Deconstructing Data Science

David Bamman, UC Berkeley

Info 290
Lecture 20: Distance models (clustering)

Apr 6, 2017

Clustering

- Clustering (and unsupervised learning more generally) finds *structure* in data, using just X

X = a set of skyscrapers

Flat Clustering

- Partitions the data into a set of *K* clusters

K-means

Algorithm 1 K-means

Data: training data $x \in \mathbb{R}^F$
Given some distance function $d(x, x') \to \mathbb{R}$
Select k initial centers $\{\mu_1, \dots, \mu_k\}$
while not converged **do**
  for $i = 1$ to N **do**
    Assign x_i to $\arg\min_c d(x_i, \mu_c)$
  end for
  for $i = 1$ to K **do**
    $\mu_i = \frac{1}{D_i} \sum_{j=1}^{D_i} x_i$
  end for
end while

Visualizing K-Means Clustering

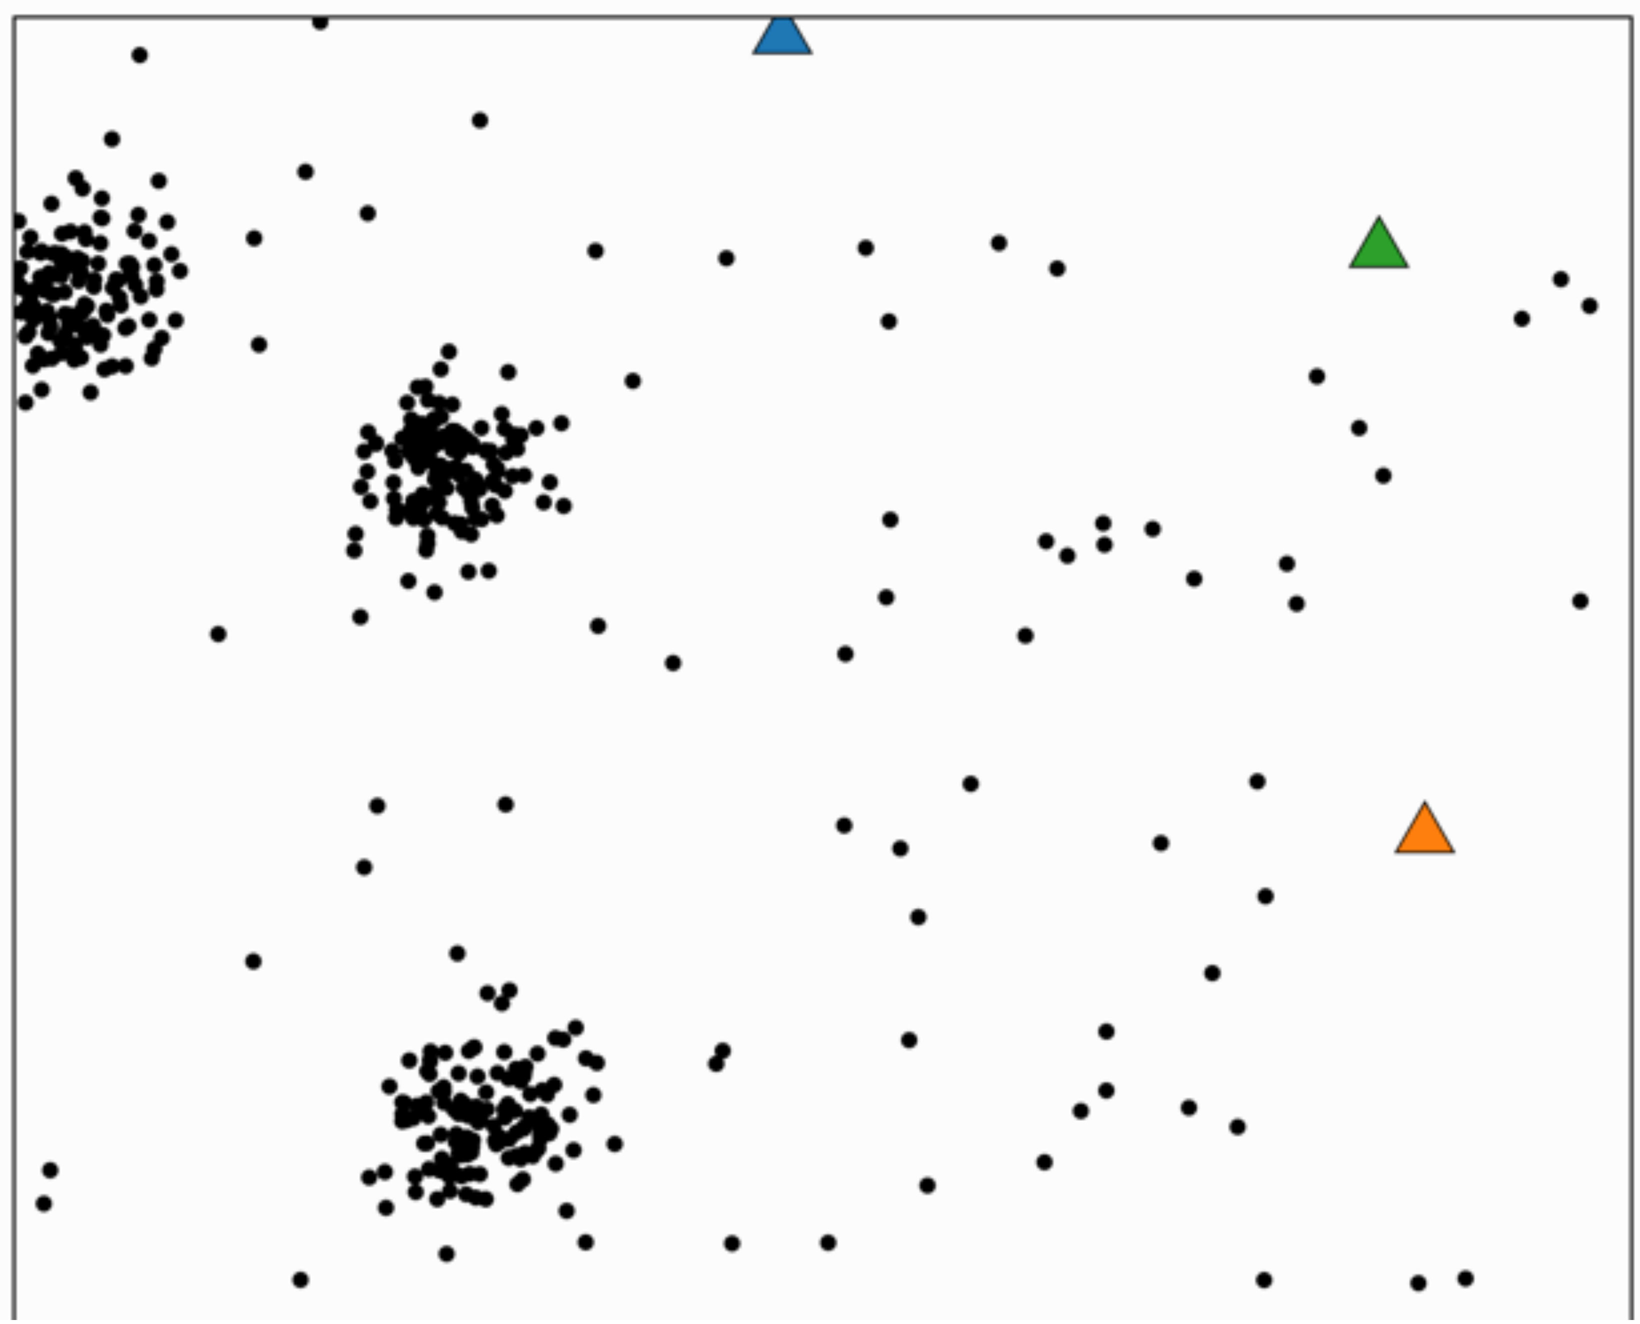

http://stanford.edu/class/ee103/visualizations/kmeans/kmeans.html

Problems

K-means

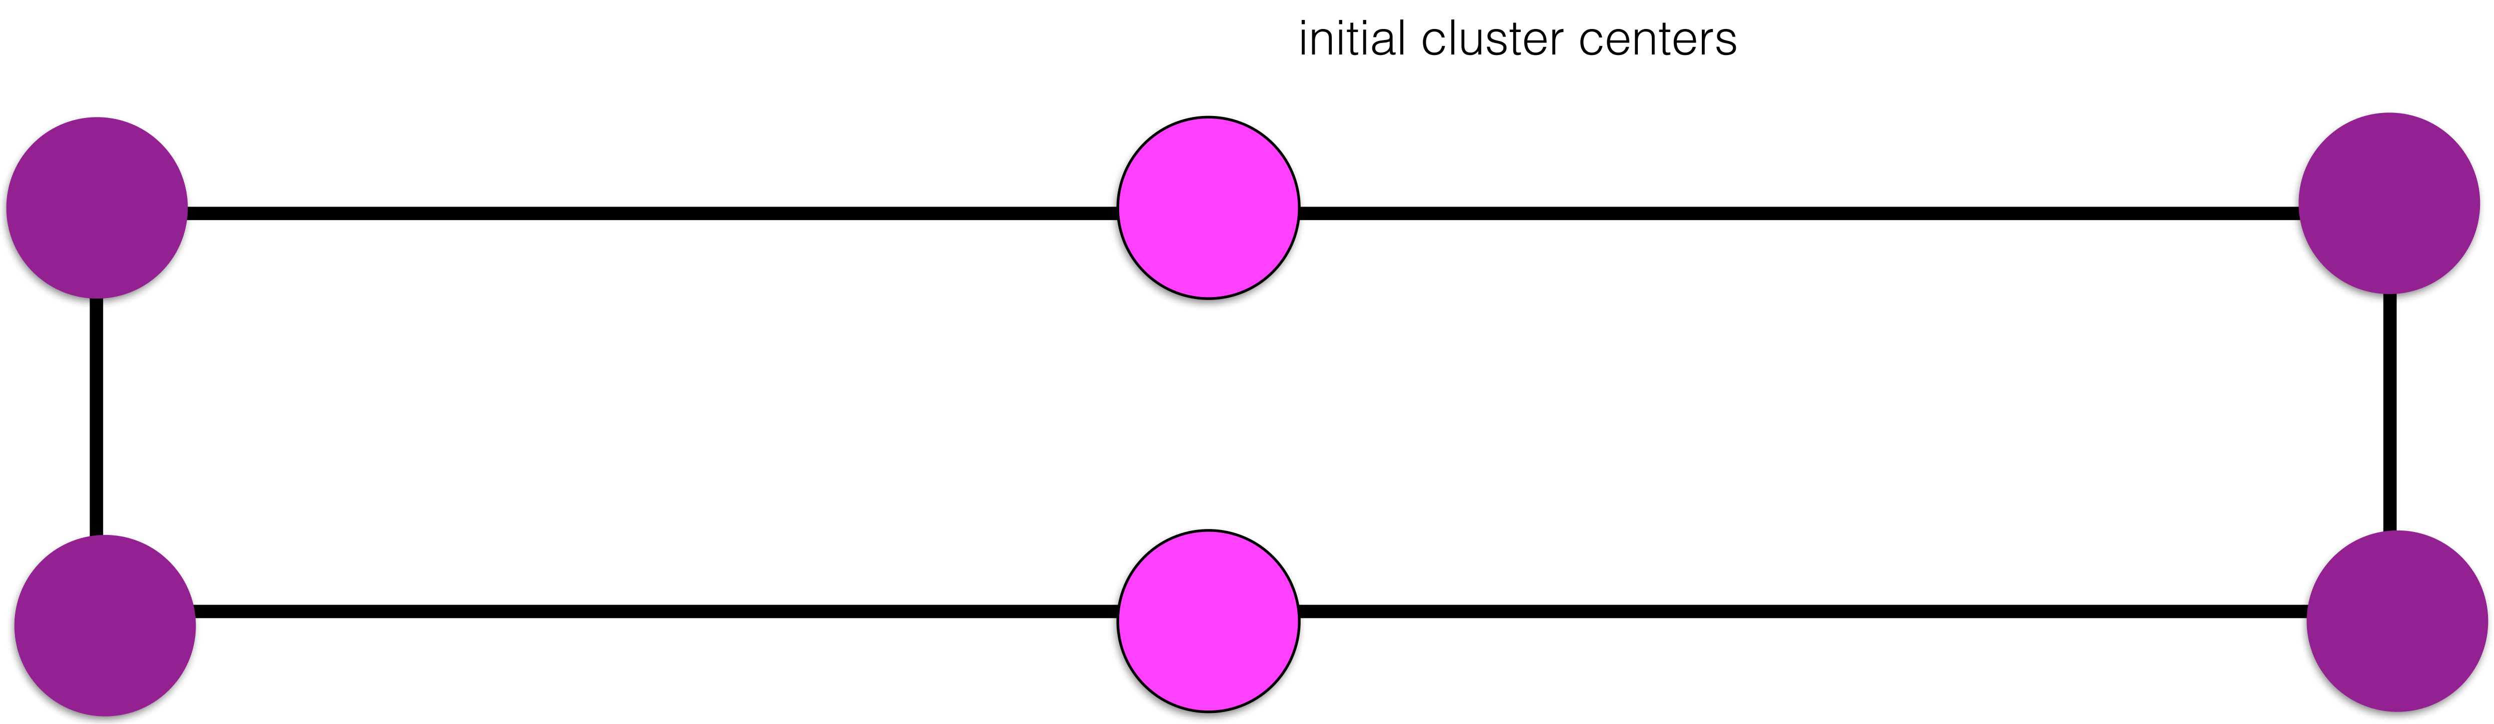

K-means++

- Improved initialization method for K-means:
    - Choose data point at random as first center
    - For all other data points x, calculate the distance D(x) between x and the nearest cluster center
    - Choose new data point x as next center, with probability proportional to $D(x)^2$
    - Repeat until *K* centers are selected

