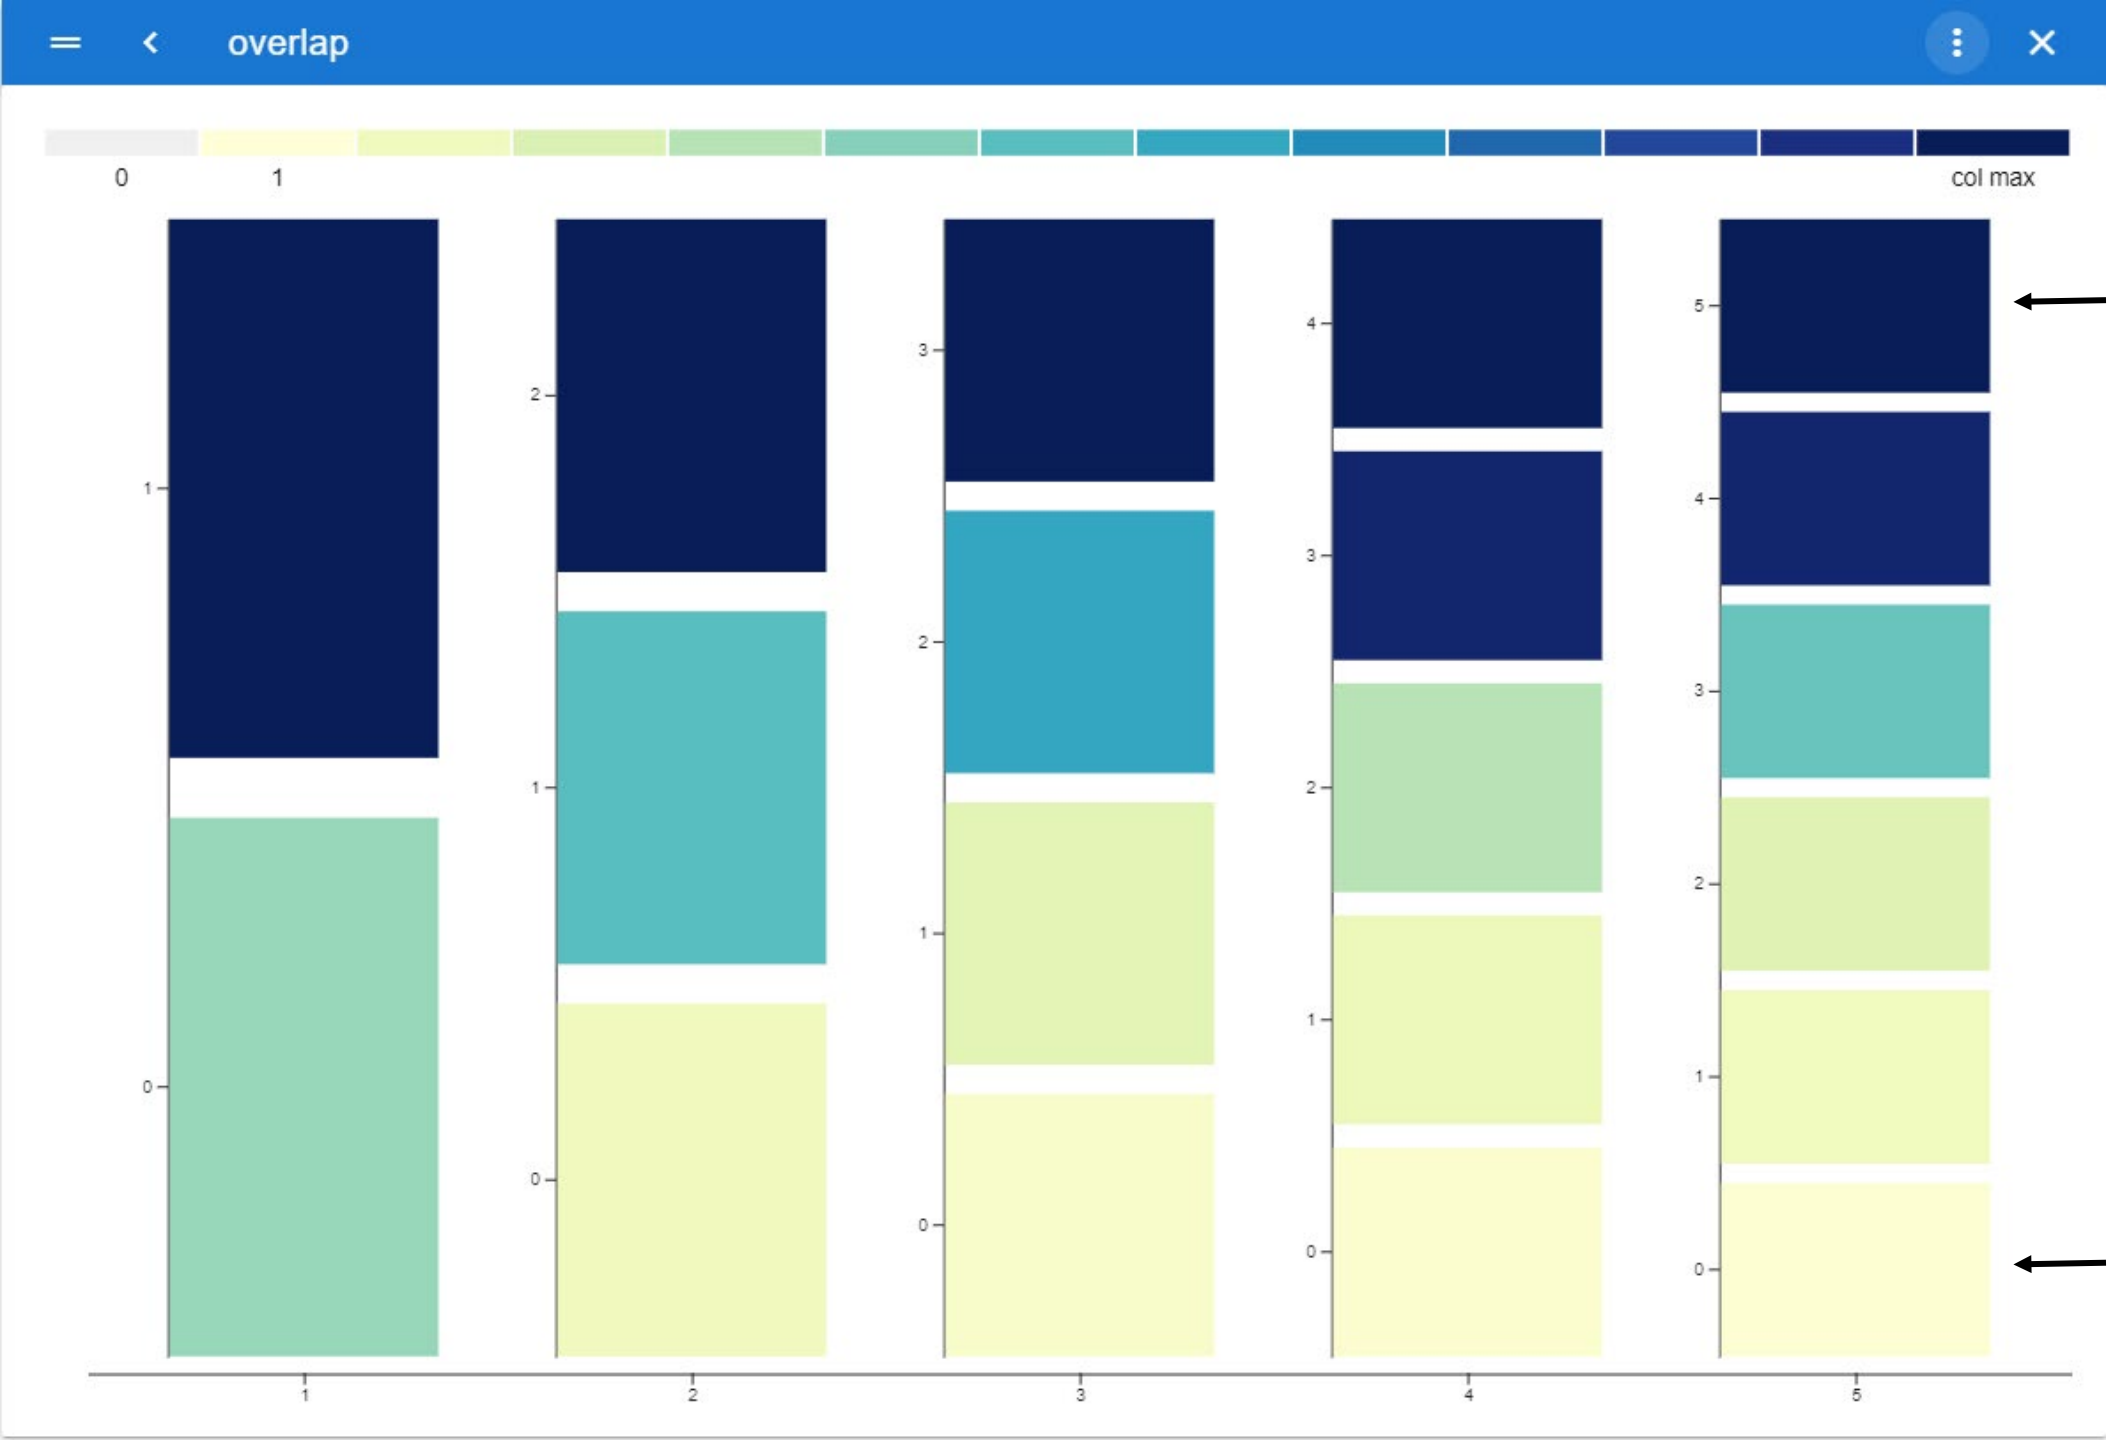
overlap
0
1
col max
327 Items
7 Items

Where are those 7?

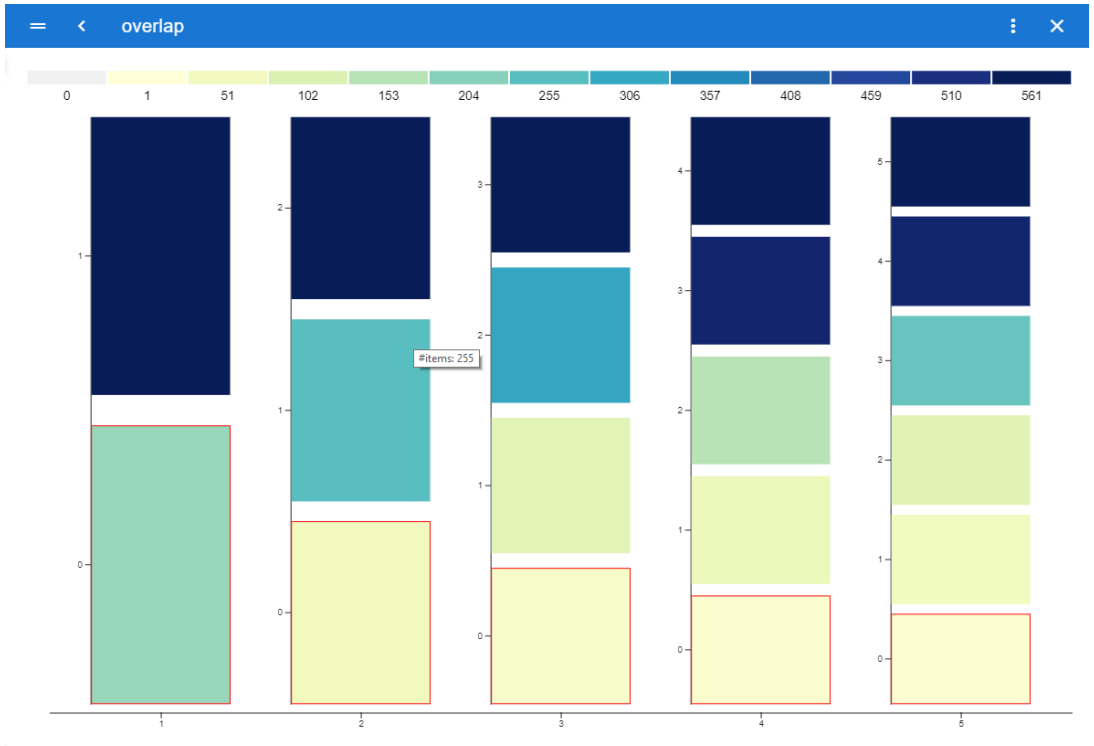

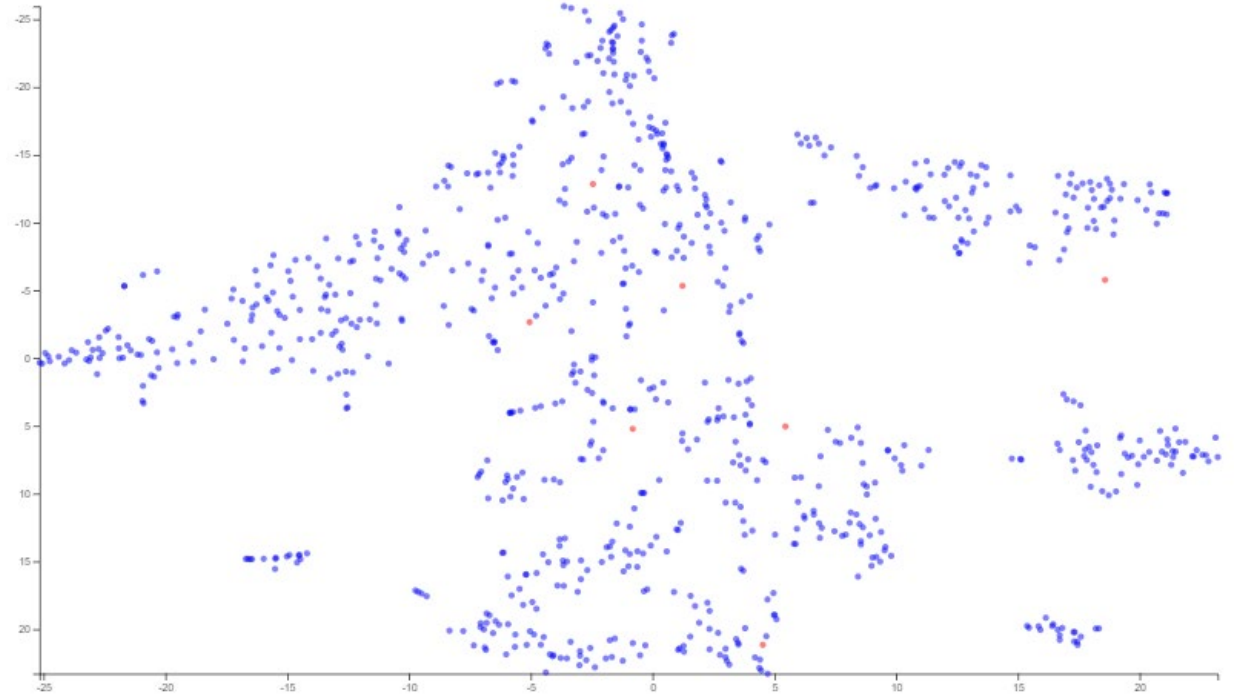

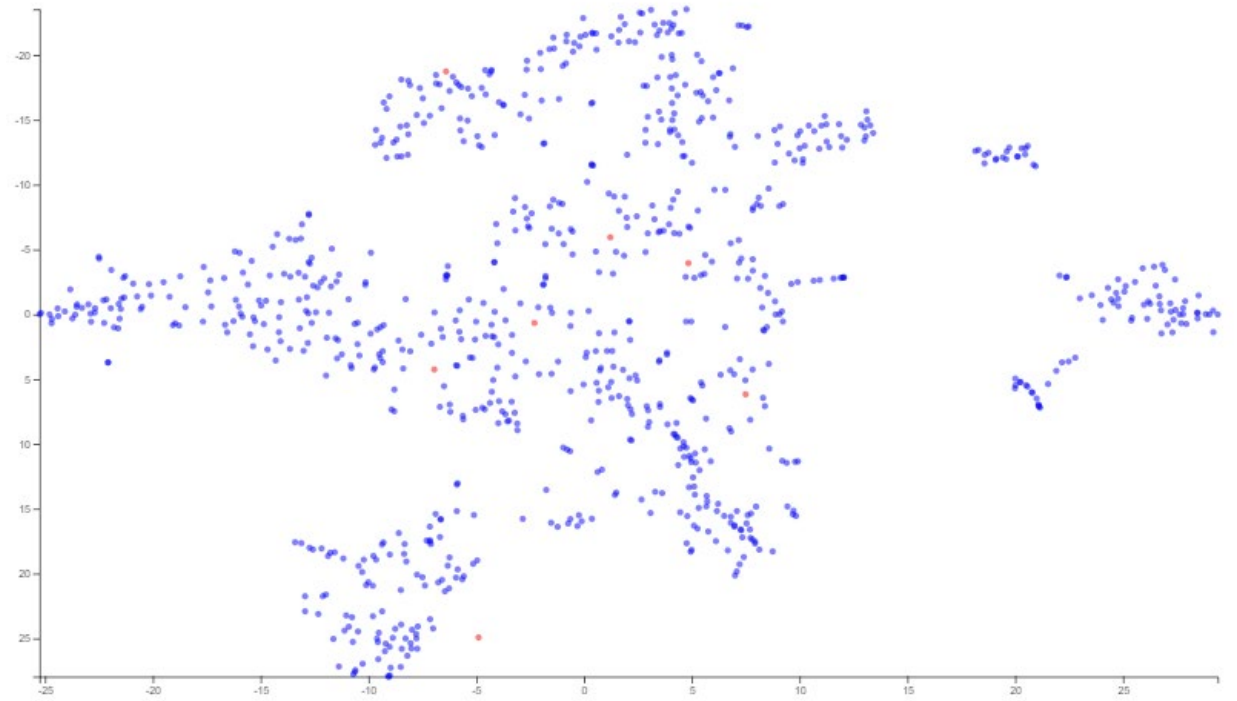

What are those 7?