K-means++

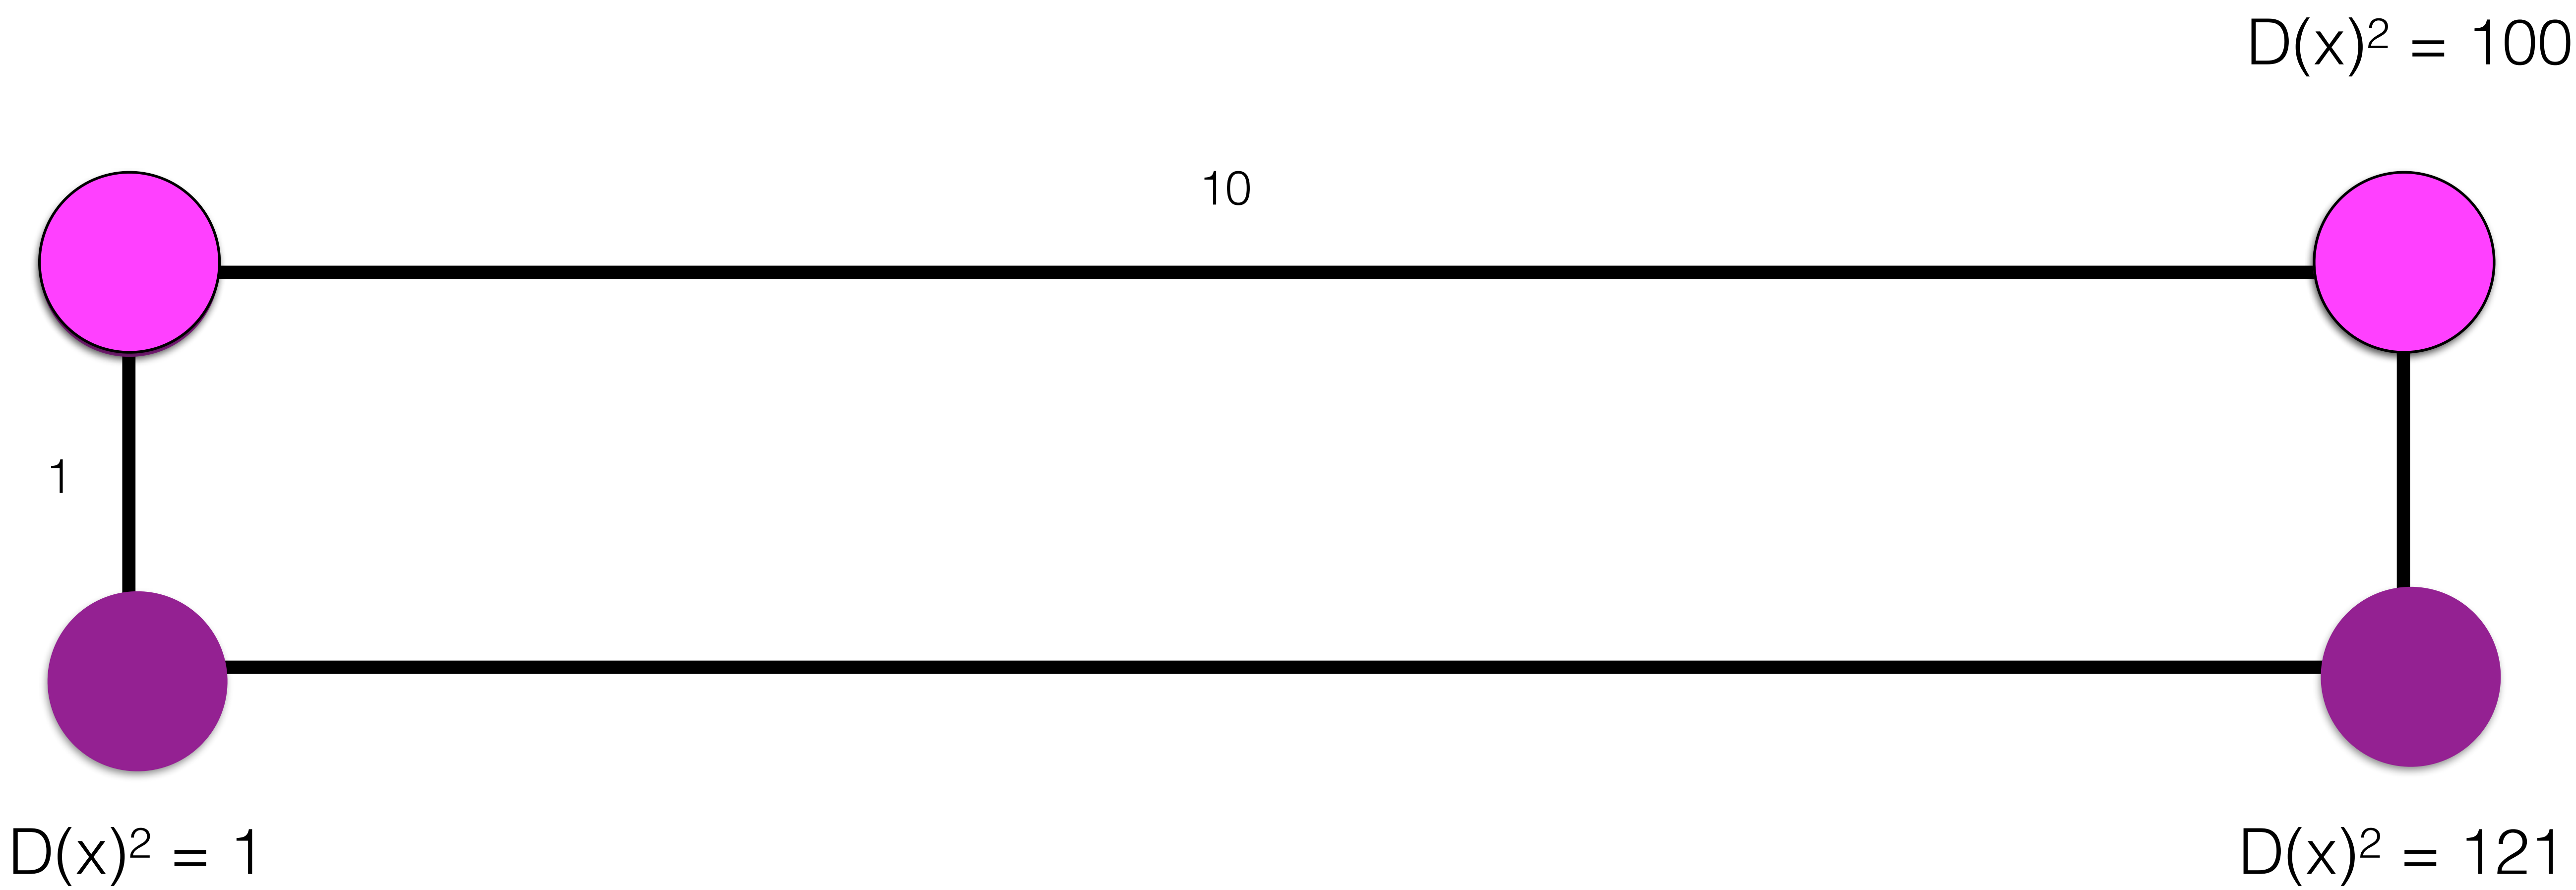

Choosing K

- how do we choose K?

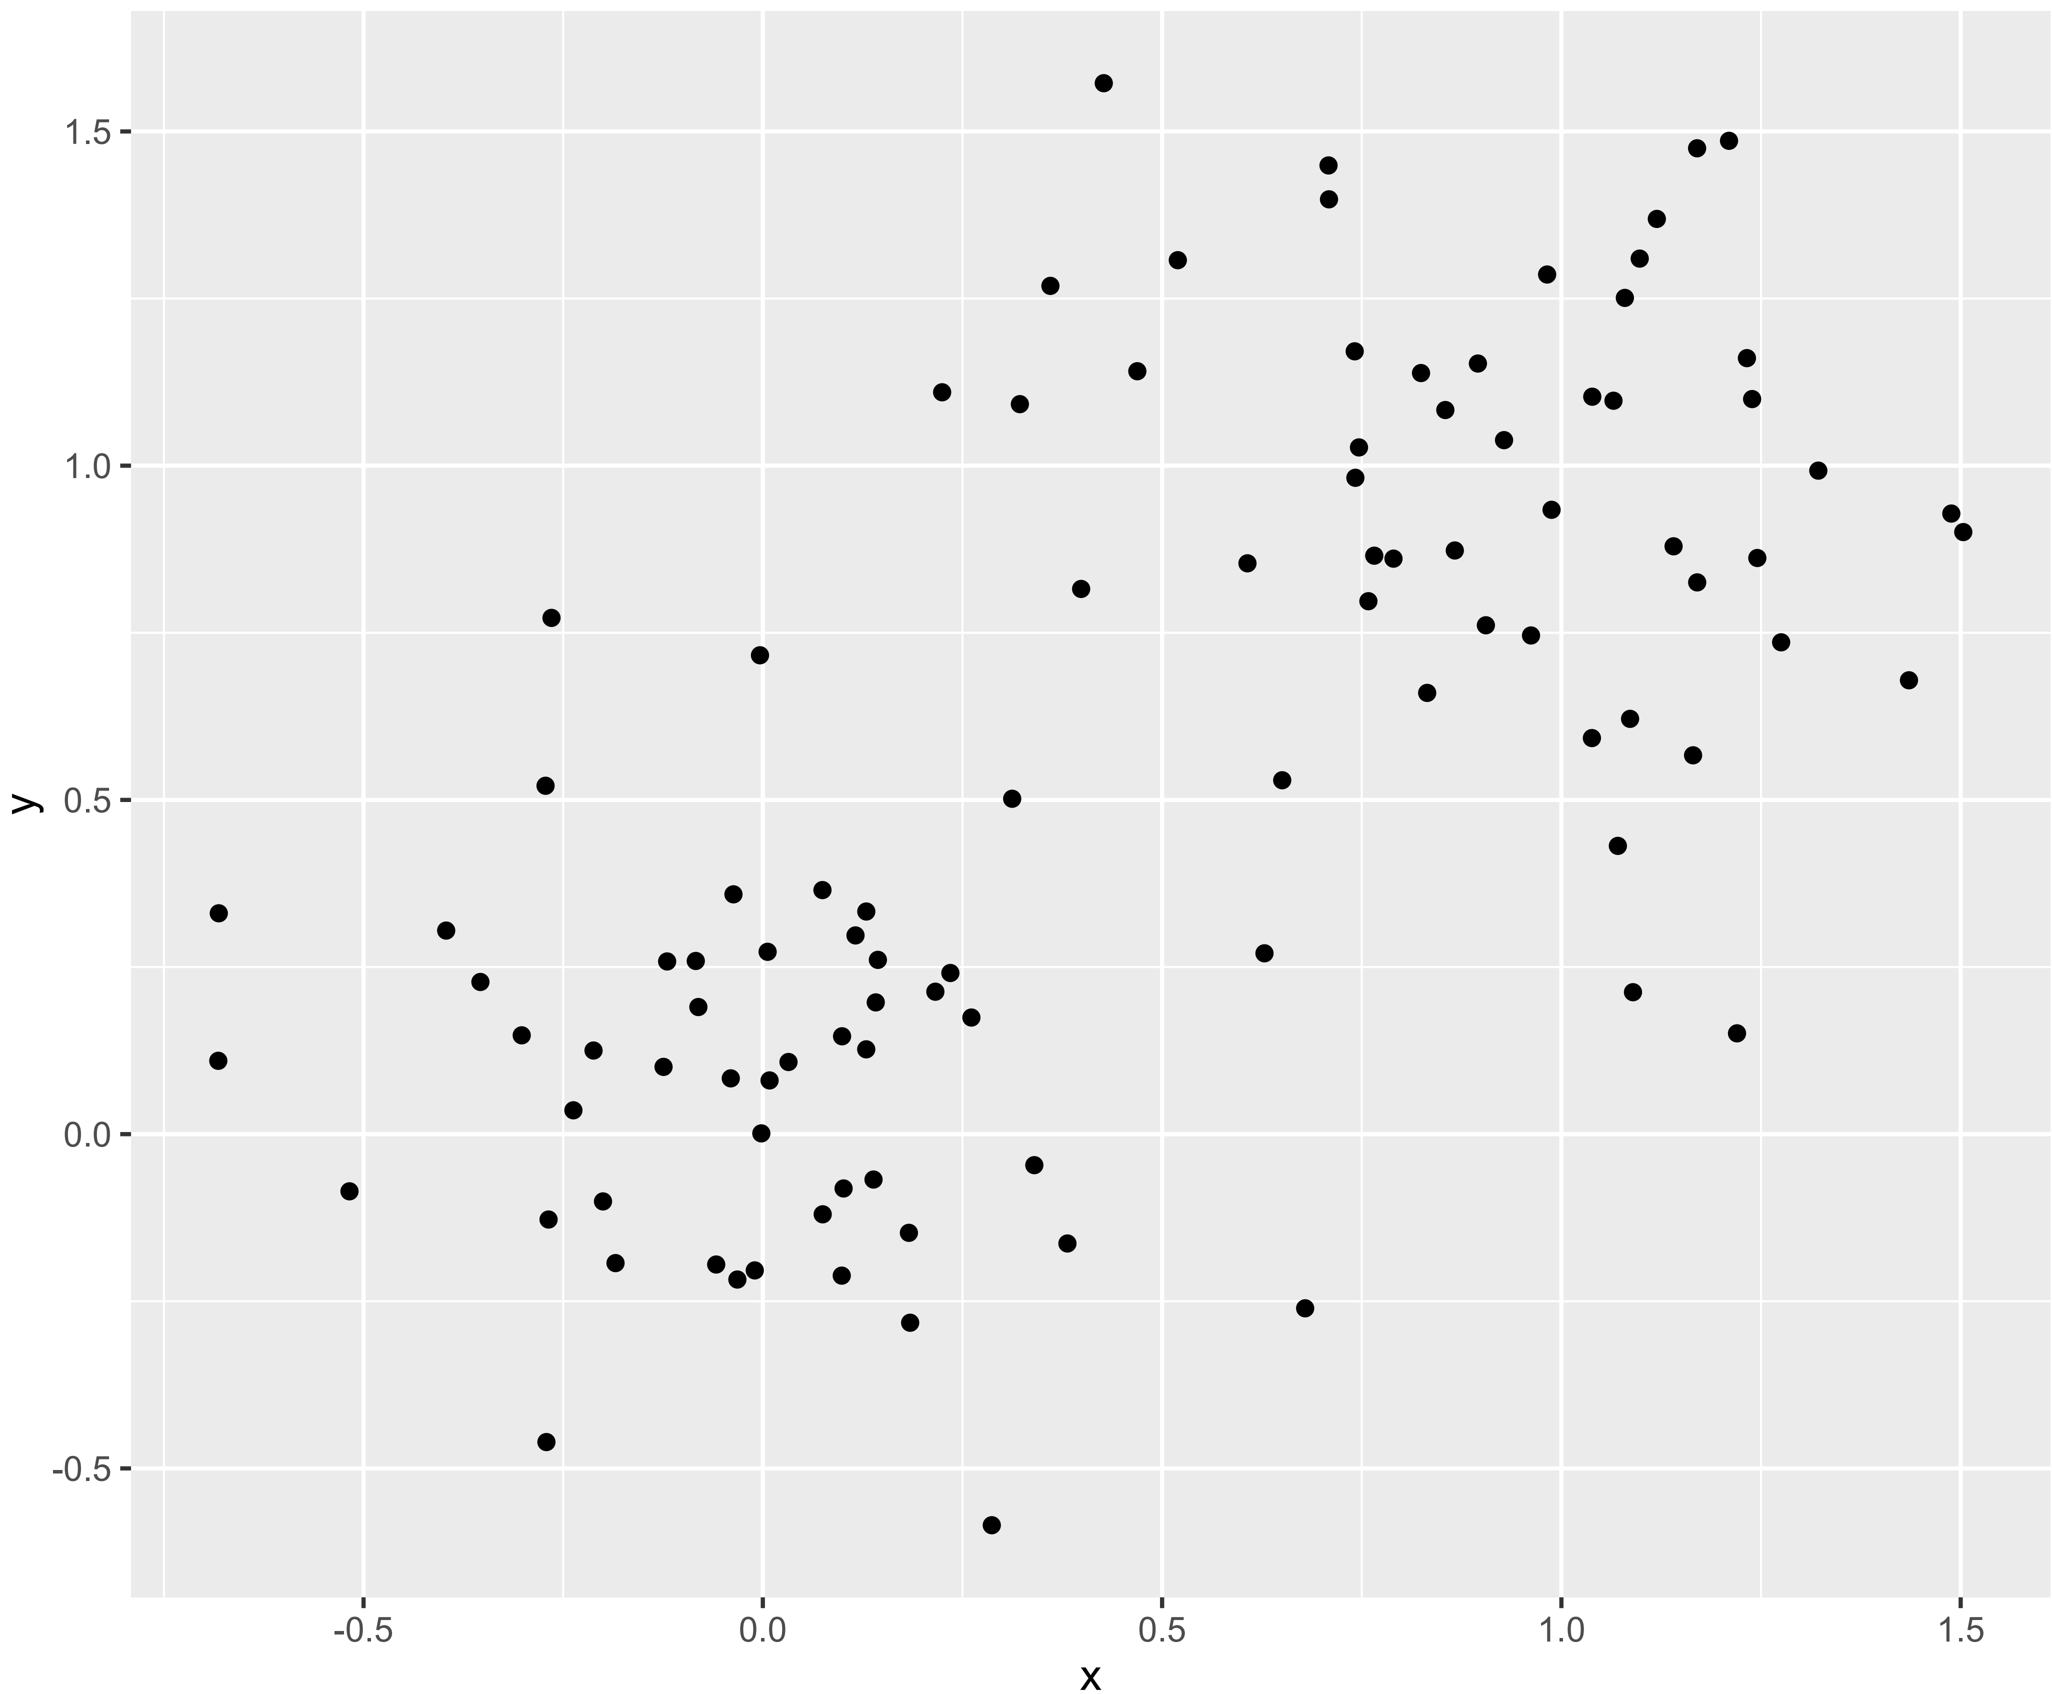
1.5
1.0
0.5
0.0
-0.5
y
-0.5
0.0
0.5
1.0
1.5
x