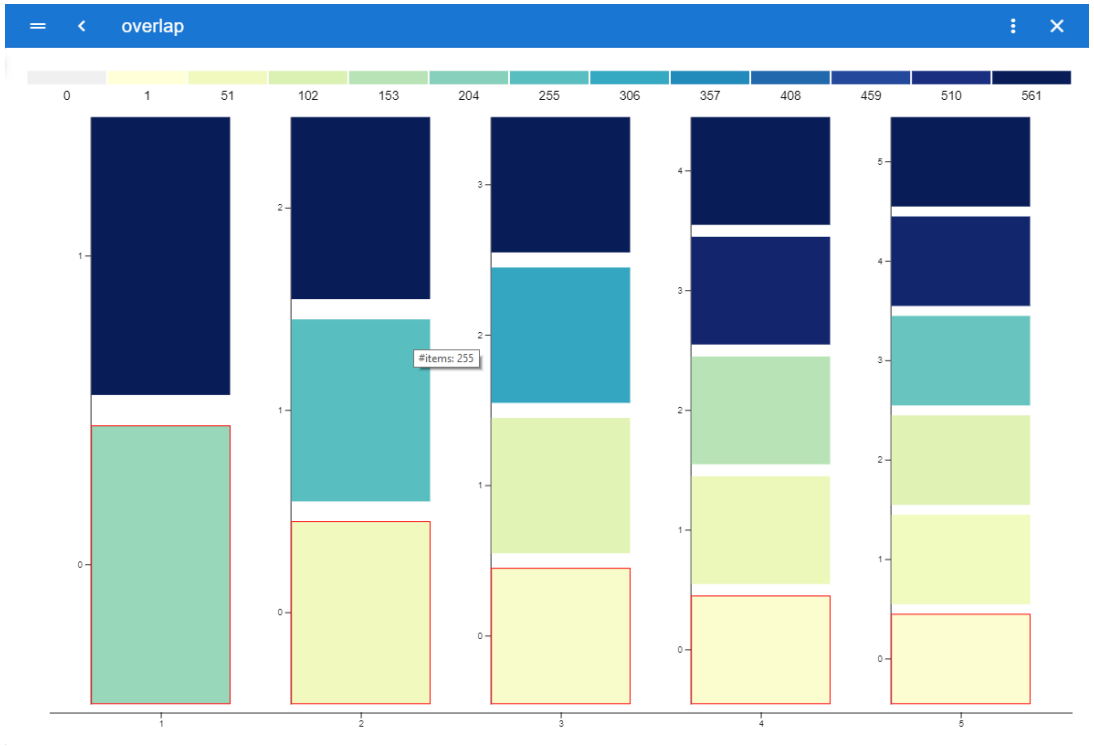

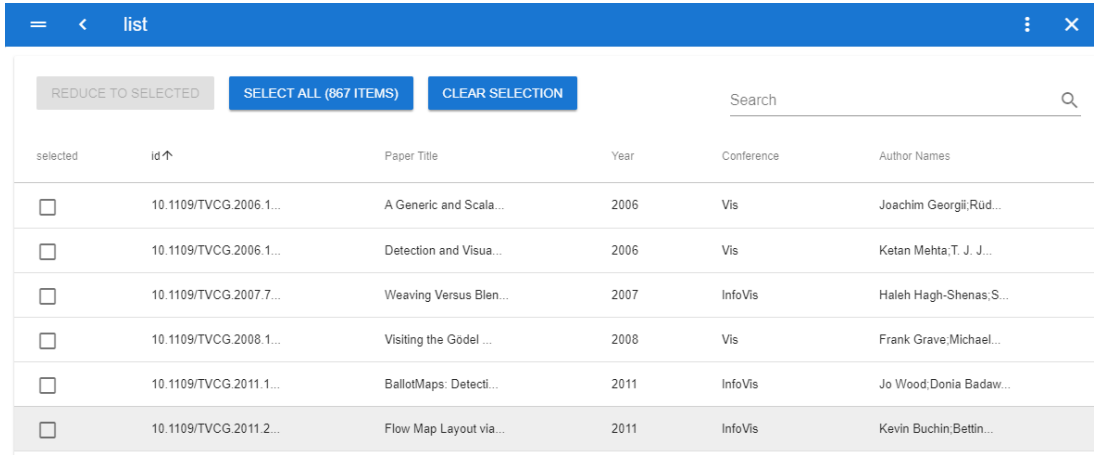

Spread view:

Where do they go? (outside of the top N)

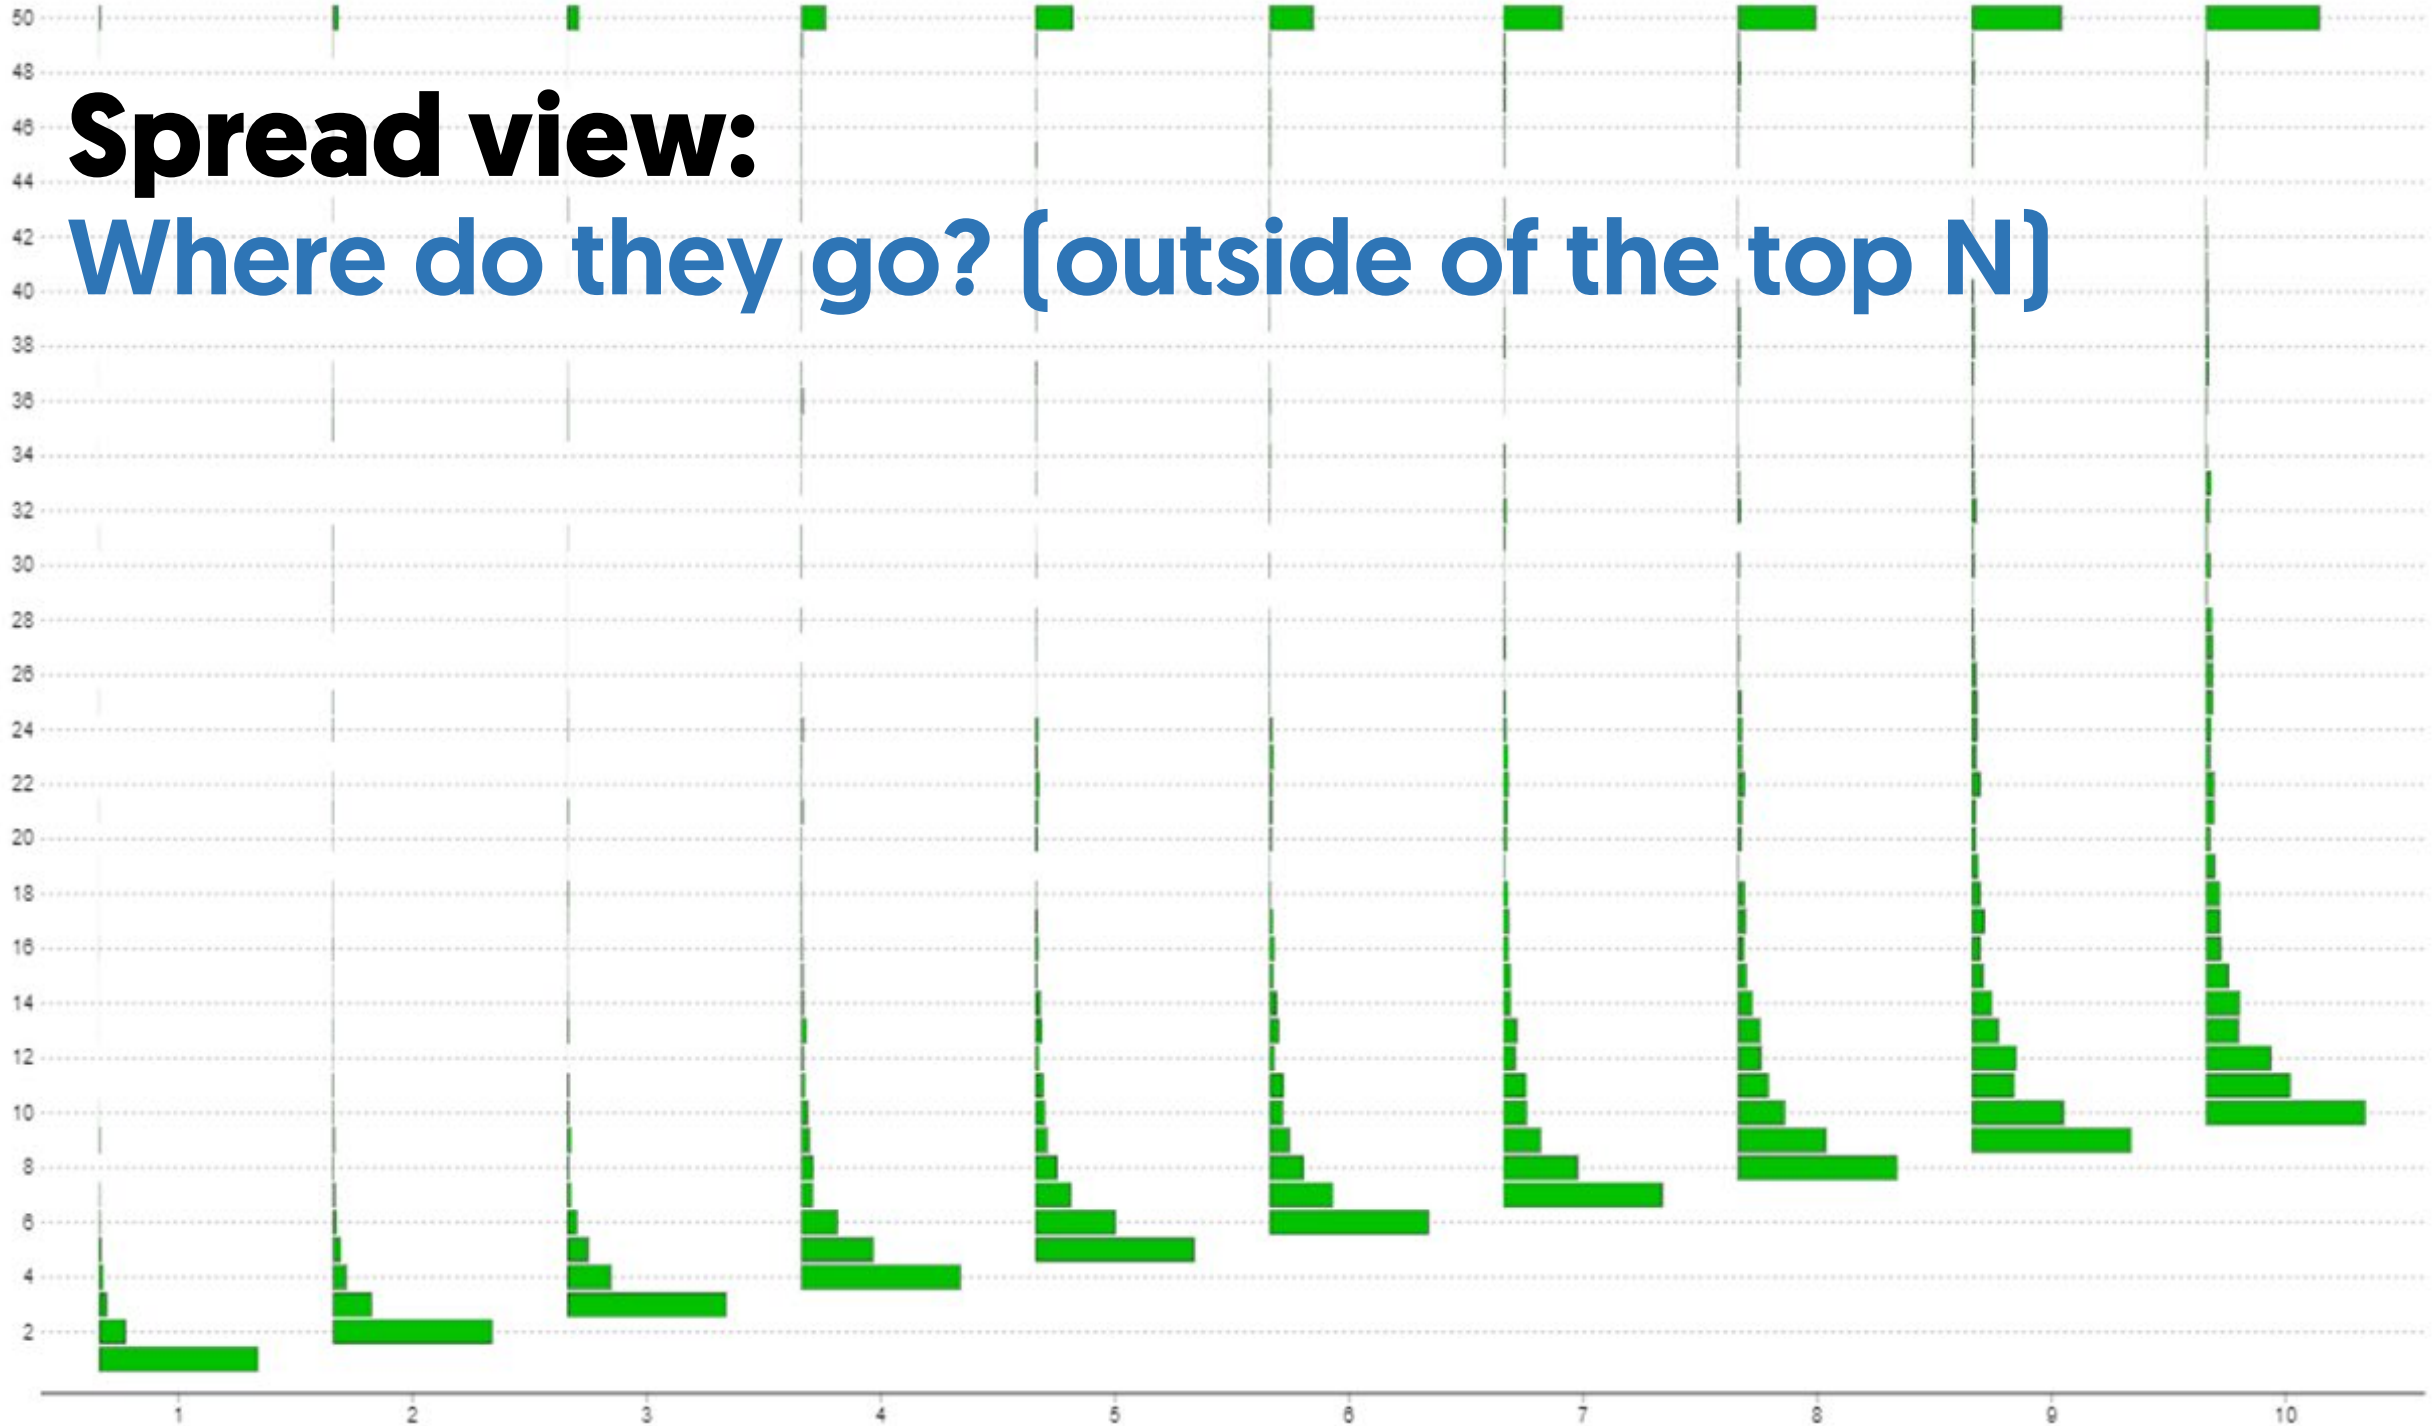

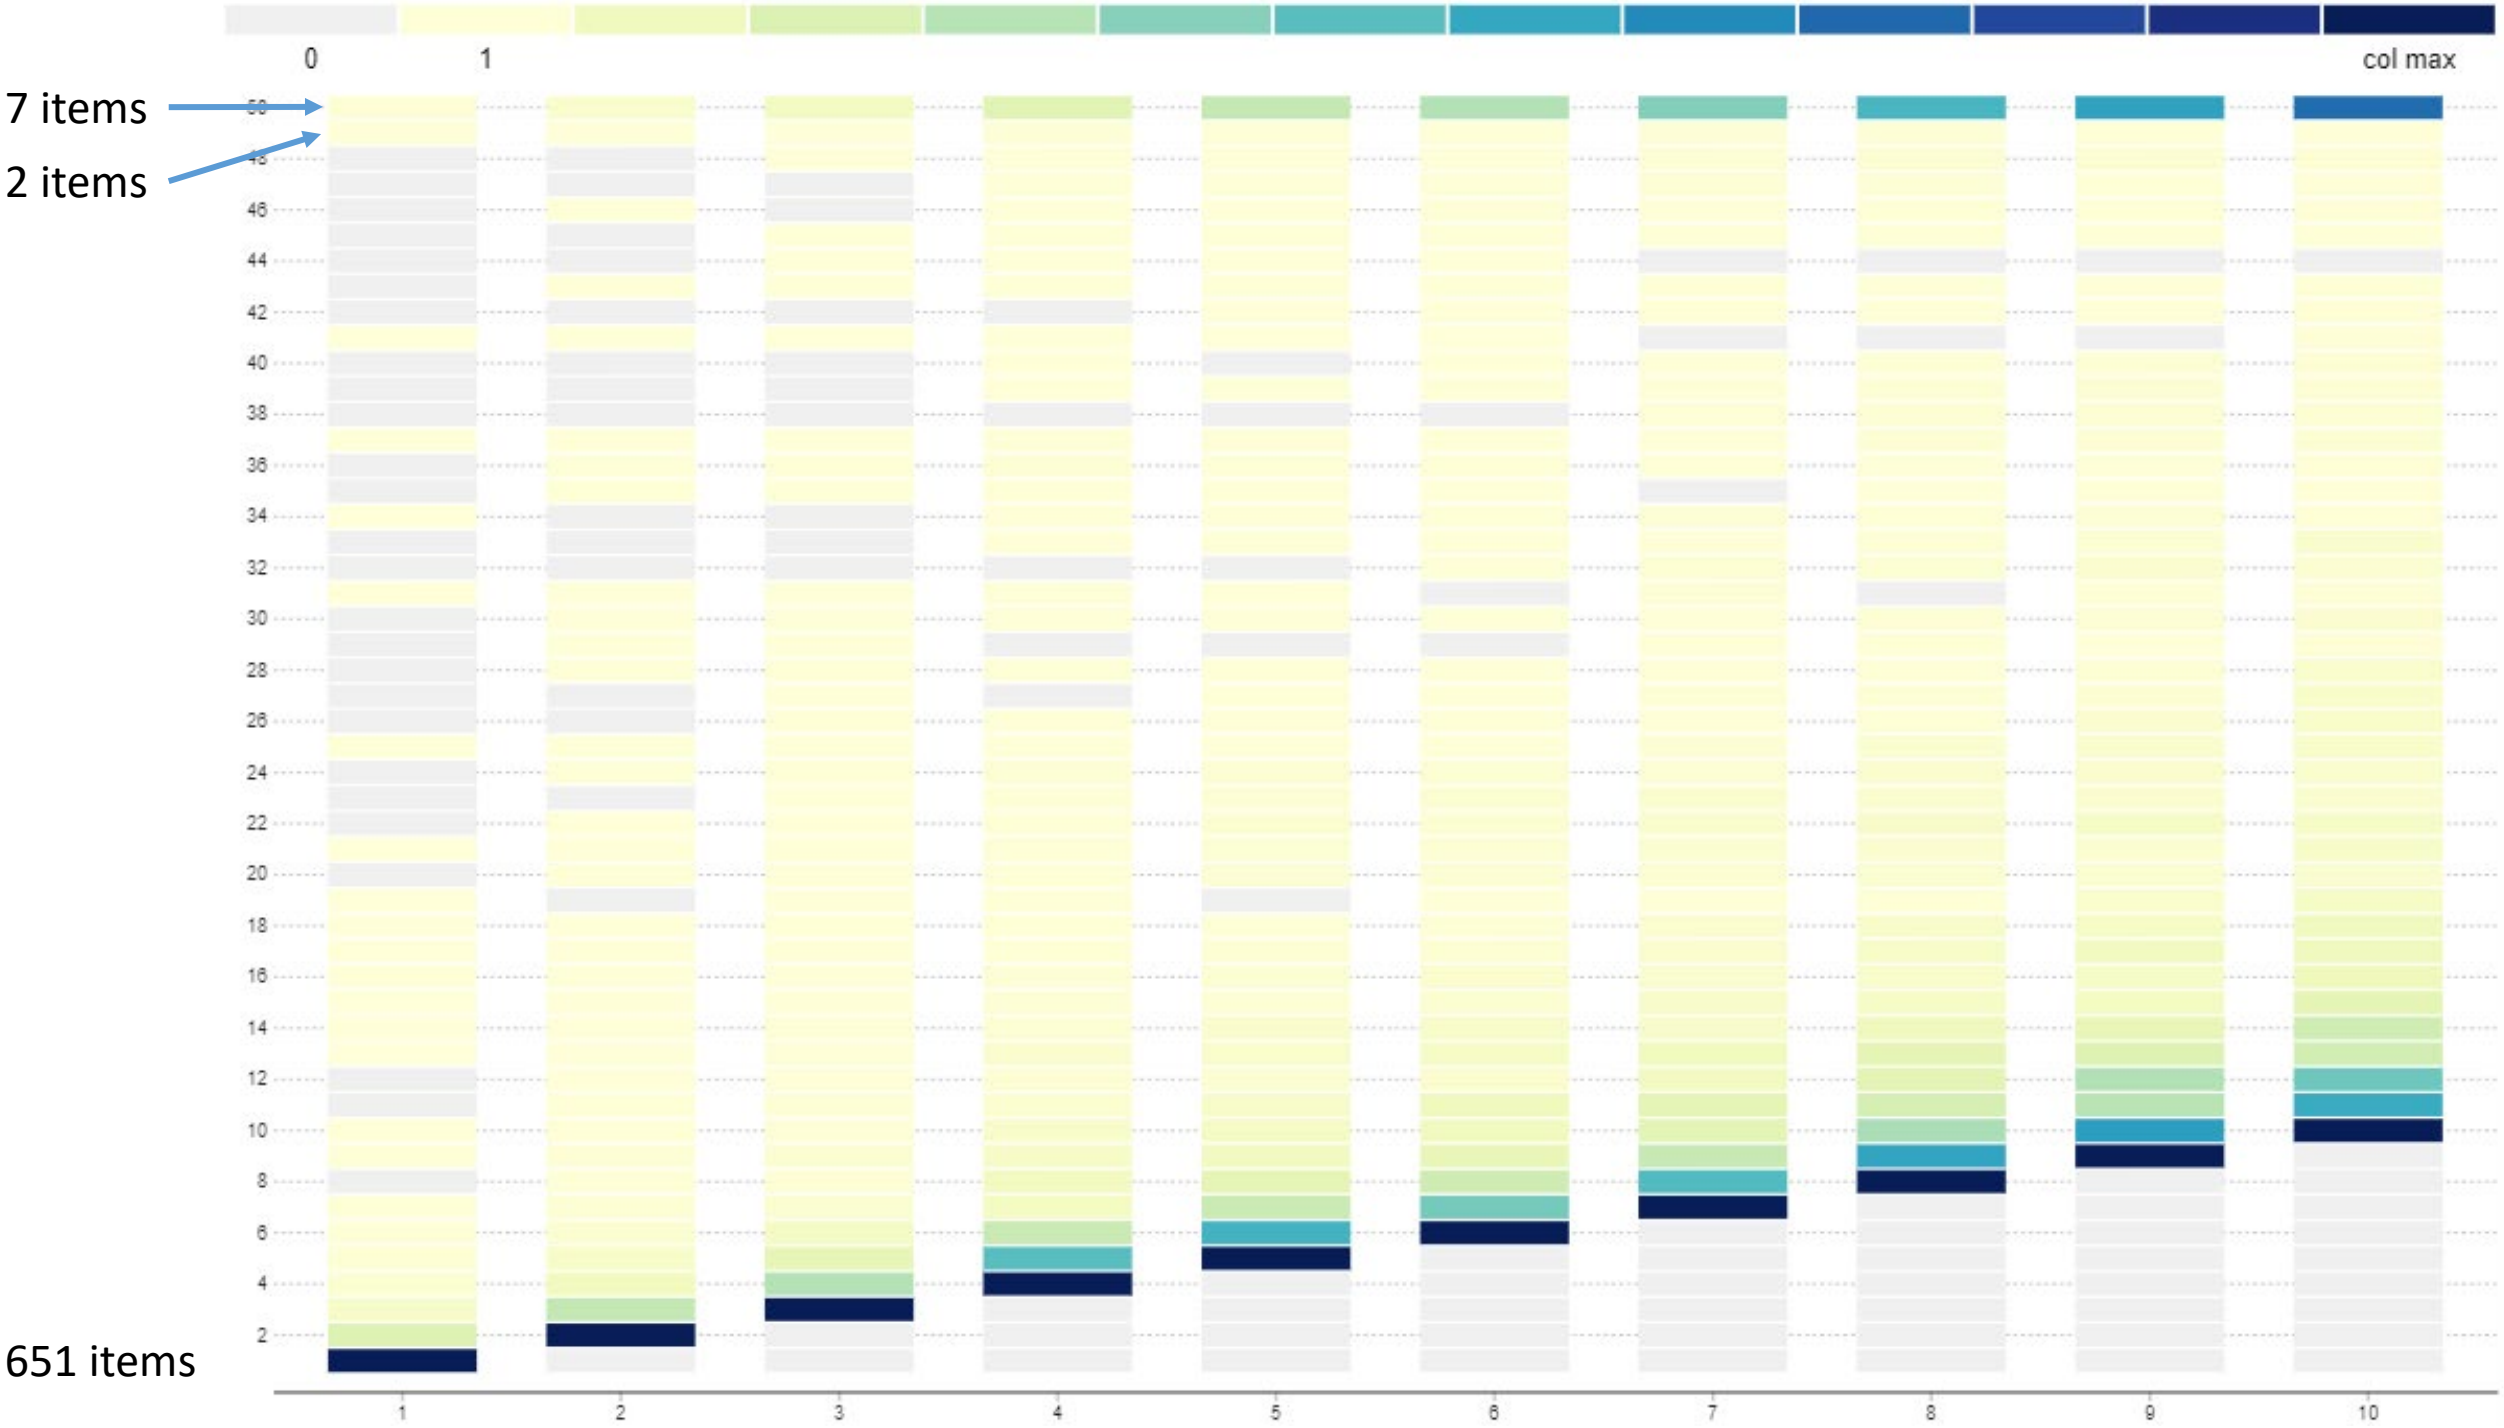
0
1
col max
7 items
2 items
651 items
50
48
46
44
42
40
38
36
34
32
30
28
26
24
22
20
18
16
14
12
10
8
6
4
2
1
2
3
4
5
6
7
8
9
10

Select 2 groups

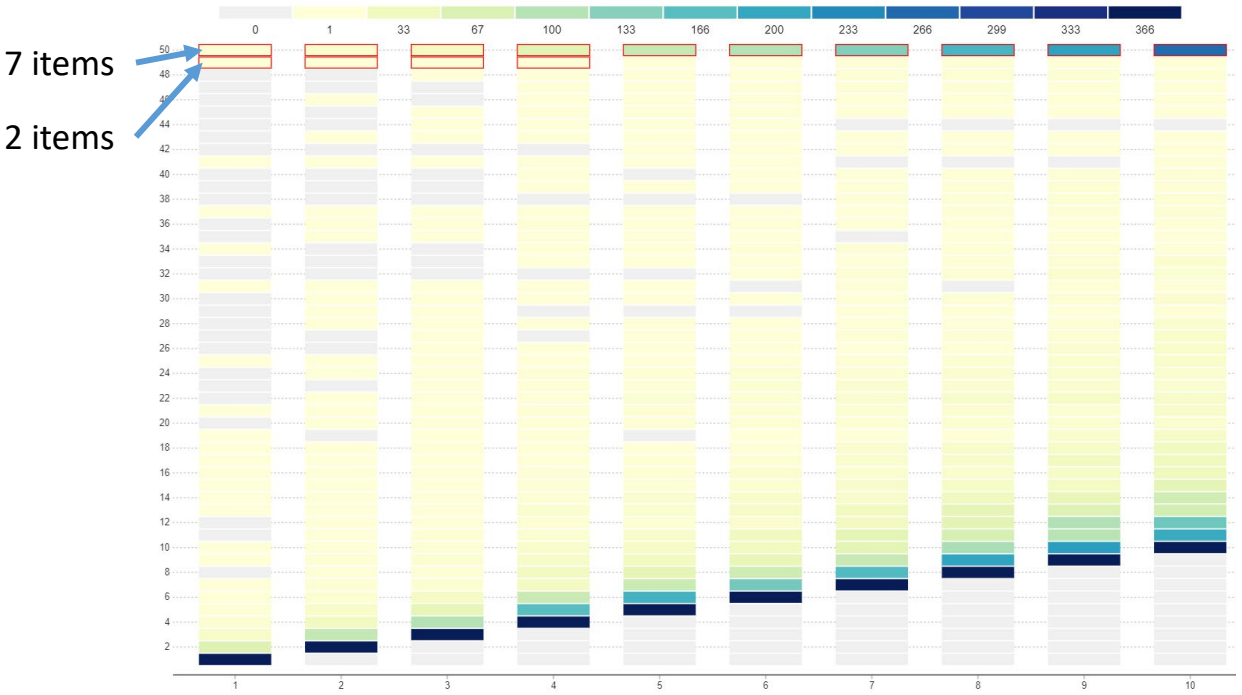

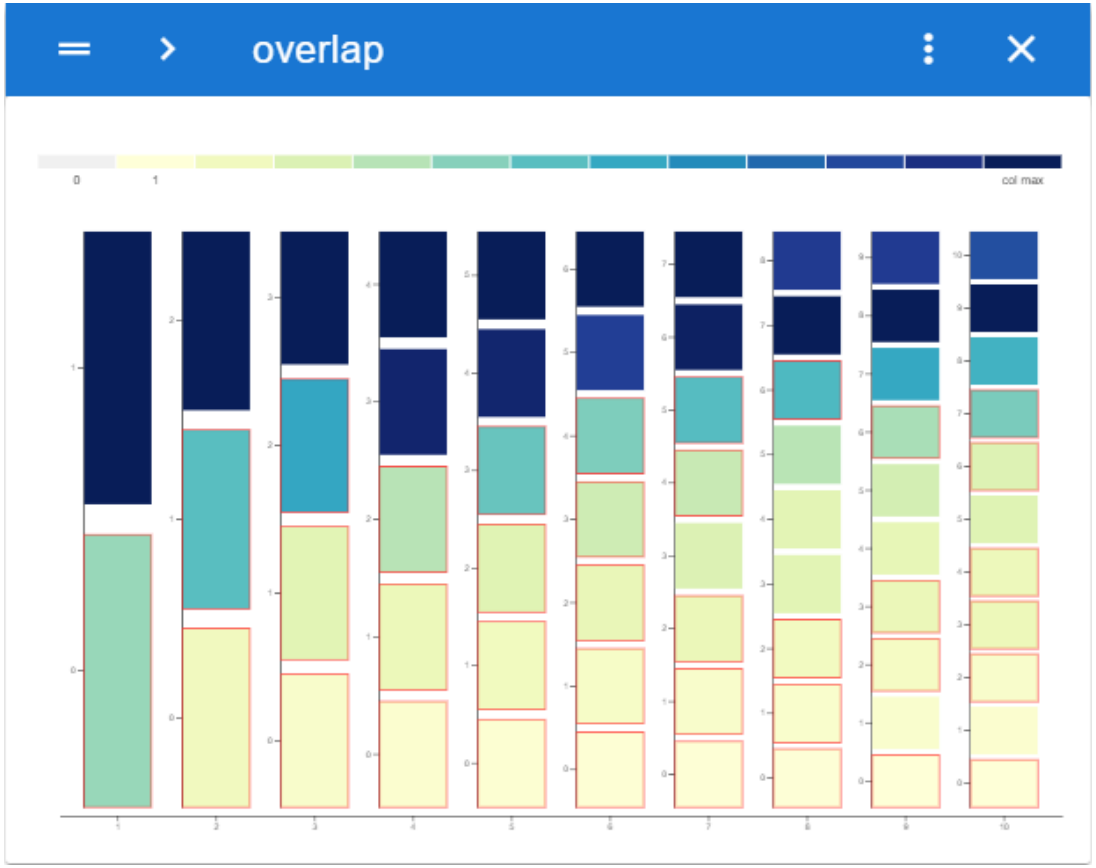

Distance view (highlight those 9)

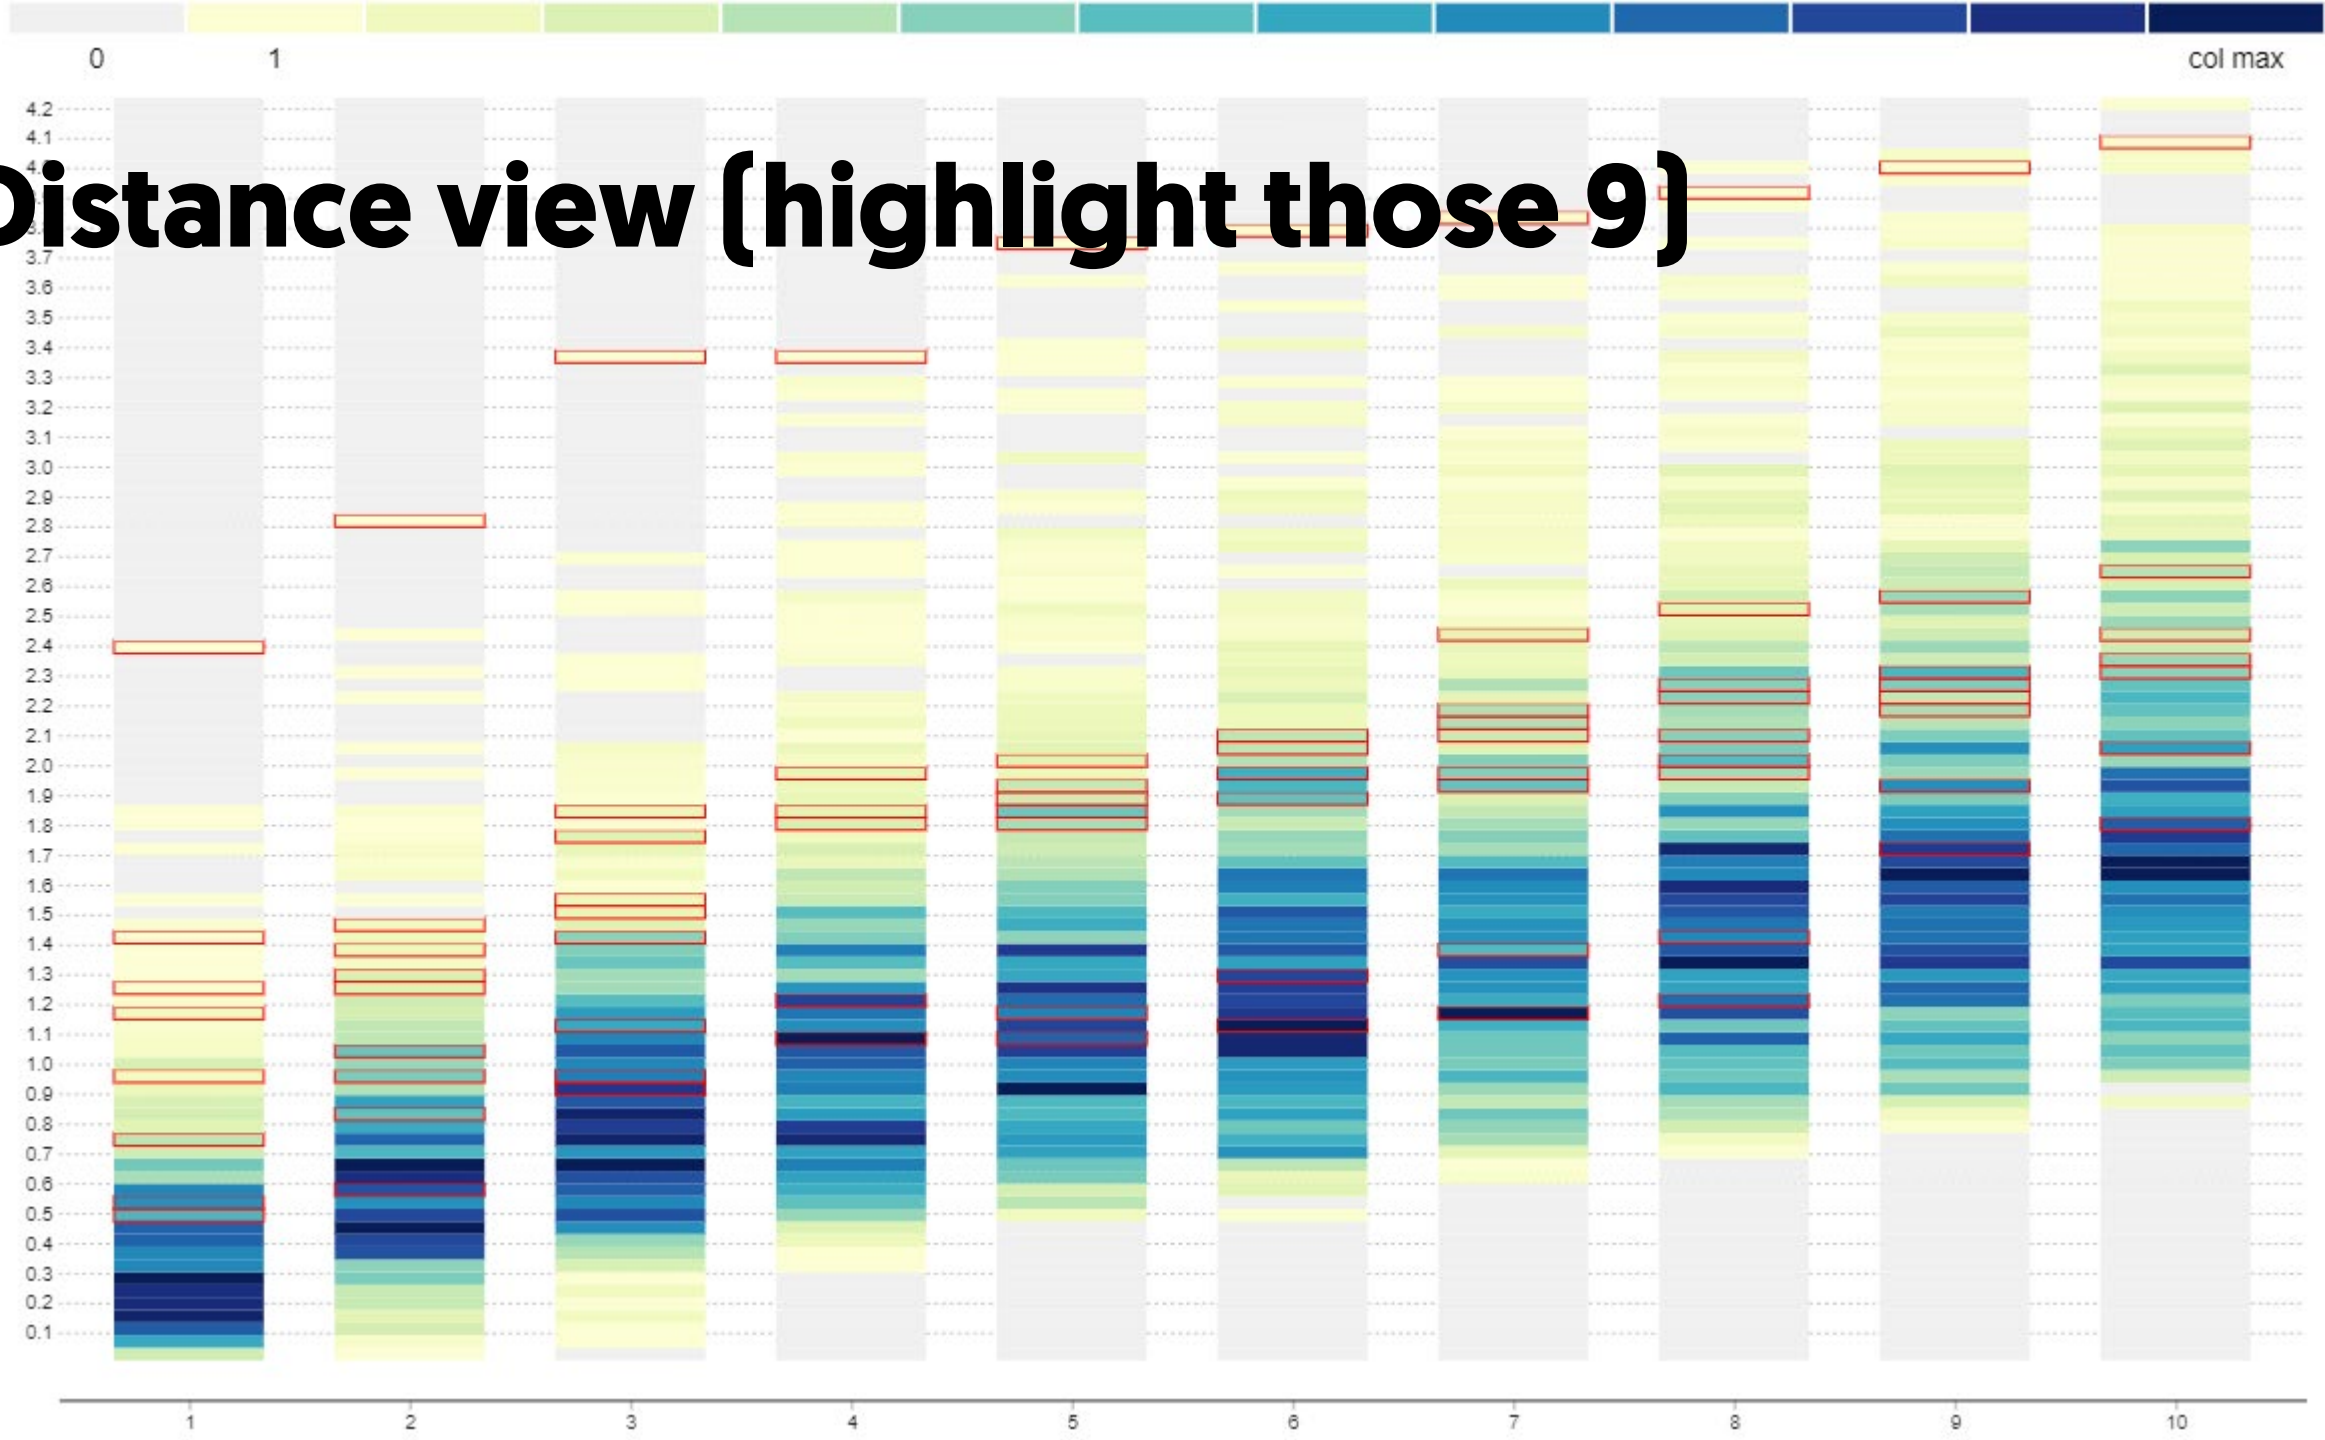

Clusters (k-Means)