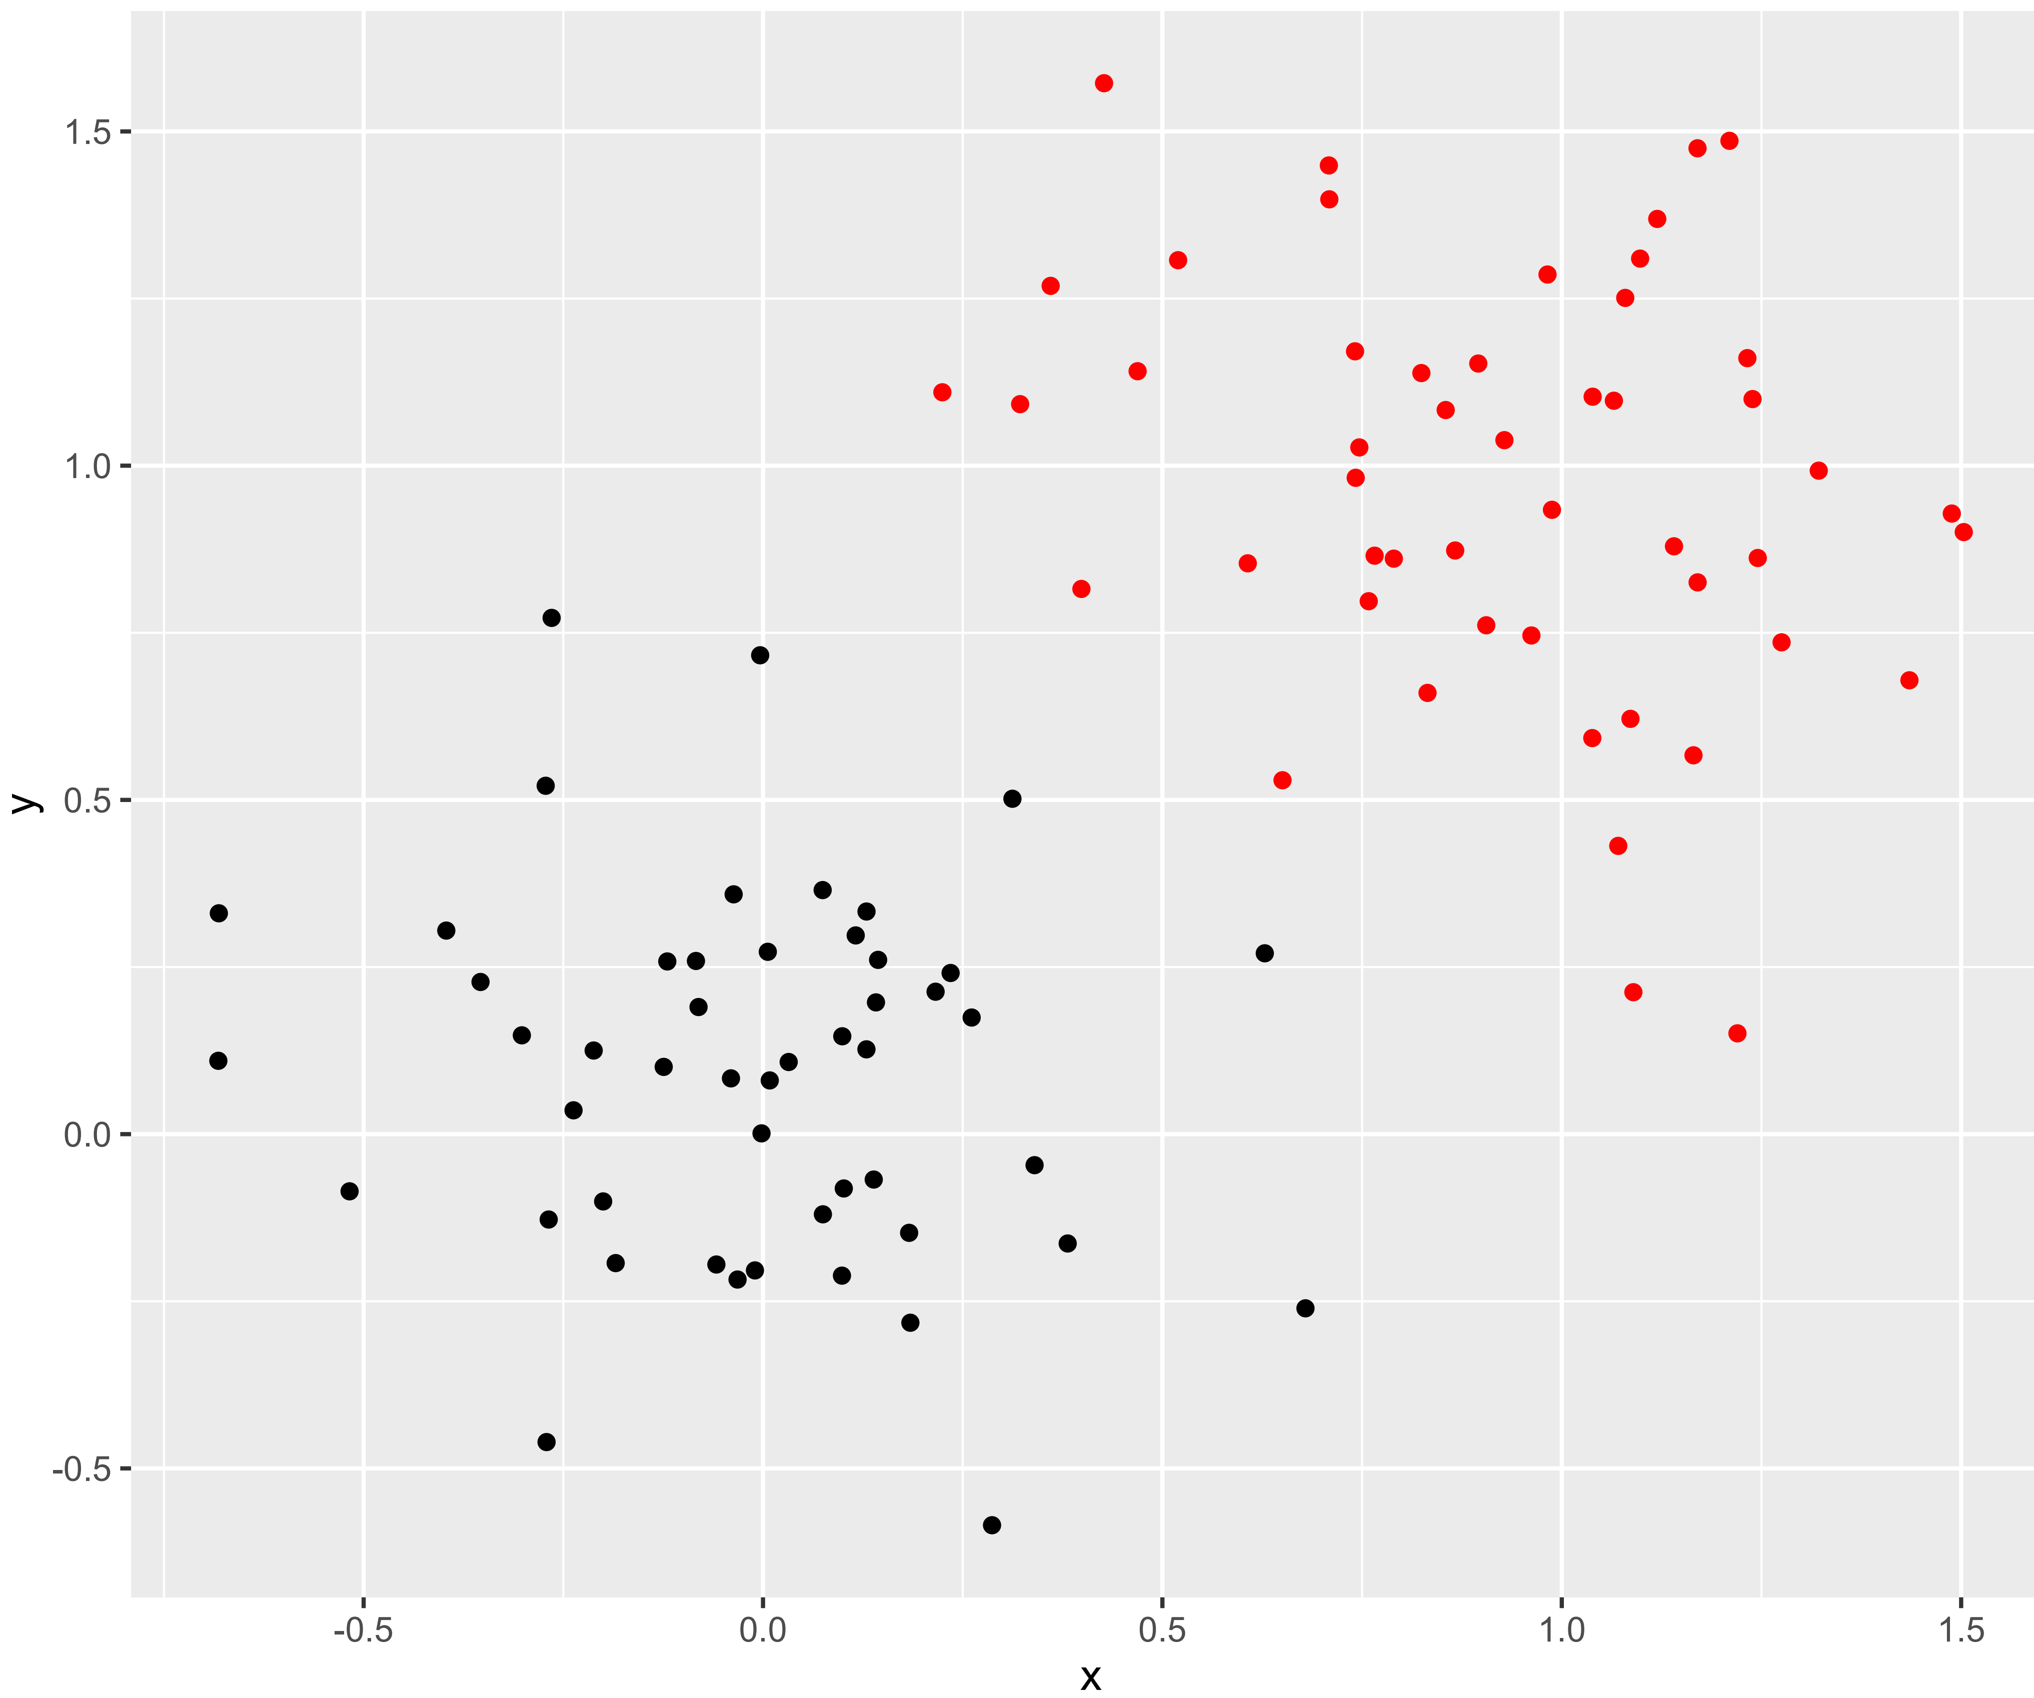

1.5
1.0
0.5
0.0
-0.5
y
-0.5
0.0
0.5
1.0
1.5
x

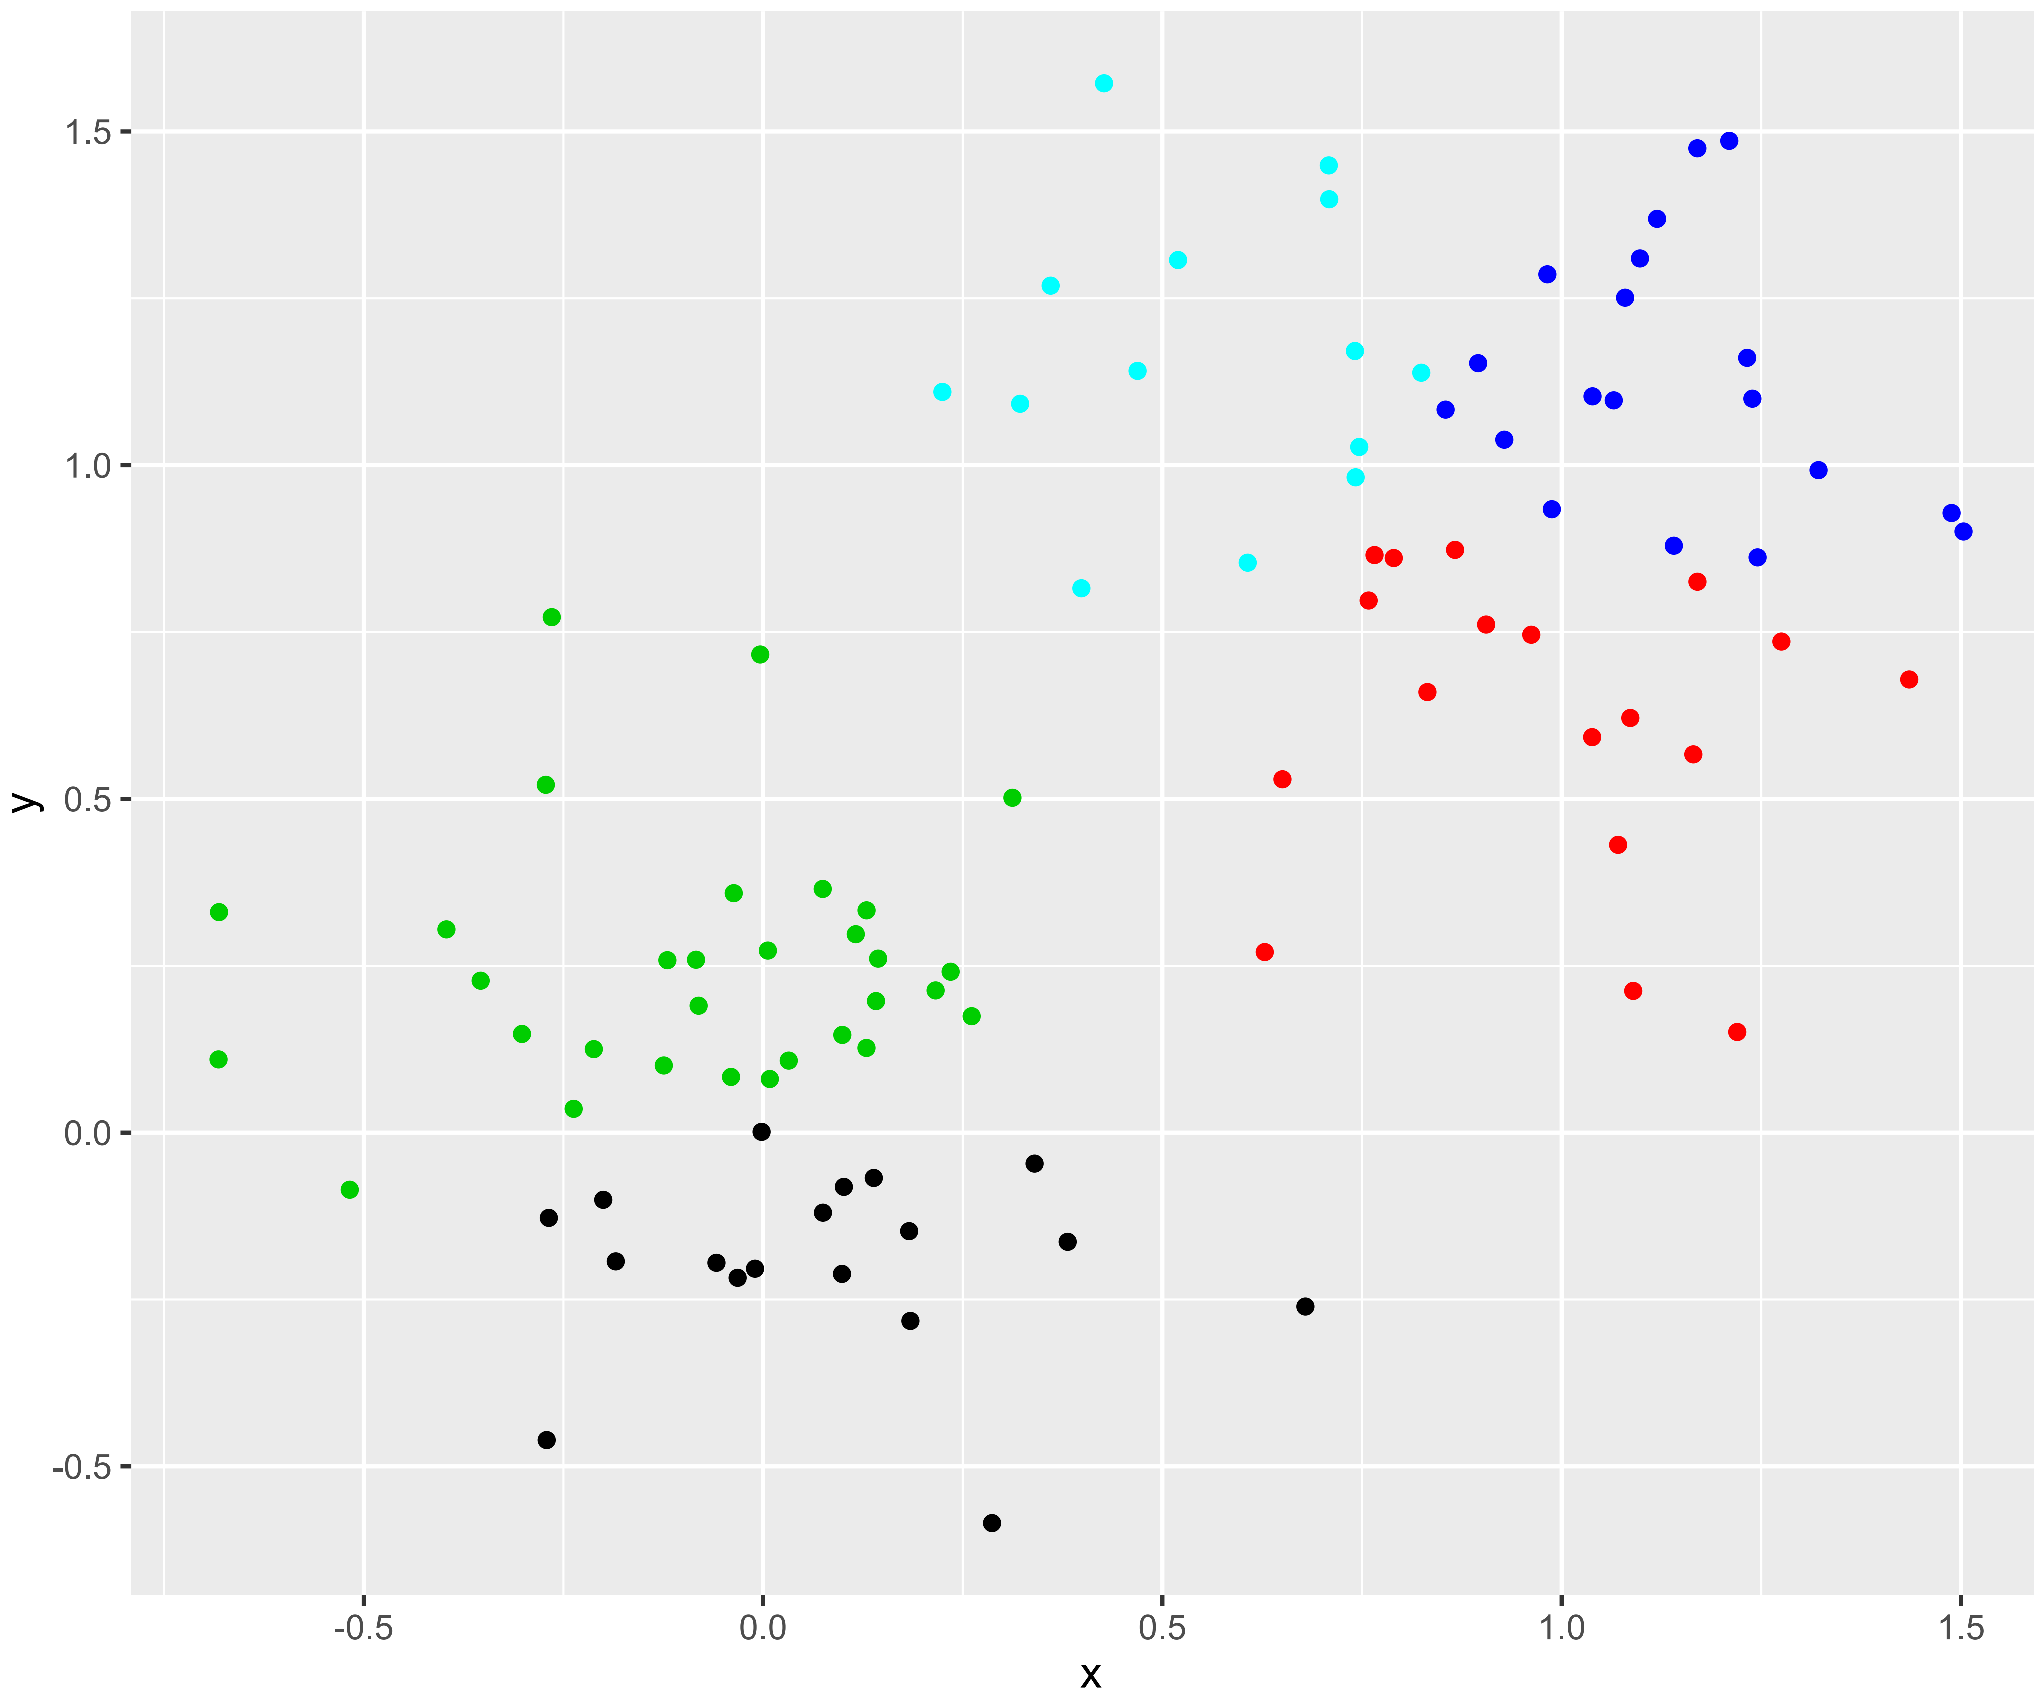

1.5
1.0
0.5
0.0
-0.5
y
-0.5
0.0
0.5
1.0
1.5