	7	2	1	9	6	5	0	8	3	4	-1	
	101	111	104	68	94	107	52	64	93	73		
7	97											97
2	2	72			18							92
1		20	78			1						99
9	2		26	68								96
6		19			76							95
5						80		14	1			95
0						23	52	1	3			79
8						1		40	44			85
3						2		9	45	15		71
4										58		58
-1												

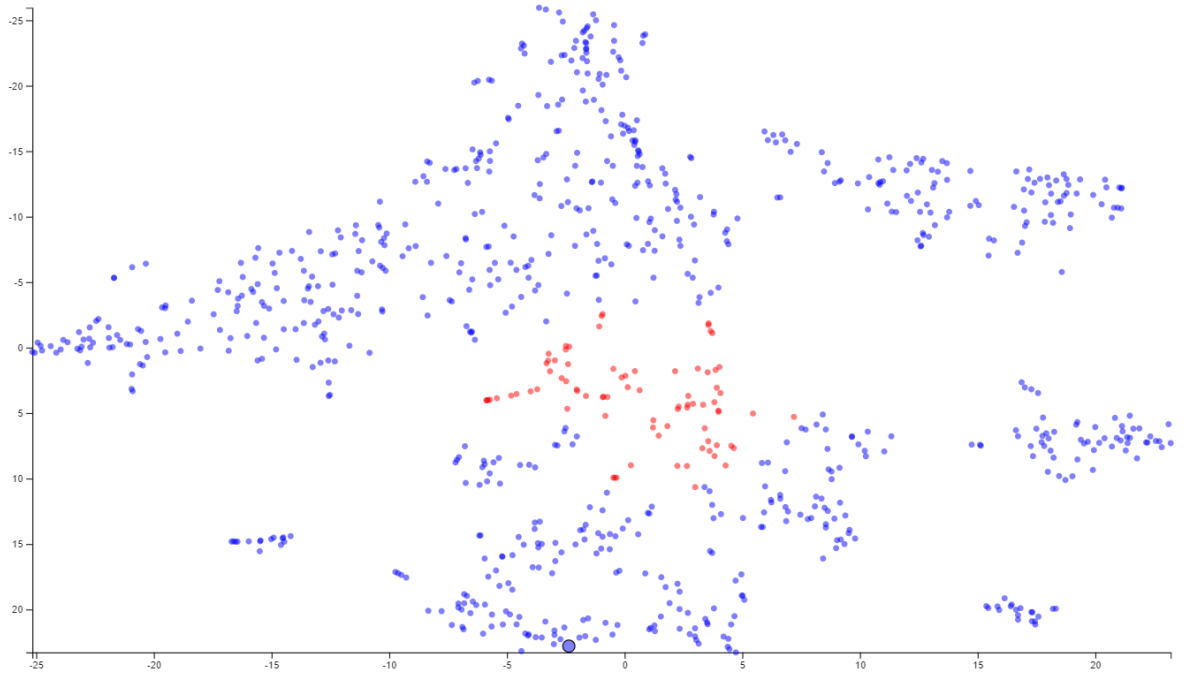

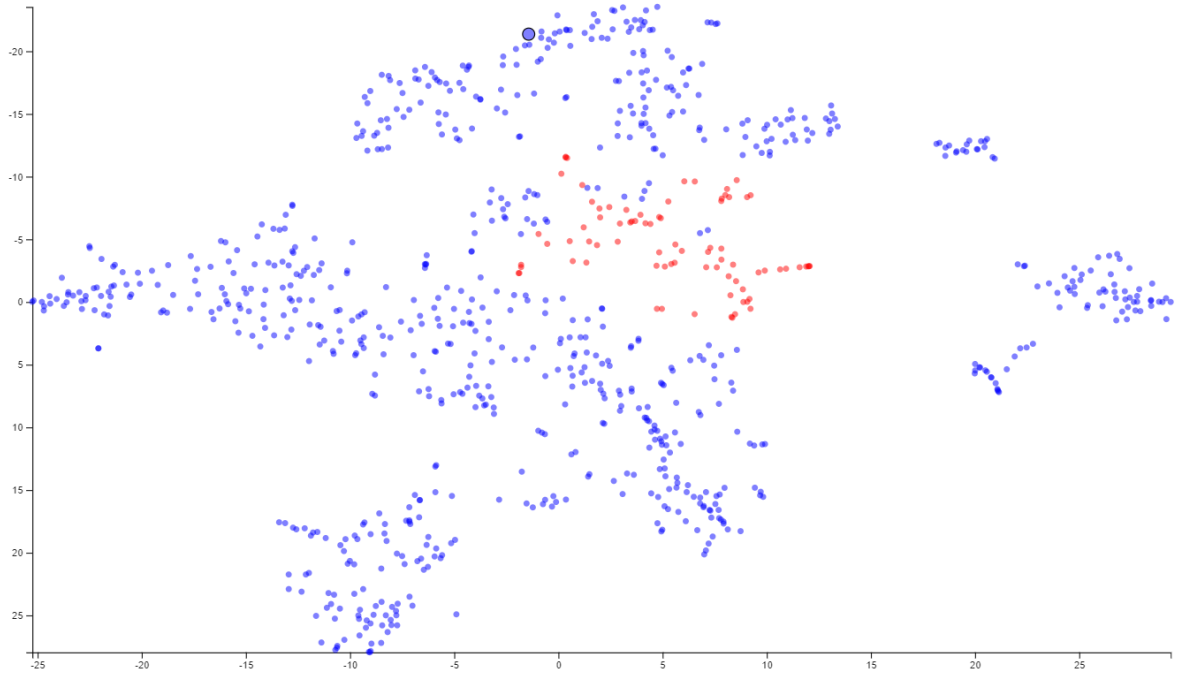

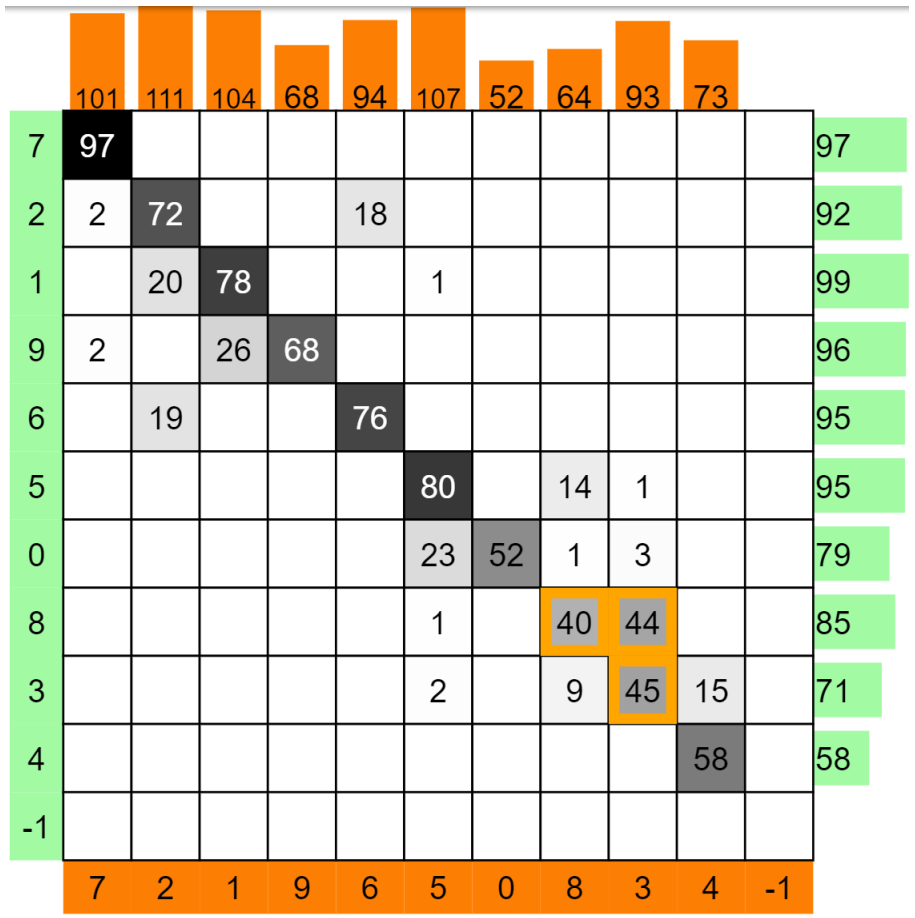
101 111 104 68 94 107 52 64 93 73
7 97 97
2 2 72 18 92
1 20 78 1 99
9 2 26 68 96
6 19 76 95
5 80 14 1 95
0 23 52 1 3 79
8 1 40 44 85
3 2 9 45 15 71
4 58 58
-1
7 2 1 9 6 5 0 8 3 4 -1

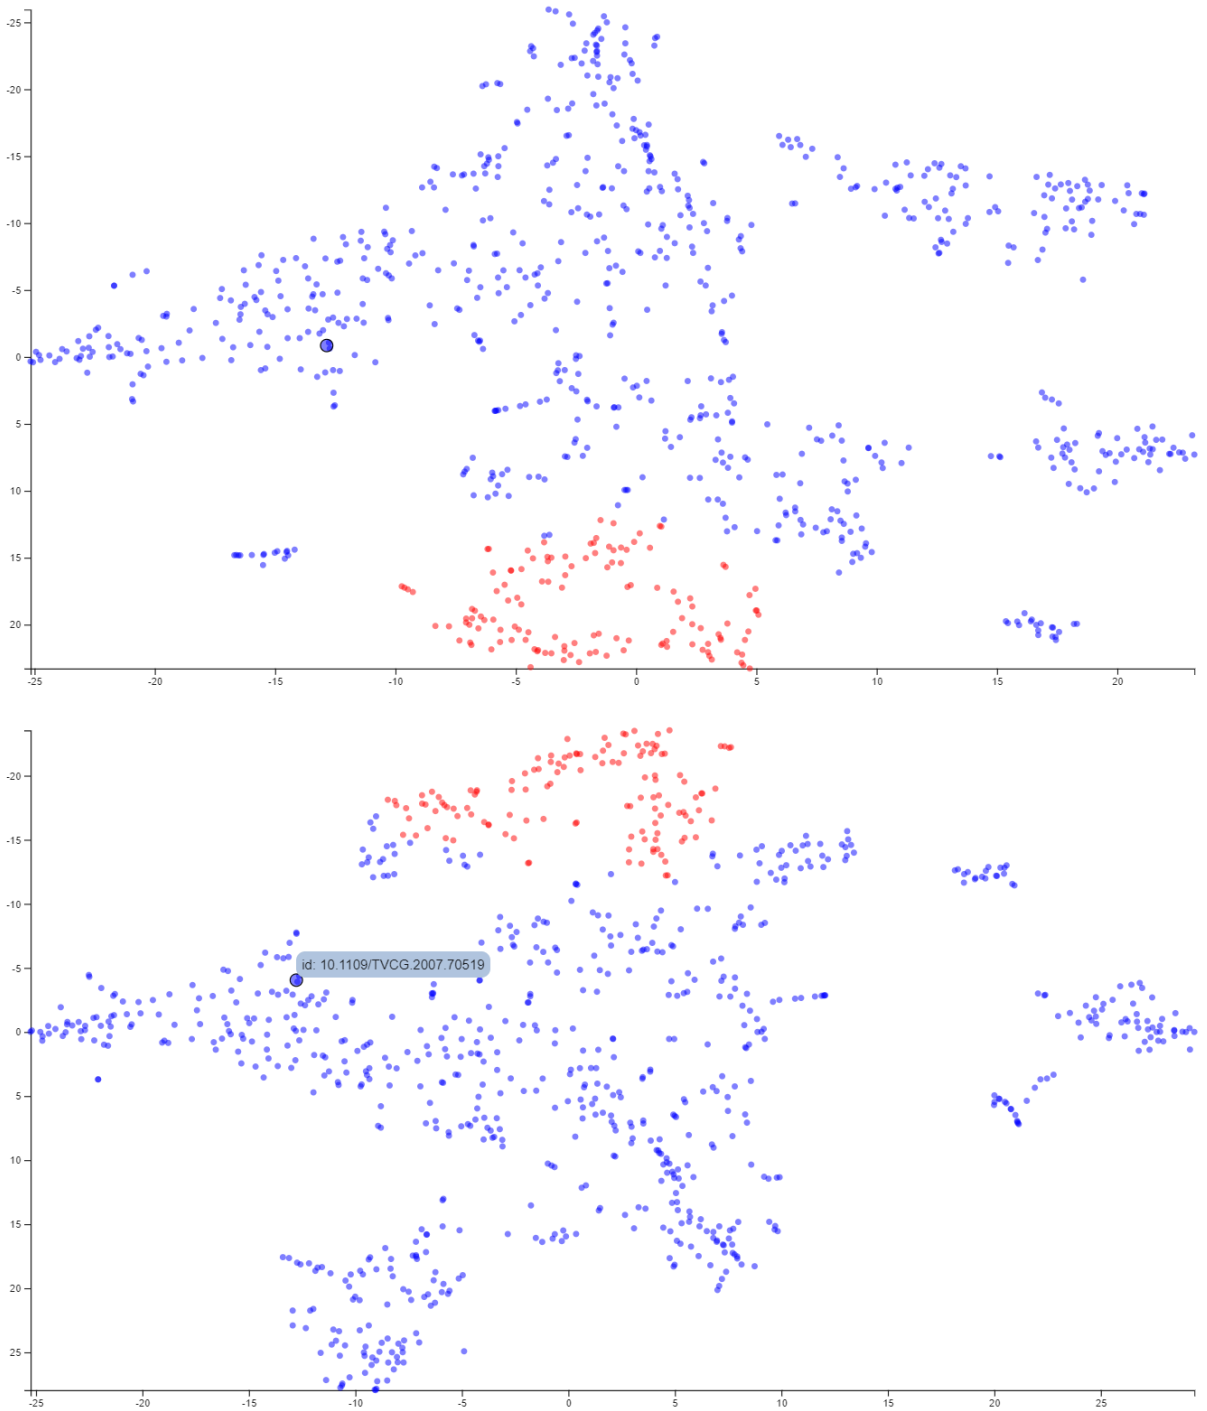
id: 10.1109/TVCG.2007.70519

Select by region

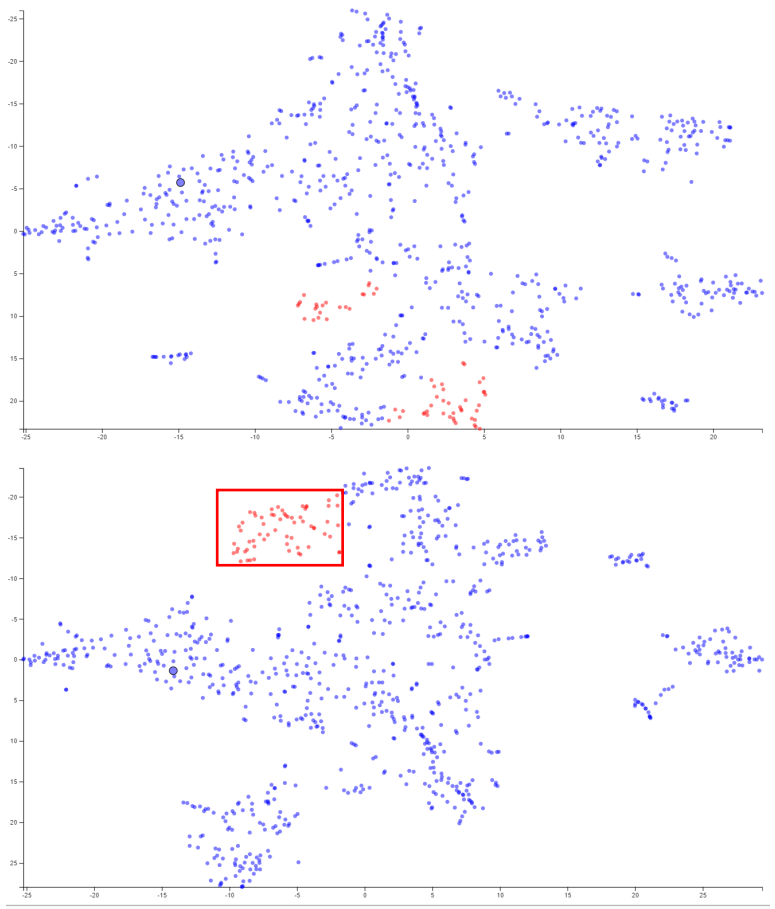

	101	111	104	68	94	107	52	64	93	73		
7	97											97
2	2	72			18							92
1		20	78			1						99
9	2		26	68								96
6		19			76							95
5						80		14	1			95
0						23	52	1	3			79
8						1		40	44			85
3						2		9	45	15		71
4										58		58
-1												
	7	2	1	9	6	5	0	8	3	4	-1	

Is this region different in terms of error?

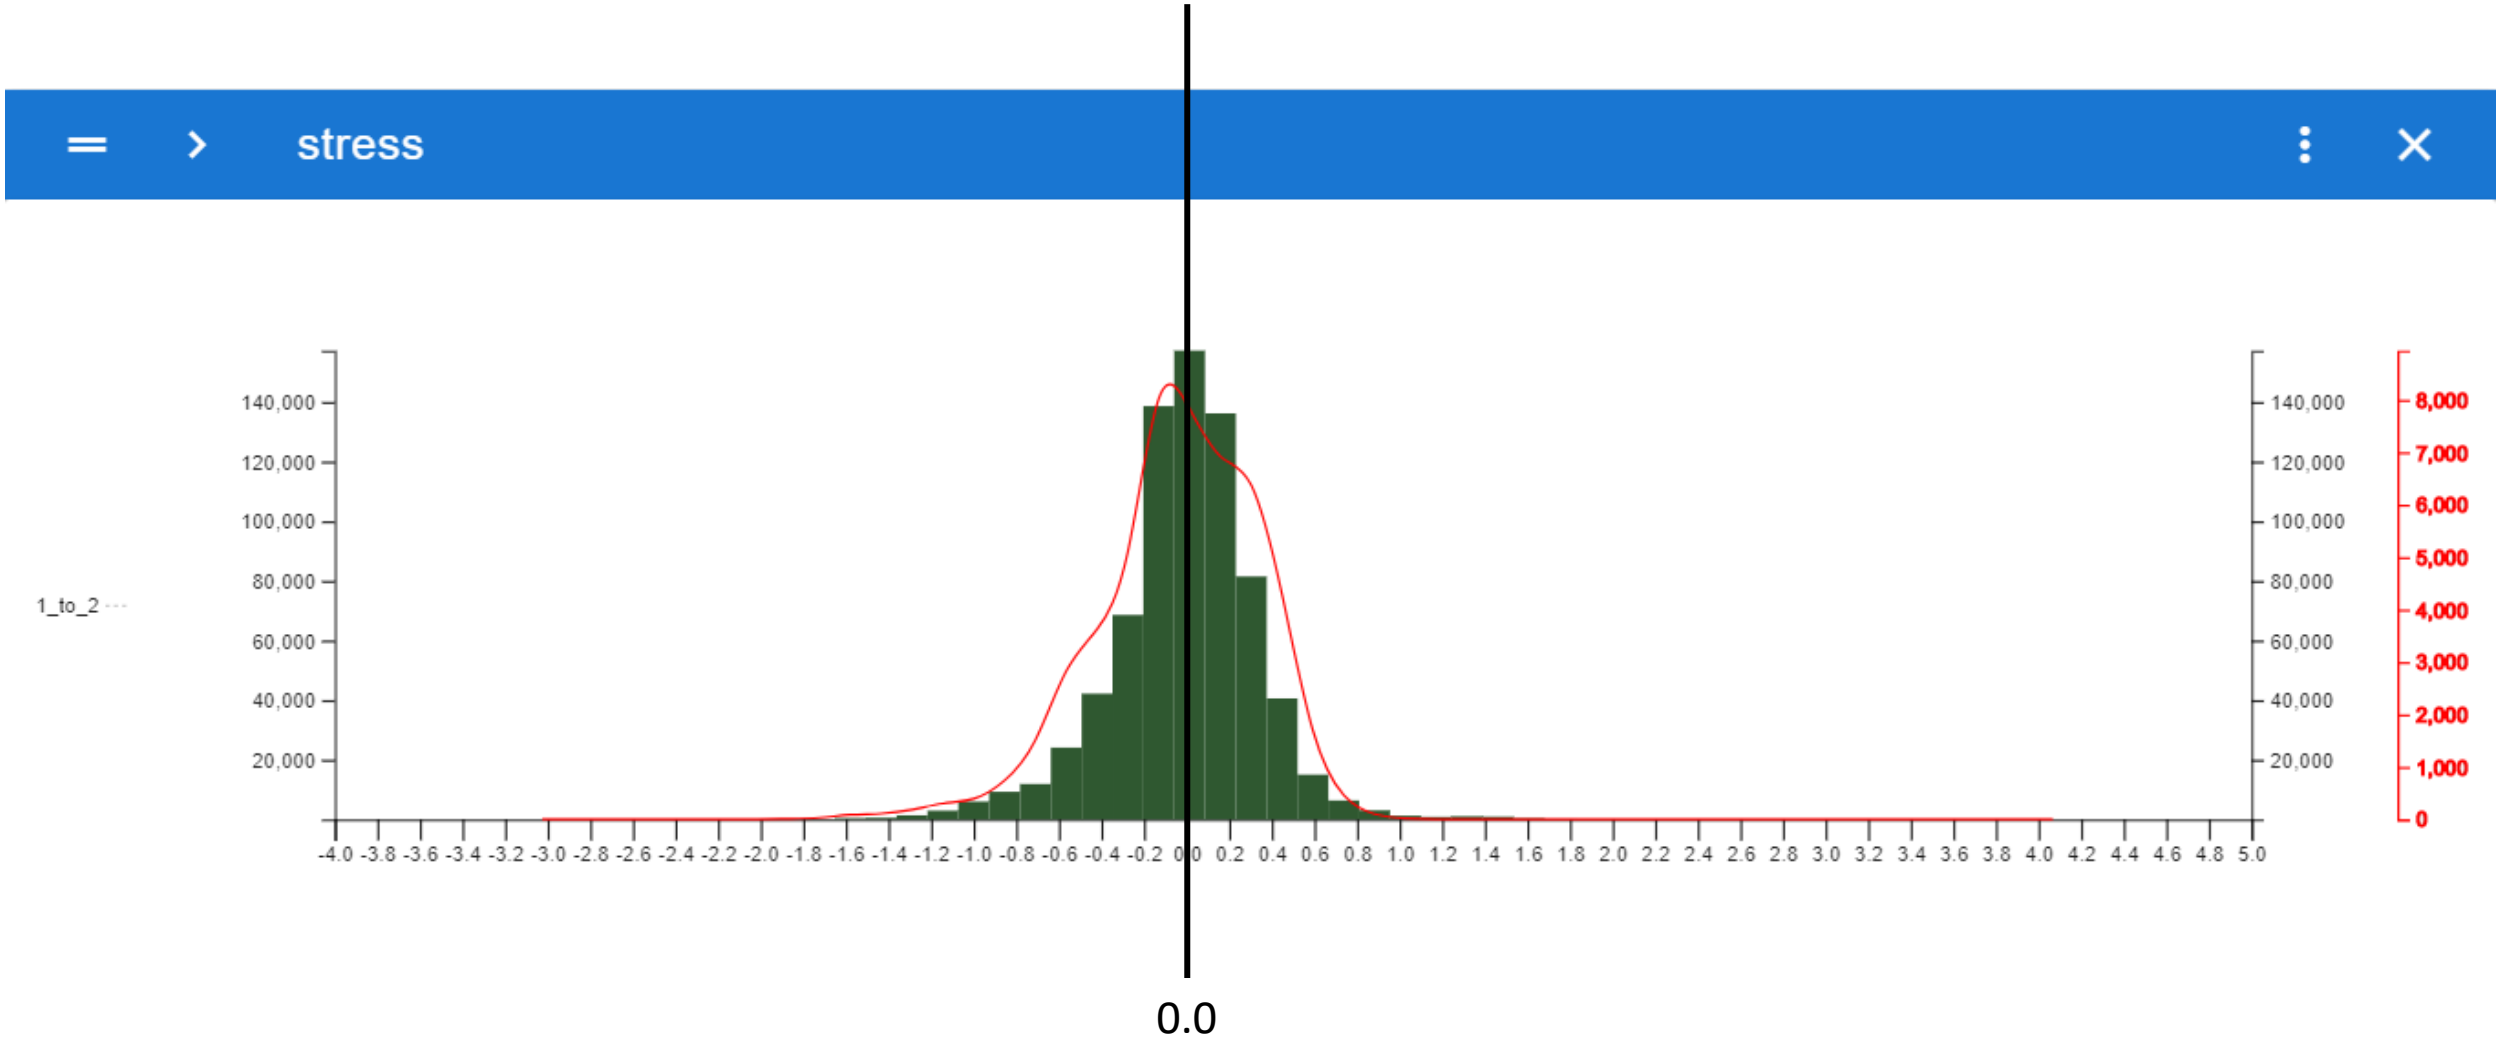

A real example:
Word Vector Embeddings, 2 Coprora

Wikipedia (modern) vs EEBO (1470-1700)

GloVE (embedding algorithm)

300 dims (key parameter of algorithm)

5000 most common words (can't show all words)

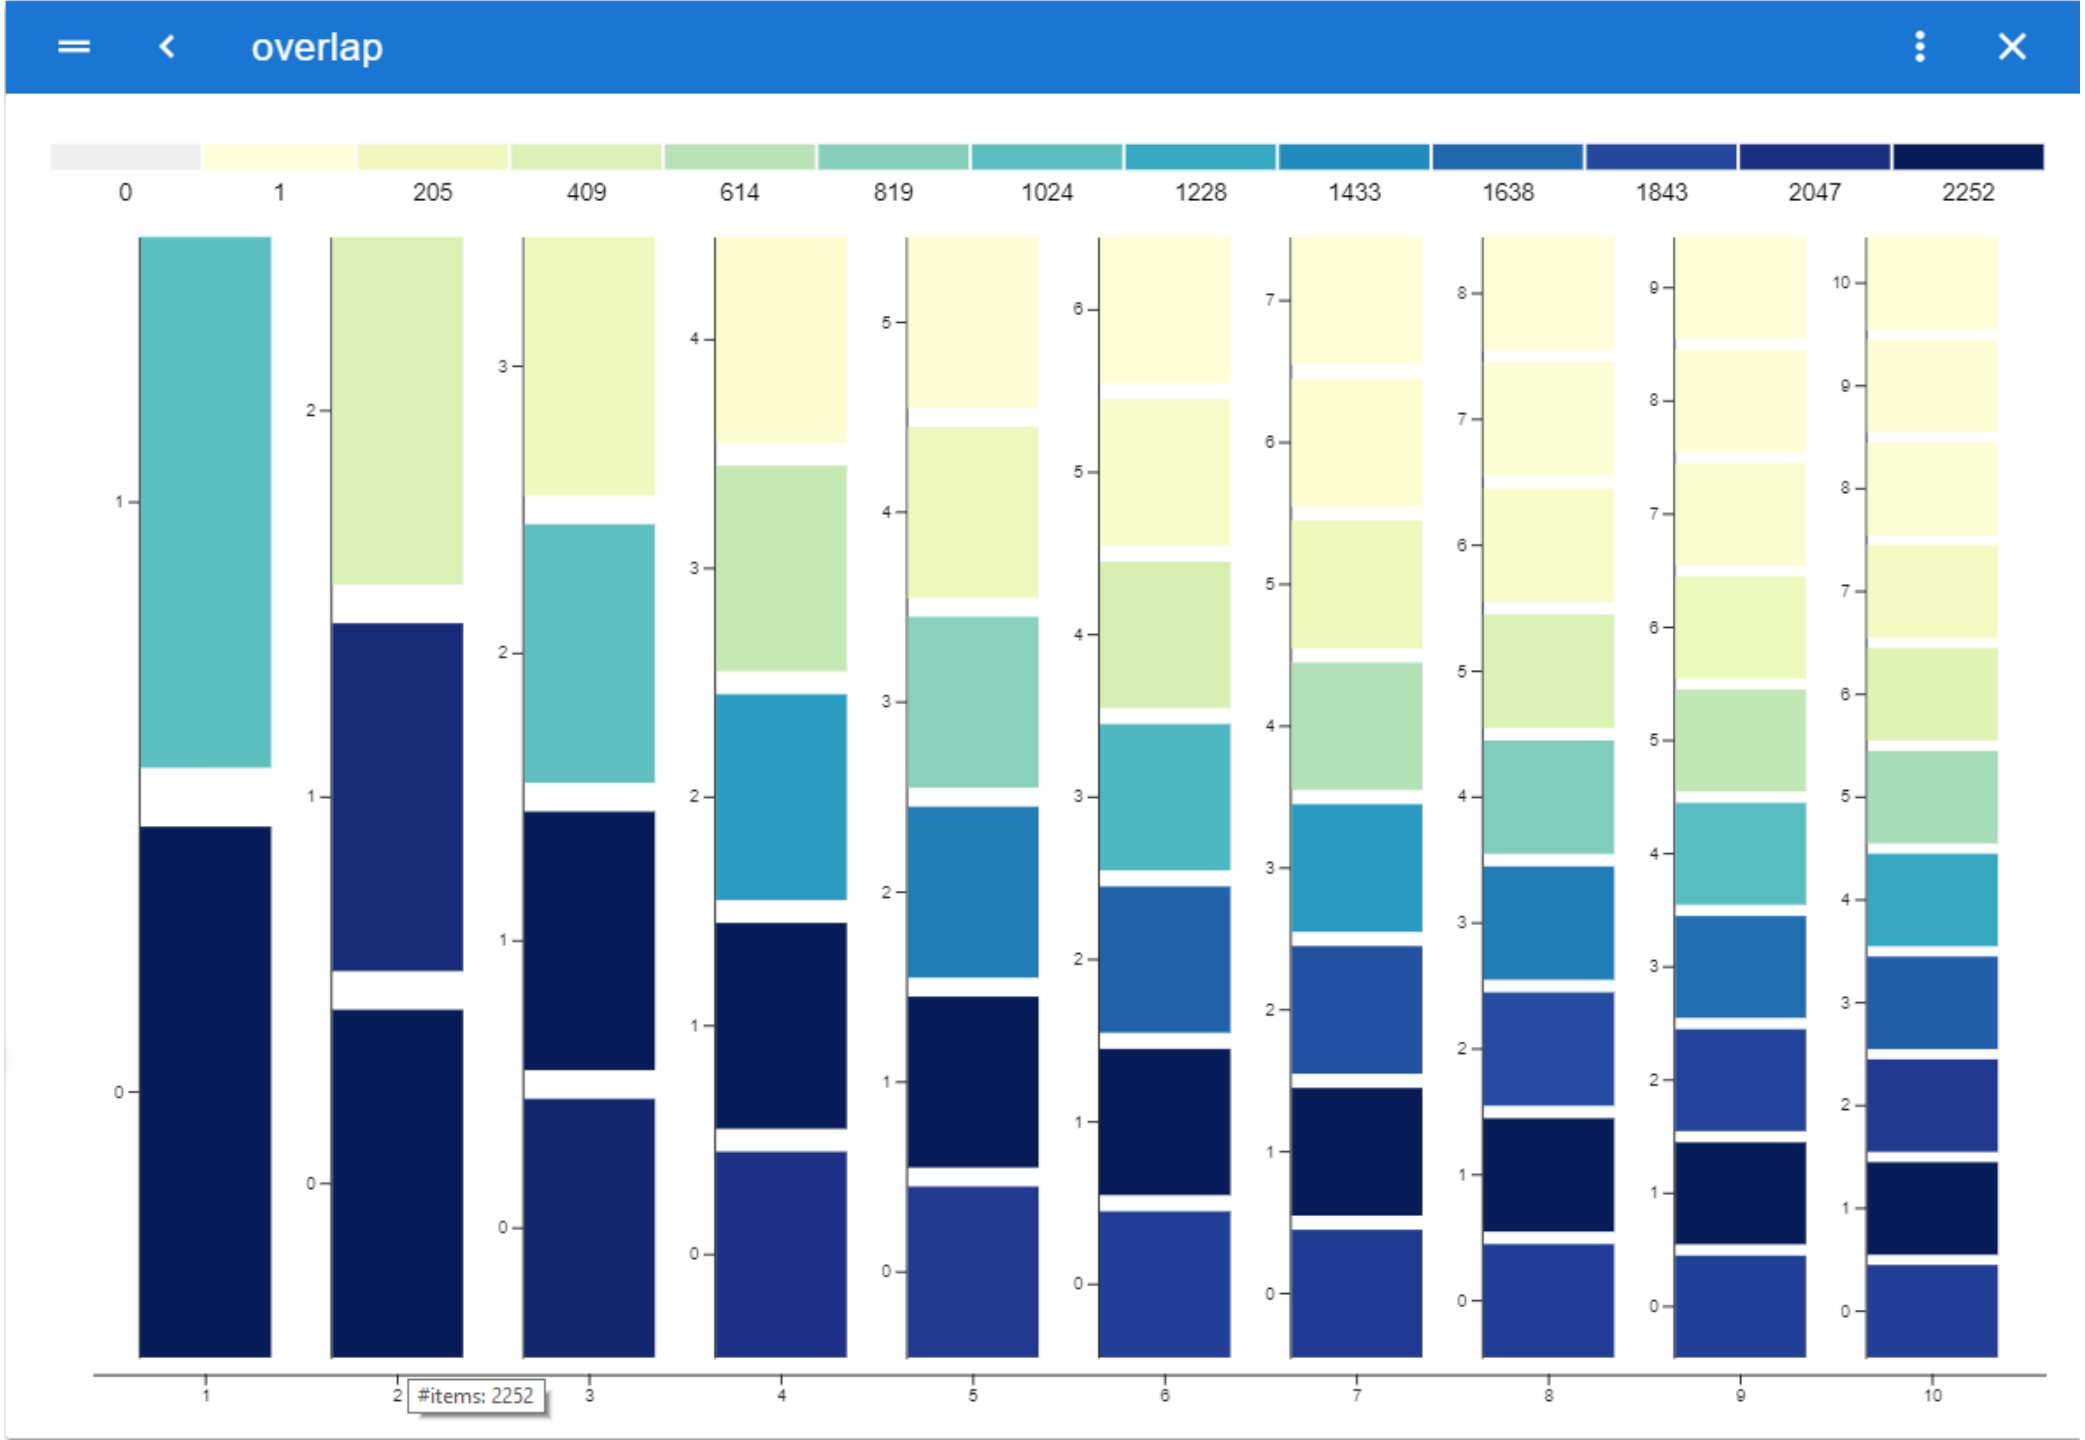
overlap
0
1
205
409
614
819
1024
1228
1433
1638
1843
2047
2252
#items: 2252

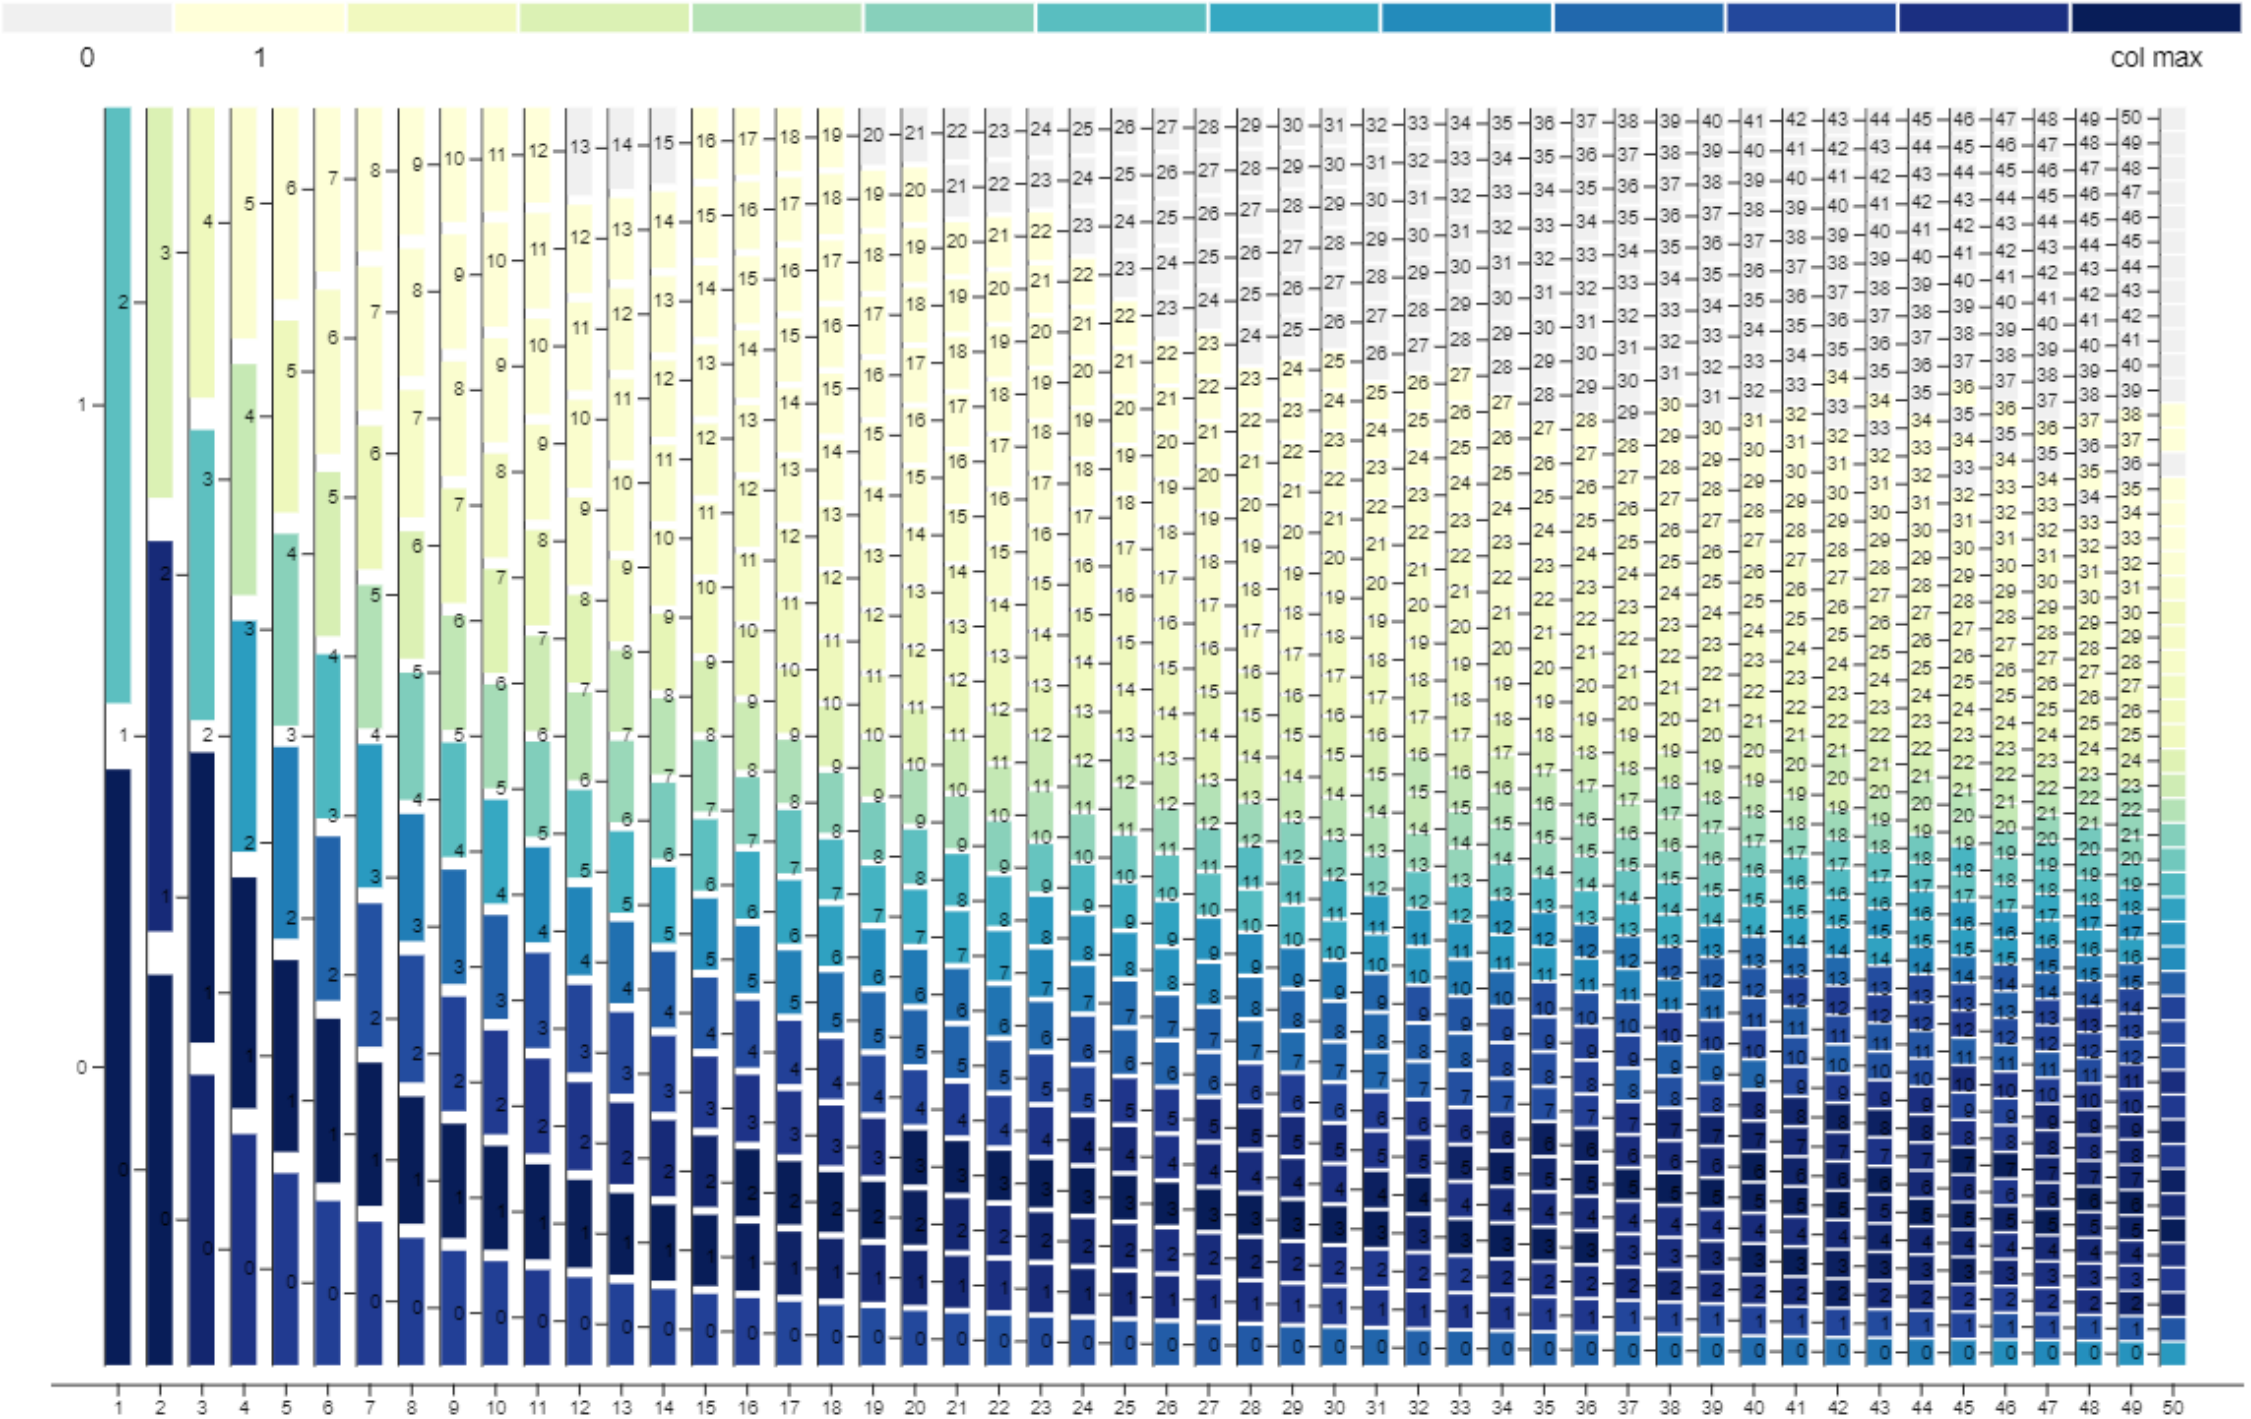
0
1
col max

Select groups with lots of overlap

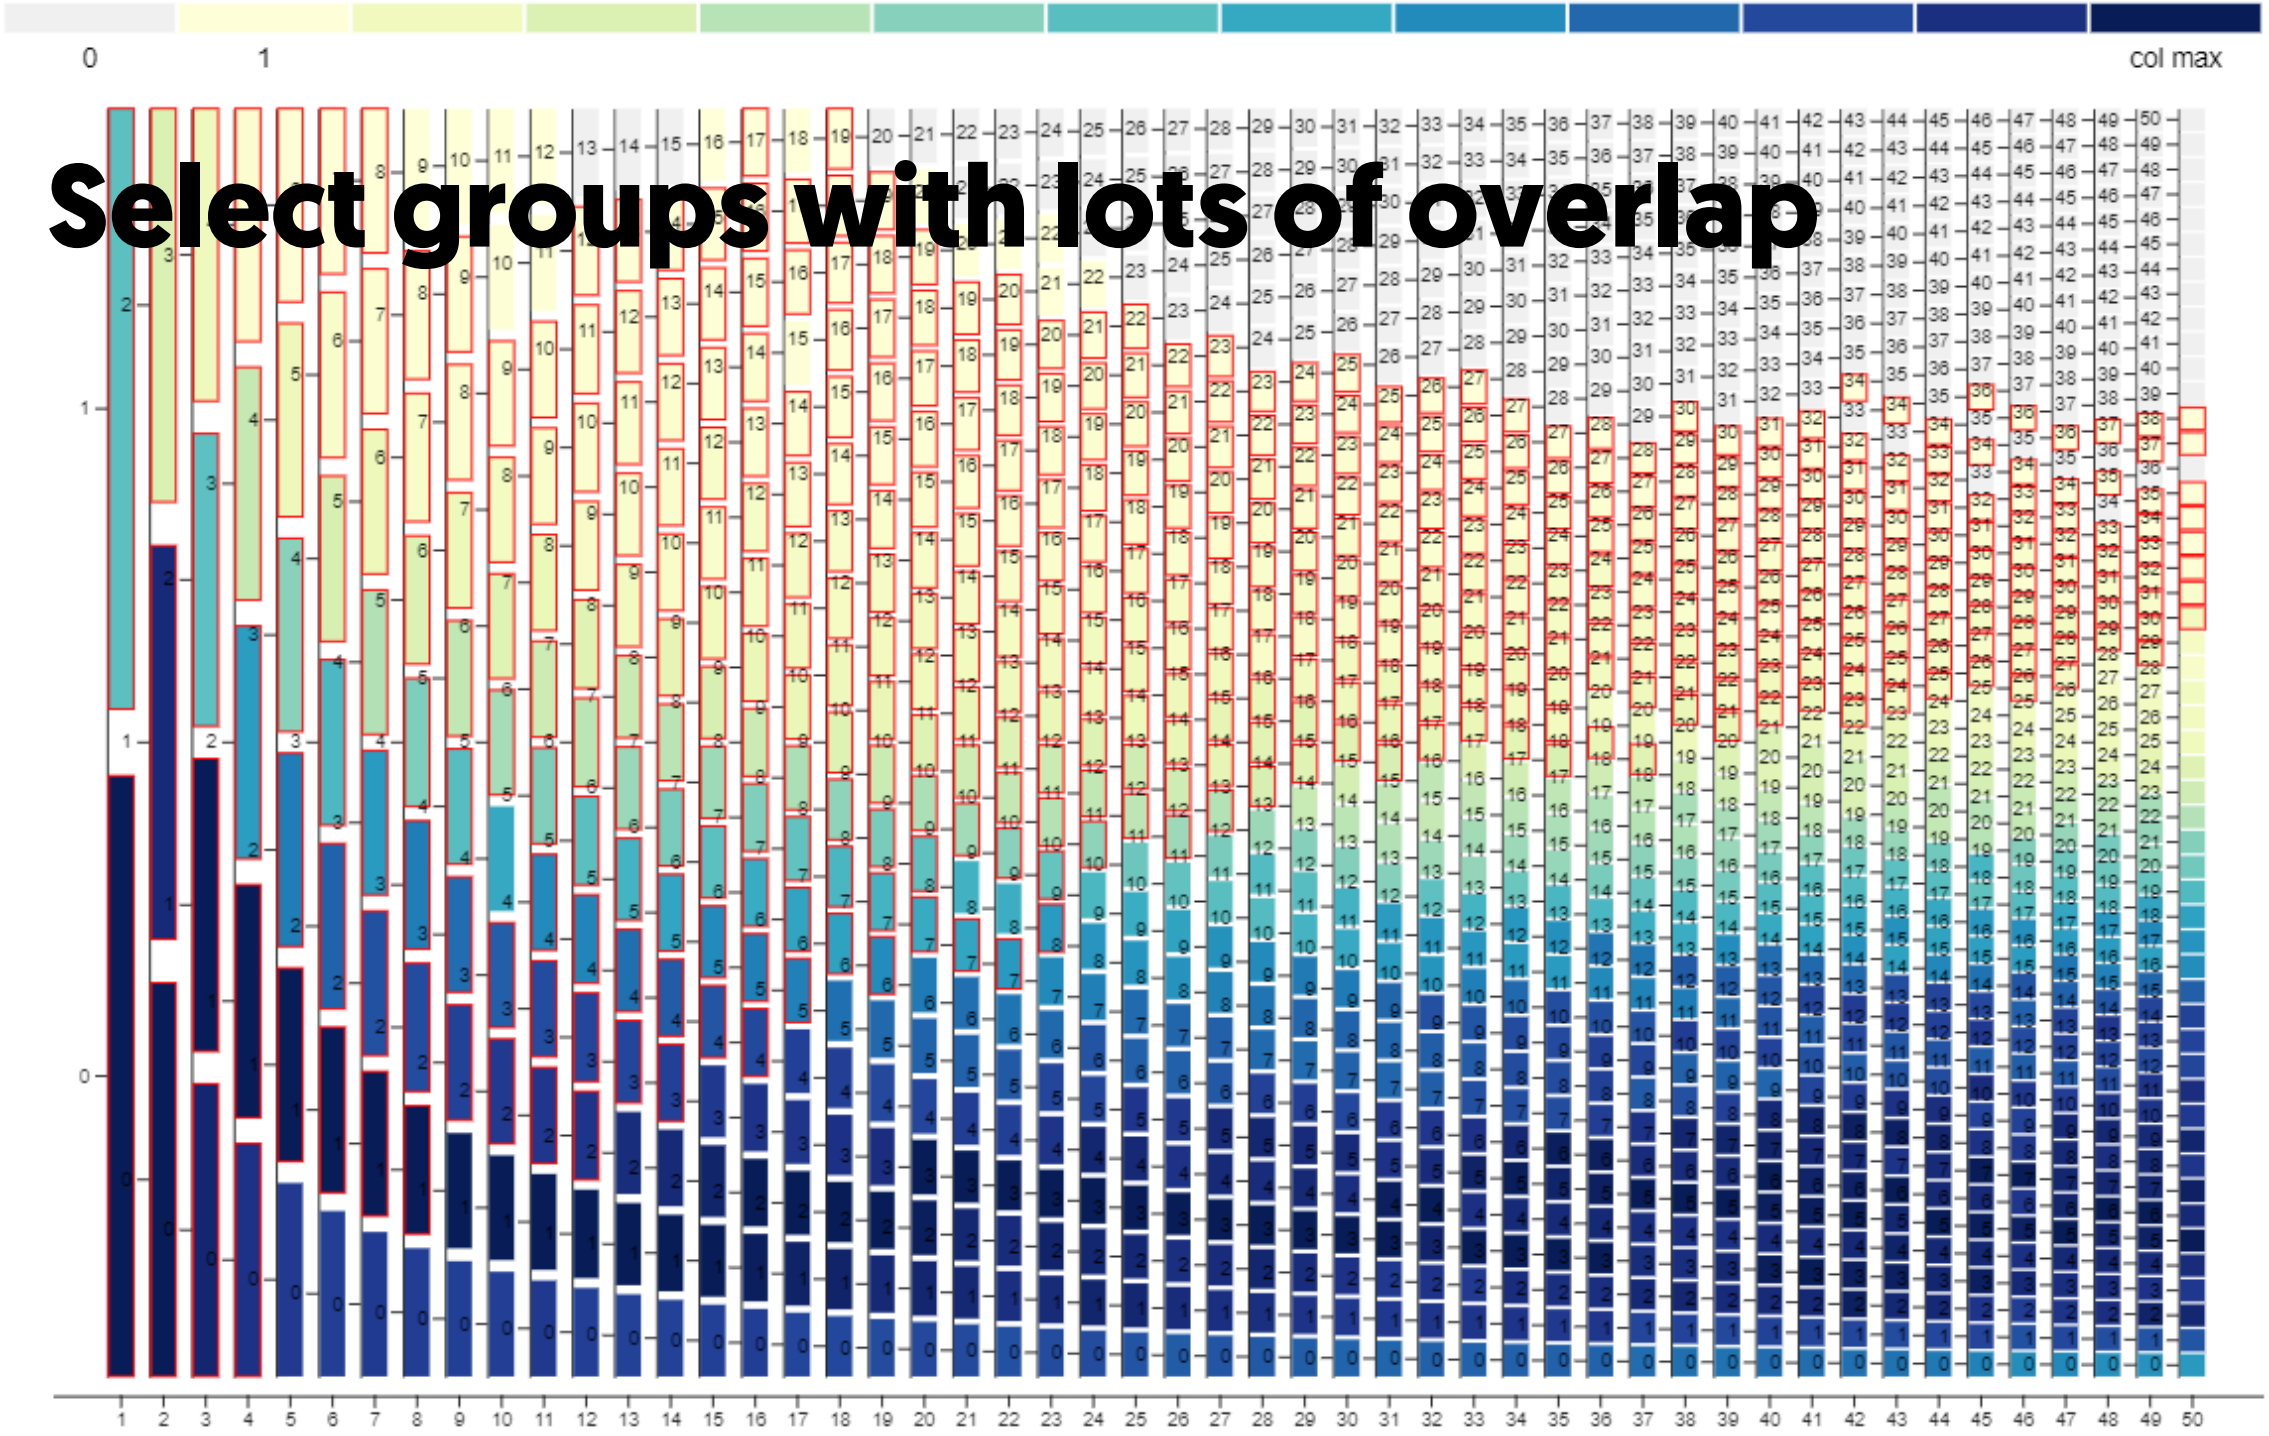

Sort by most frequent words (from 41) numbers, days, proper names, ...

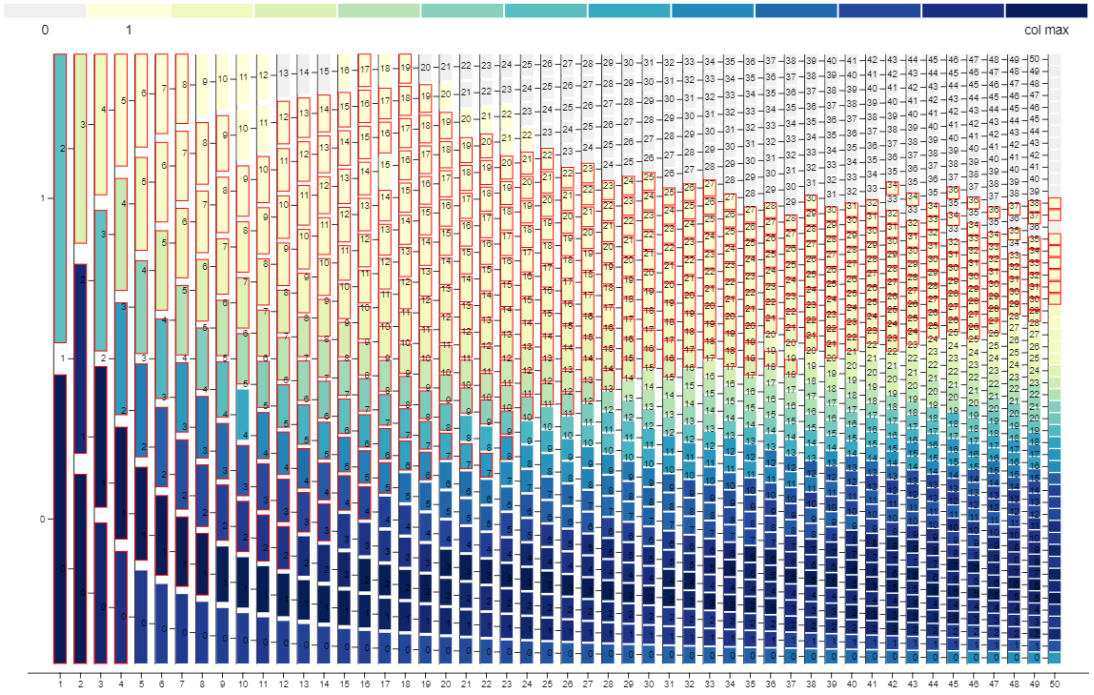

selected	id	freq 1 ↓
☐	three	1645302
☐	five	684045
☐	six	517393
☐	wife	403581
☐	seven	341921
☐	daughter	338759
☐	brother	323089
☐	ten	295505
☐	sister	192456
☐	younger	129124
☐	wilson	112785
☐	sunday	104629
☐	davis	102854
☐	twelve	97969
☐	hundred	96069
☐	norway	95728
☐	nor	79225
☐	saturday	77023
☐	twenty	76052
☐	anne	74218
☐	daughters	72731
☐	uncle	68453
☐	walker	67263
☐	margaret	66402
☐	harris	64530

Select no overlap in top 50

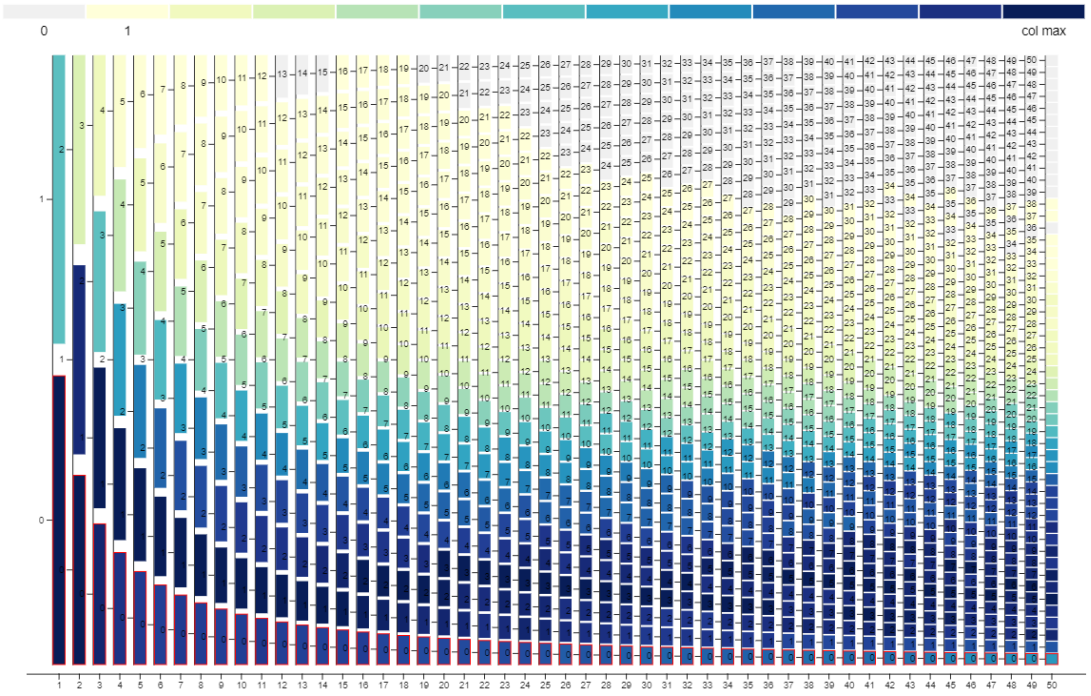

selected	id	freq 1	freq 2 ↓
☐	cast	178596	331481
☐	vice	112847	75918
☐	bid	47908	70815
☑	cup	517385	69281
☐	abuse	50778	66484
☐	native	278967	48654
☐	austin	62922	37401
☐	ep	79033	36968
☐	venture	51796	35277
☐	ed	55568	30072
☐	car	296375	29654
☐	maine	54443	28673
☐	key	202014	23934
☐	construction	368816	18778
☐	folk	82224	15644
☐	fa	48580	15324
☐	pat	35395	15062
☐	quantum	35699	14628

Cup?

neighborlist

embedding 1 SELECT ALL Search

rank	dist	id	freq 1	freq 2
0	0.0	cup	517385	69281
1	0.28	champions	129125	7498
2	0.33	winning	256468	7353
3	0.34	win	402616	41611
4	0.34	winners	83205	250

1-5 of 11

Rows per page: 5

embedding 2 SELECT ALL Search

rank	dist	id	freq 1	freq 2
0	0.0	cup	517385	69281
1	0.33	drink	34292	196687
2	0.34	wine	72208	229395
3	0.45	bowl	103446	5999
4	0.47	drinking	39741	41776

1-5 of 11

Rows per page: 5

Another example:
Build topic models with abstracts?

Vispub data (871 papers)

Abstracts vs. Full Text

NMF, 10 topics

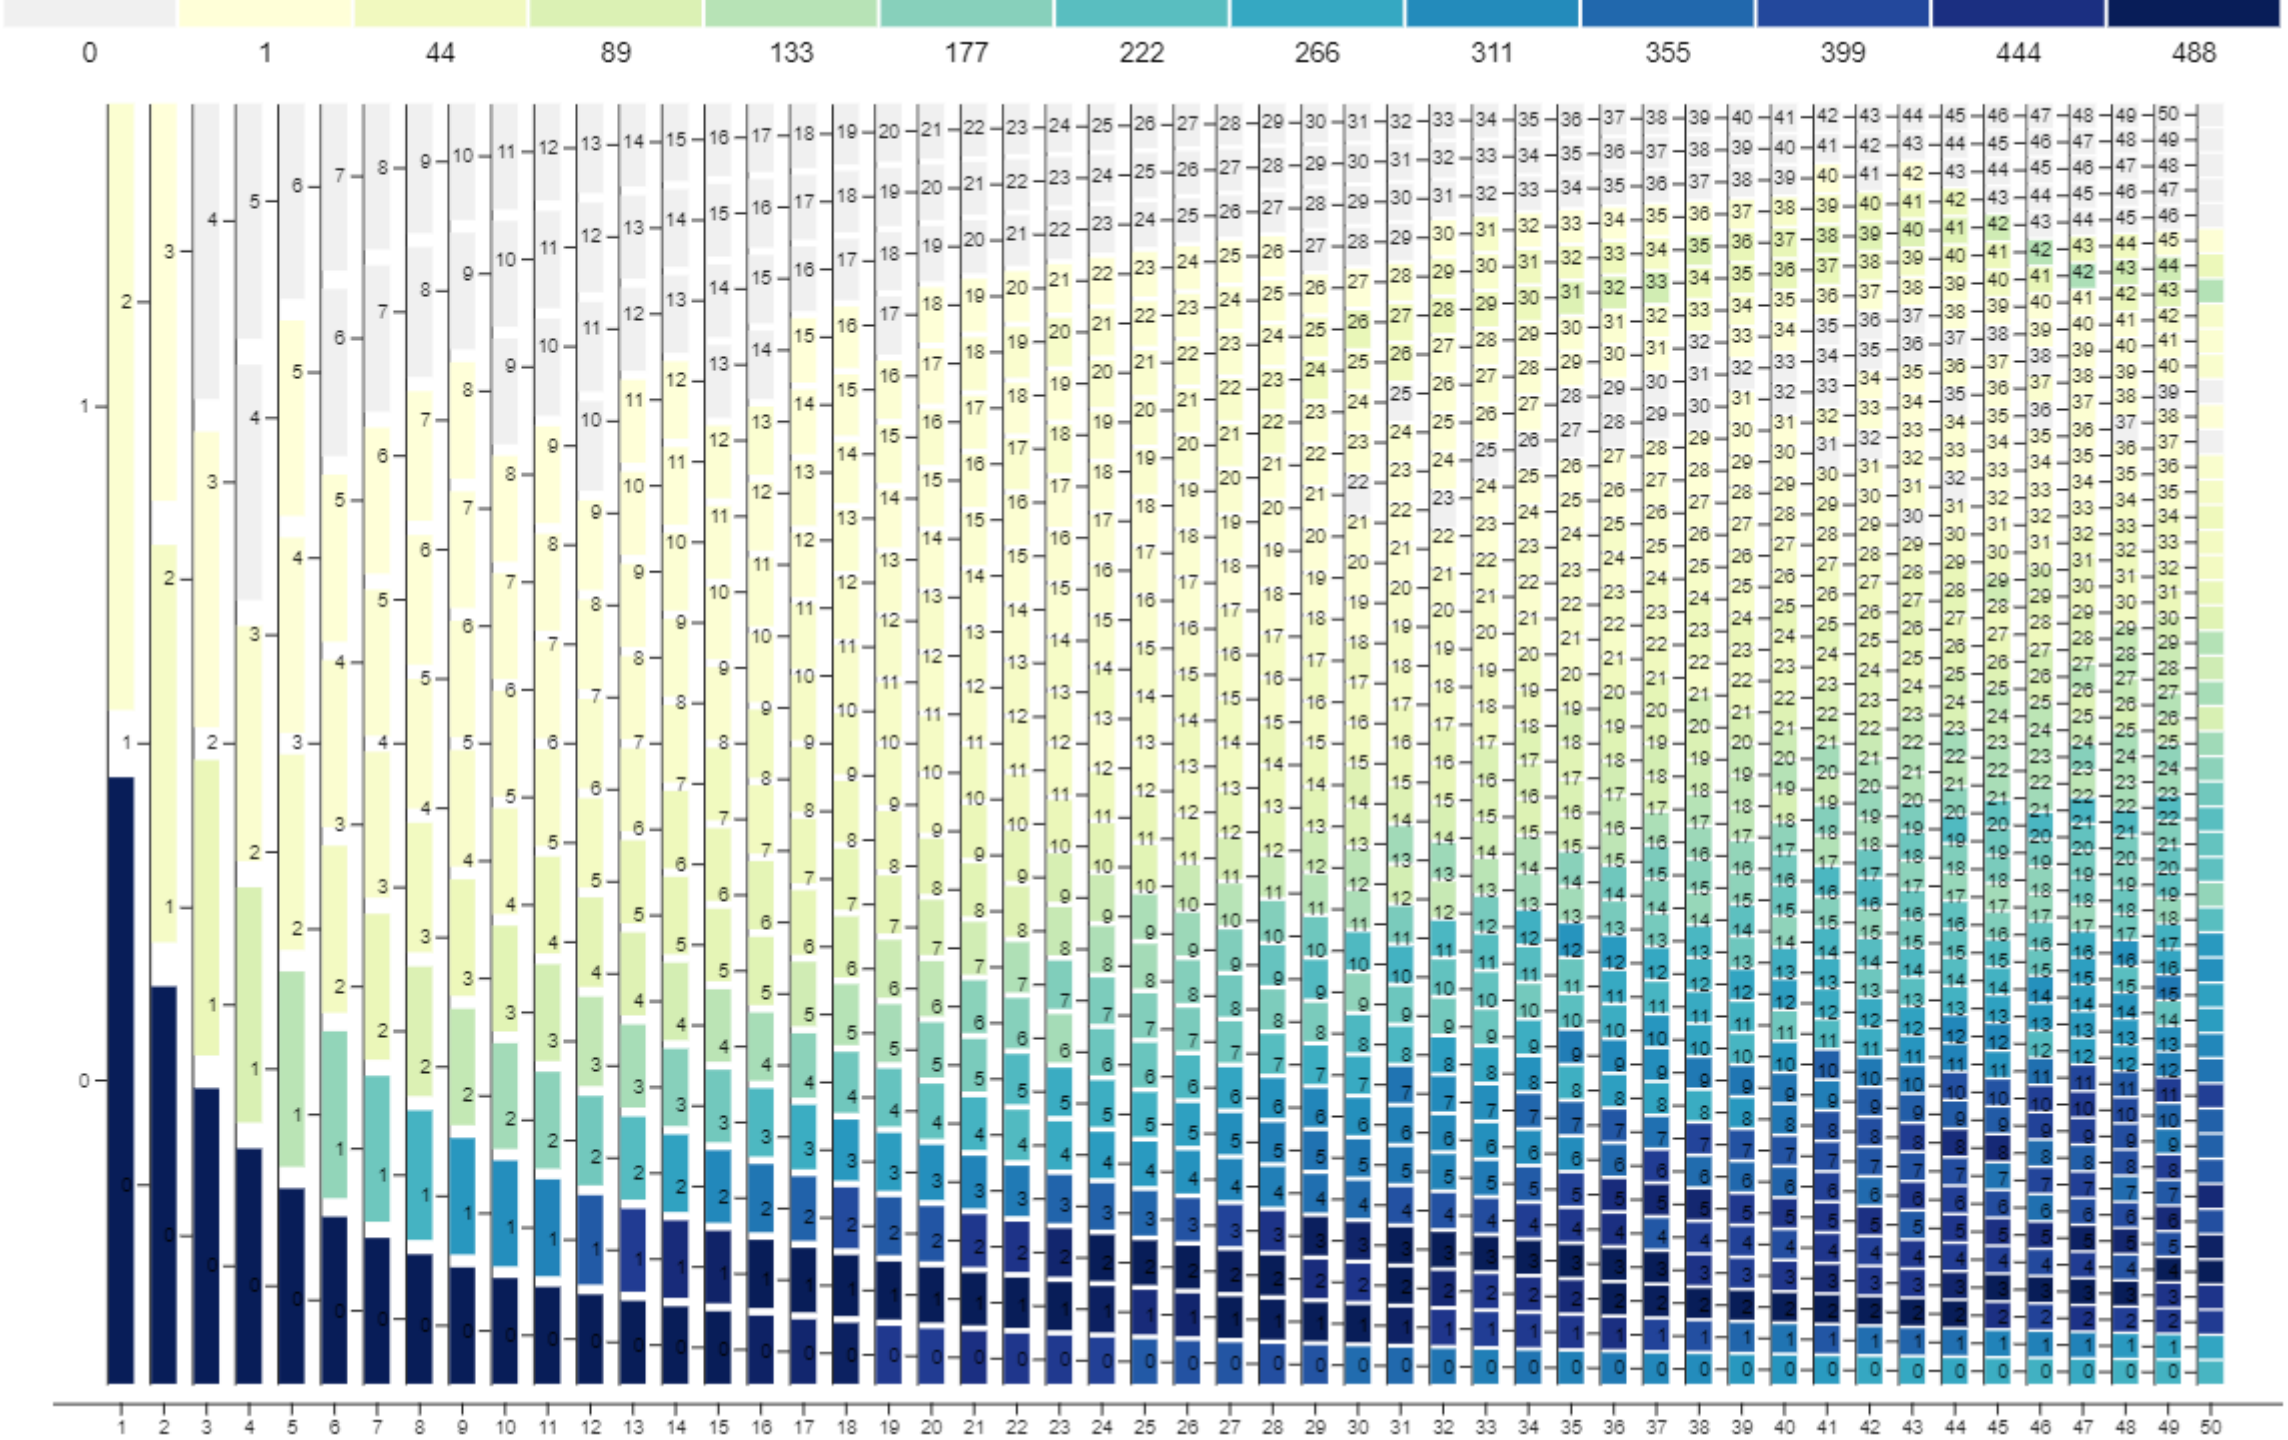
0
1
44
89
133
177
222
266
311
355
399
444
488

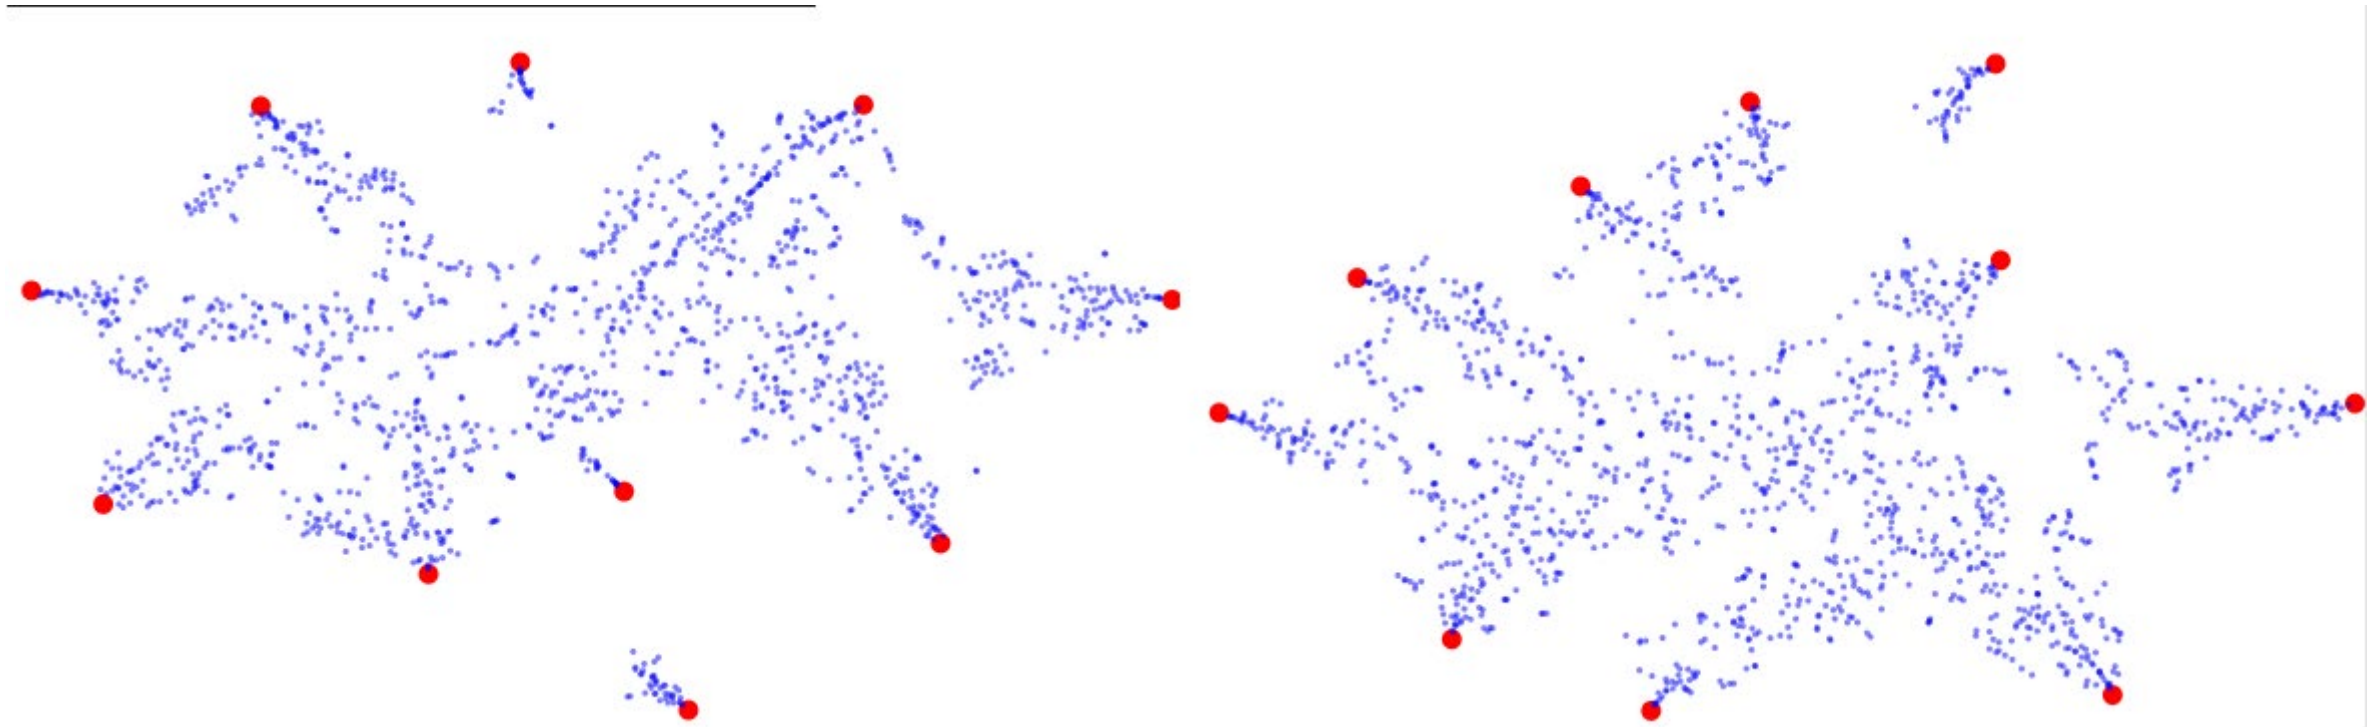

Pick a cluster – it does correspond

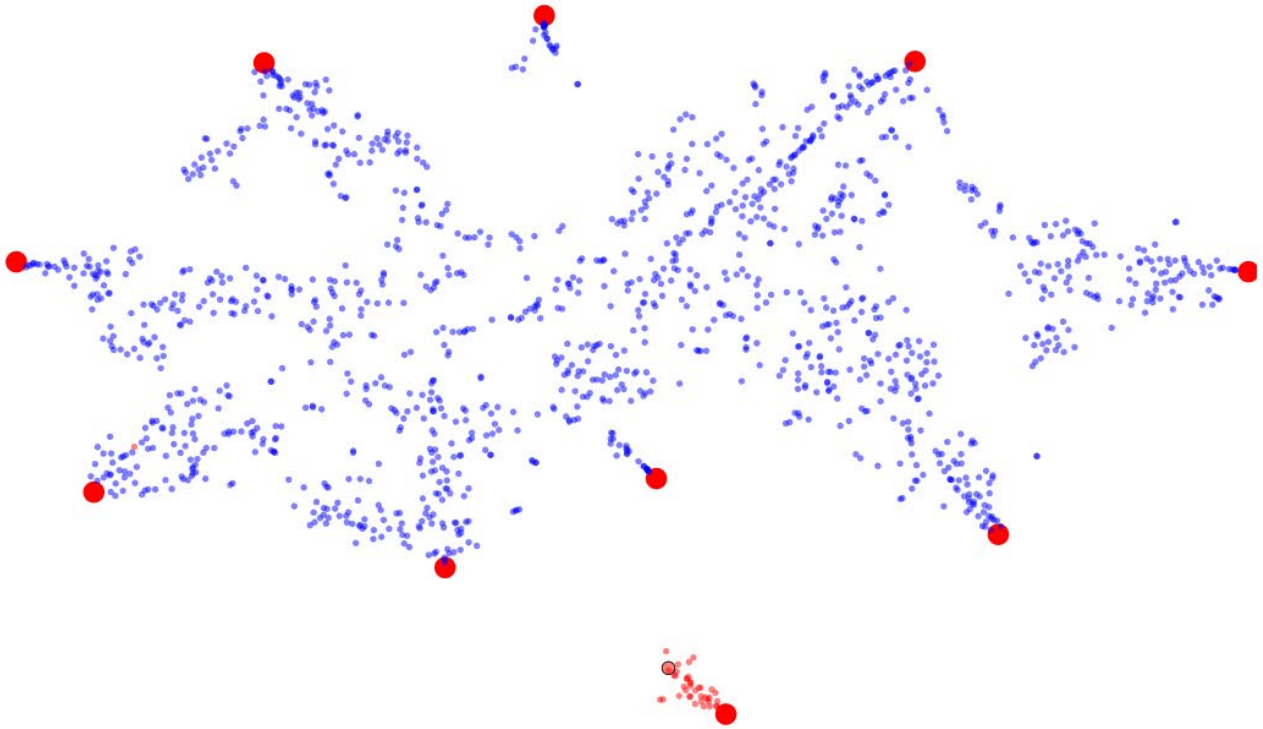

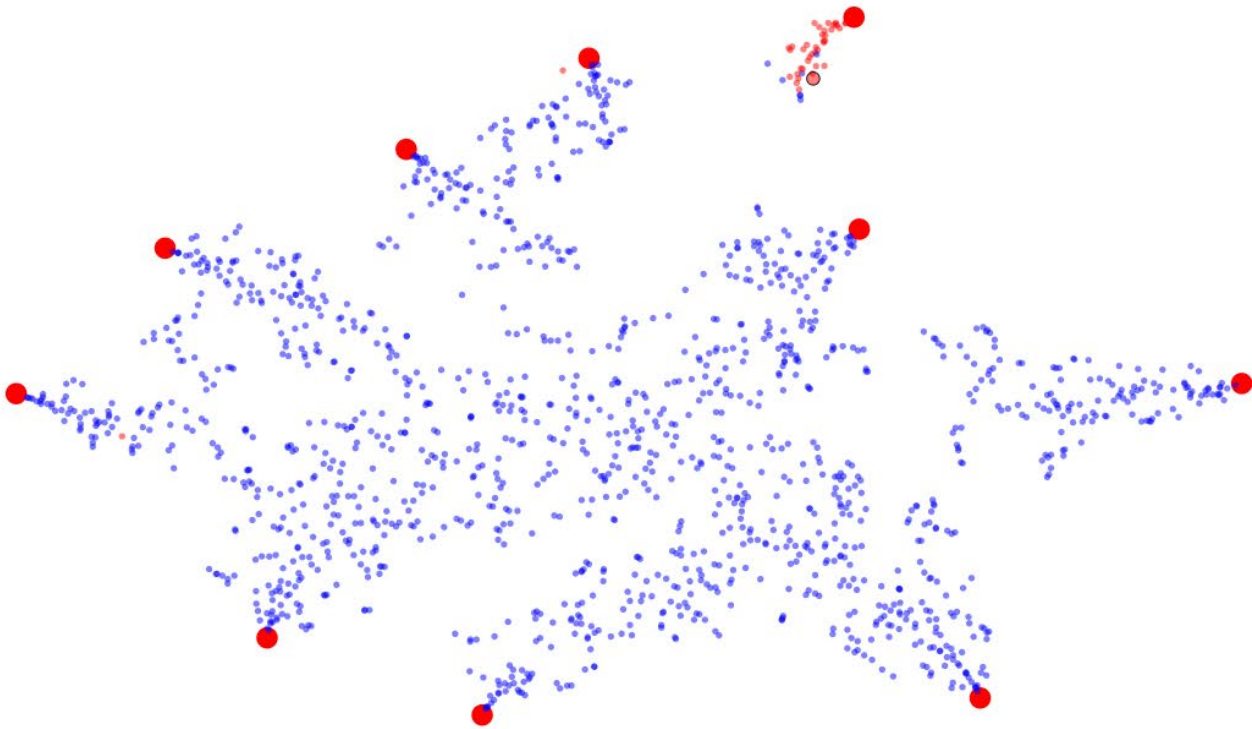

Most of these have lots of overlap

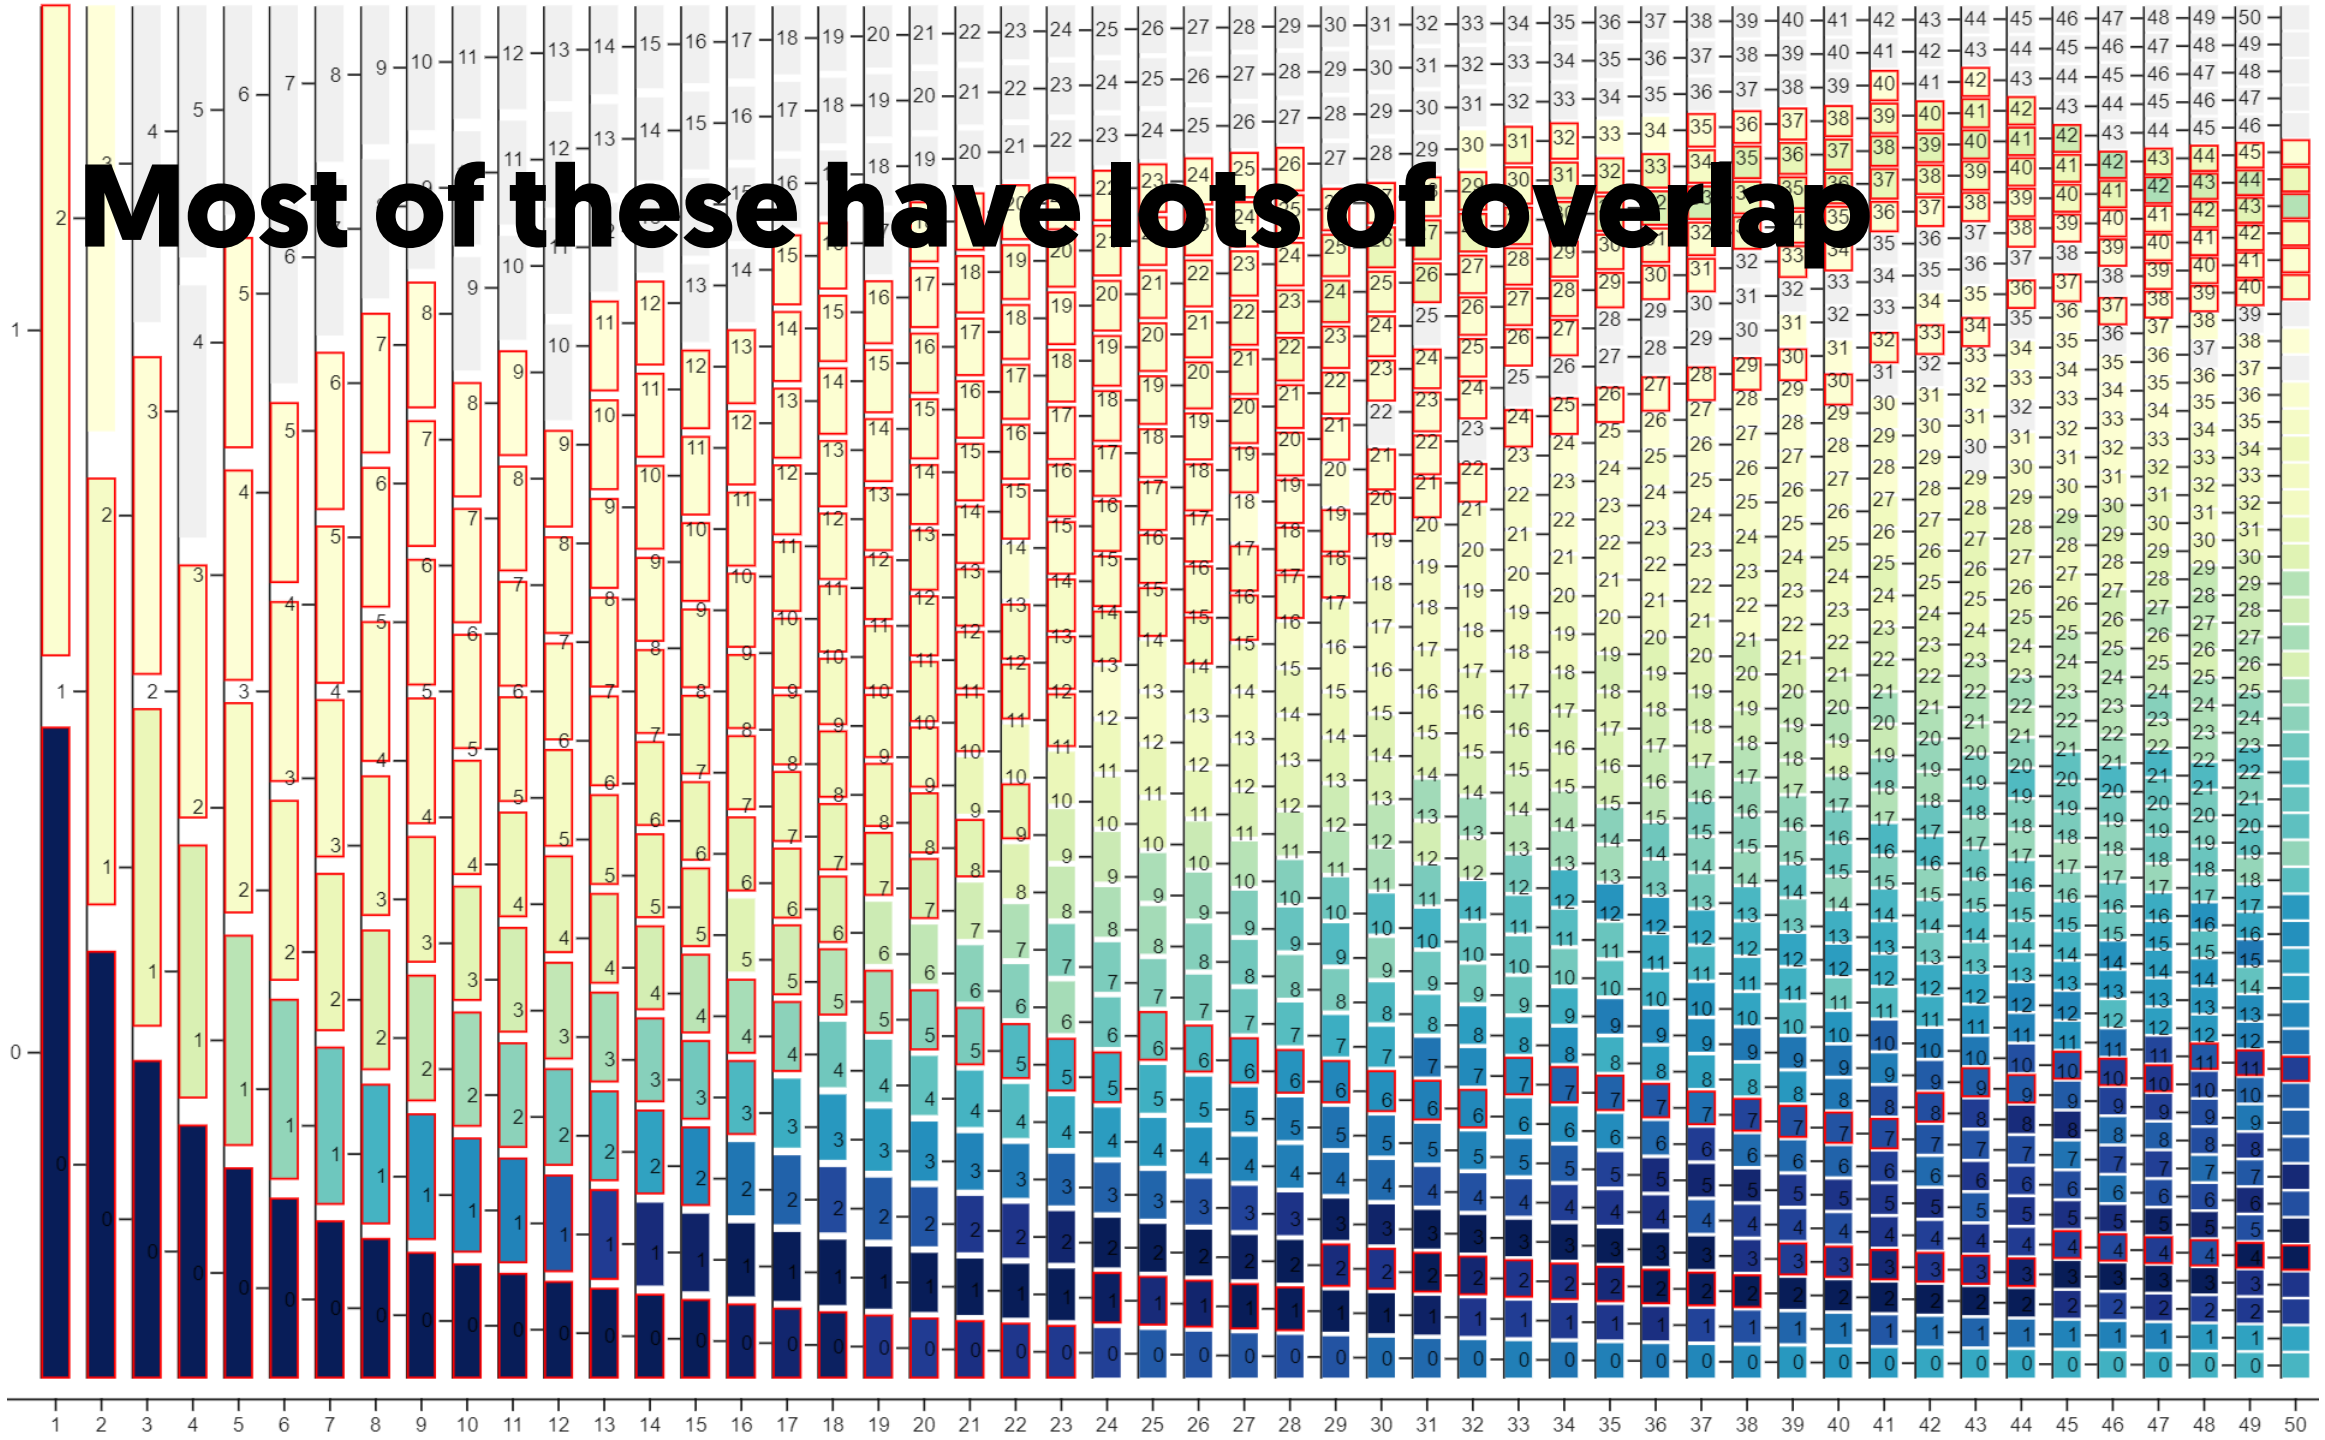

Cluster doesn’t match
left: abstracts about GPU/rendering
right: papers more about domain

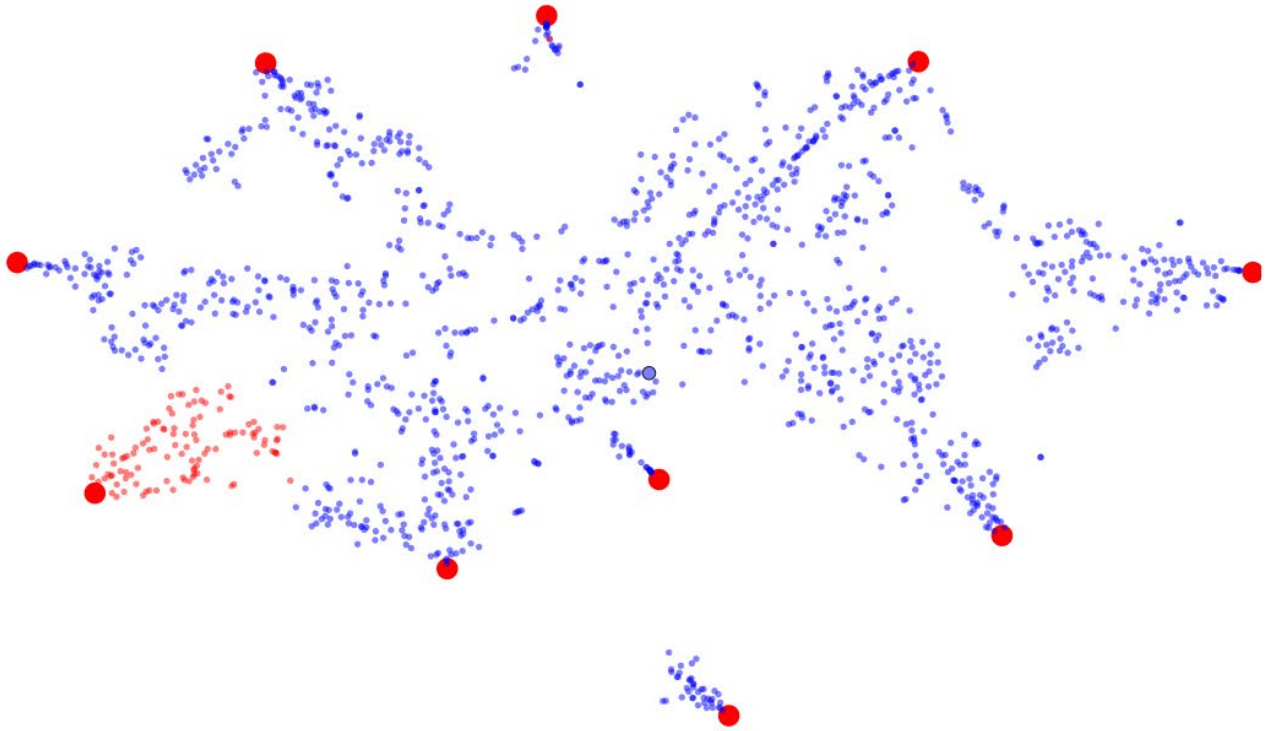

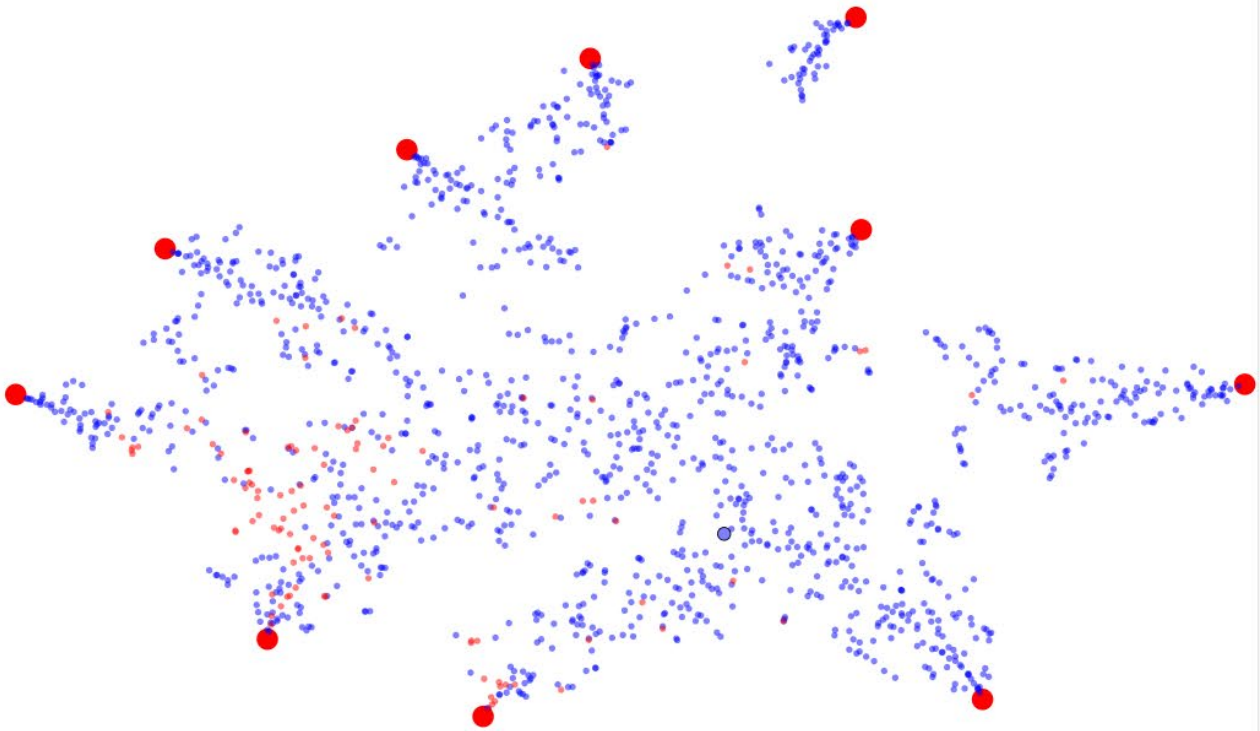

Hypothesis:
Abstracts are different than papers

Limitations

Binary comparison – less good for parameter tuning

need new designs?

Many complex views that need to be combined

address usability – through task-driven pre-arrangements?

Scalability of implementation and visual designs

Validation on real problems with real users

Summary

Embeddings for Text Analysis

If we can design interpretation tools

Compare document and word embeddings to interpret them

Specialized tools for comparison of embeddings

Task-centric design process

Comparison as Analysis Approach

Because we have a design process

An approach to thinking about design for comparison

Examples using text analysis with embedding comparison

4 considerations of comparison

Thanks!

To you for listening

To my students and collaborators

(too many to list)

To our sponsors

NSF, NIH, DARPA, Mellon Foundation

Michael Gleicher

http://pages.cs.wisc.edu/~gleicher