x

The “elbow”

Core idea: clusters should minimize the within-cluster variance

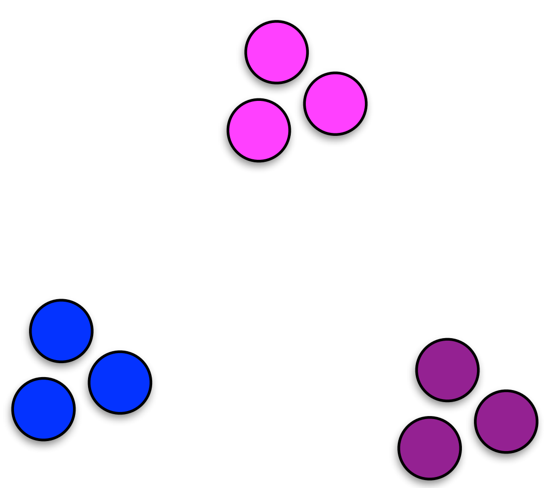

good

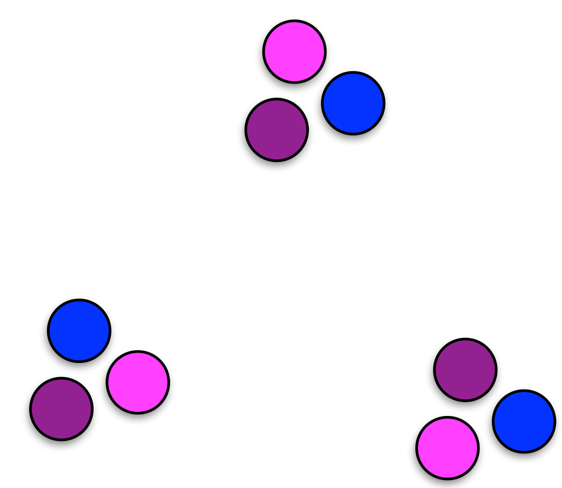

bad

The "elbow"

Core idea: clusters should minimize the within-cluster variance

within-cluster
sum of squares

$$\sum_{i=1}^{F}(x_i - \mu_i)^2$$

for each cluster

The “elbow”

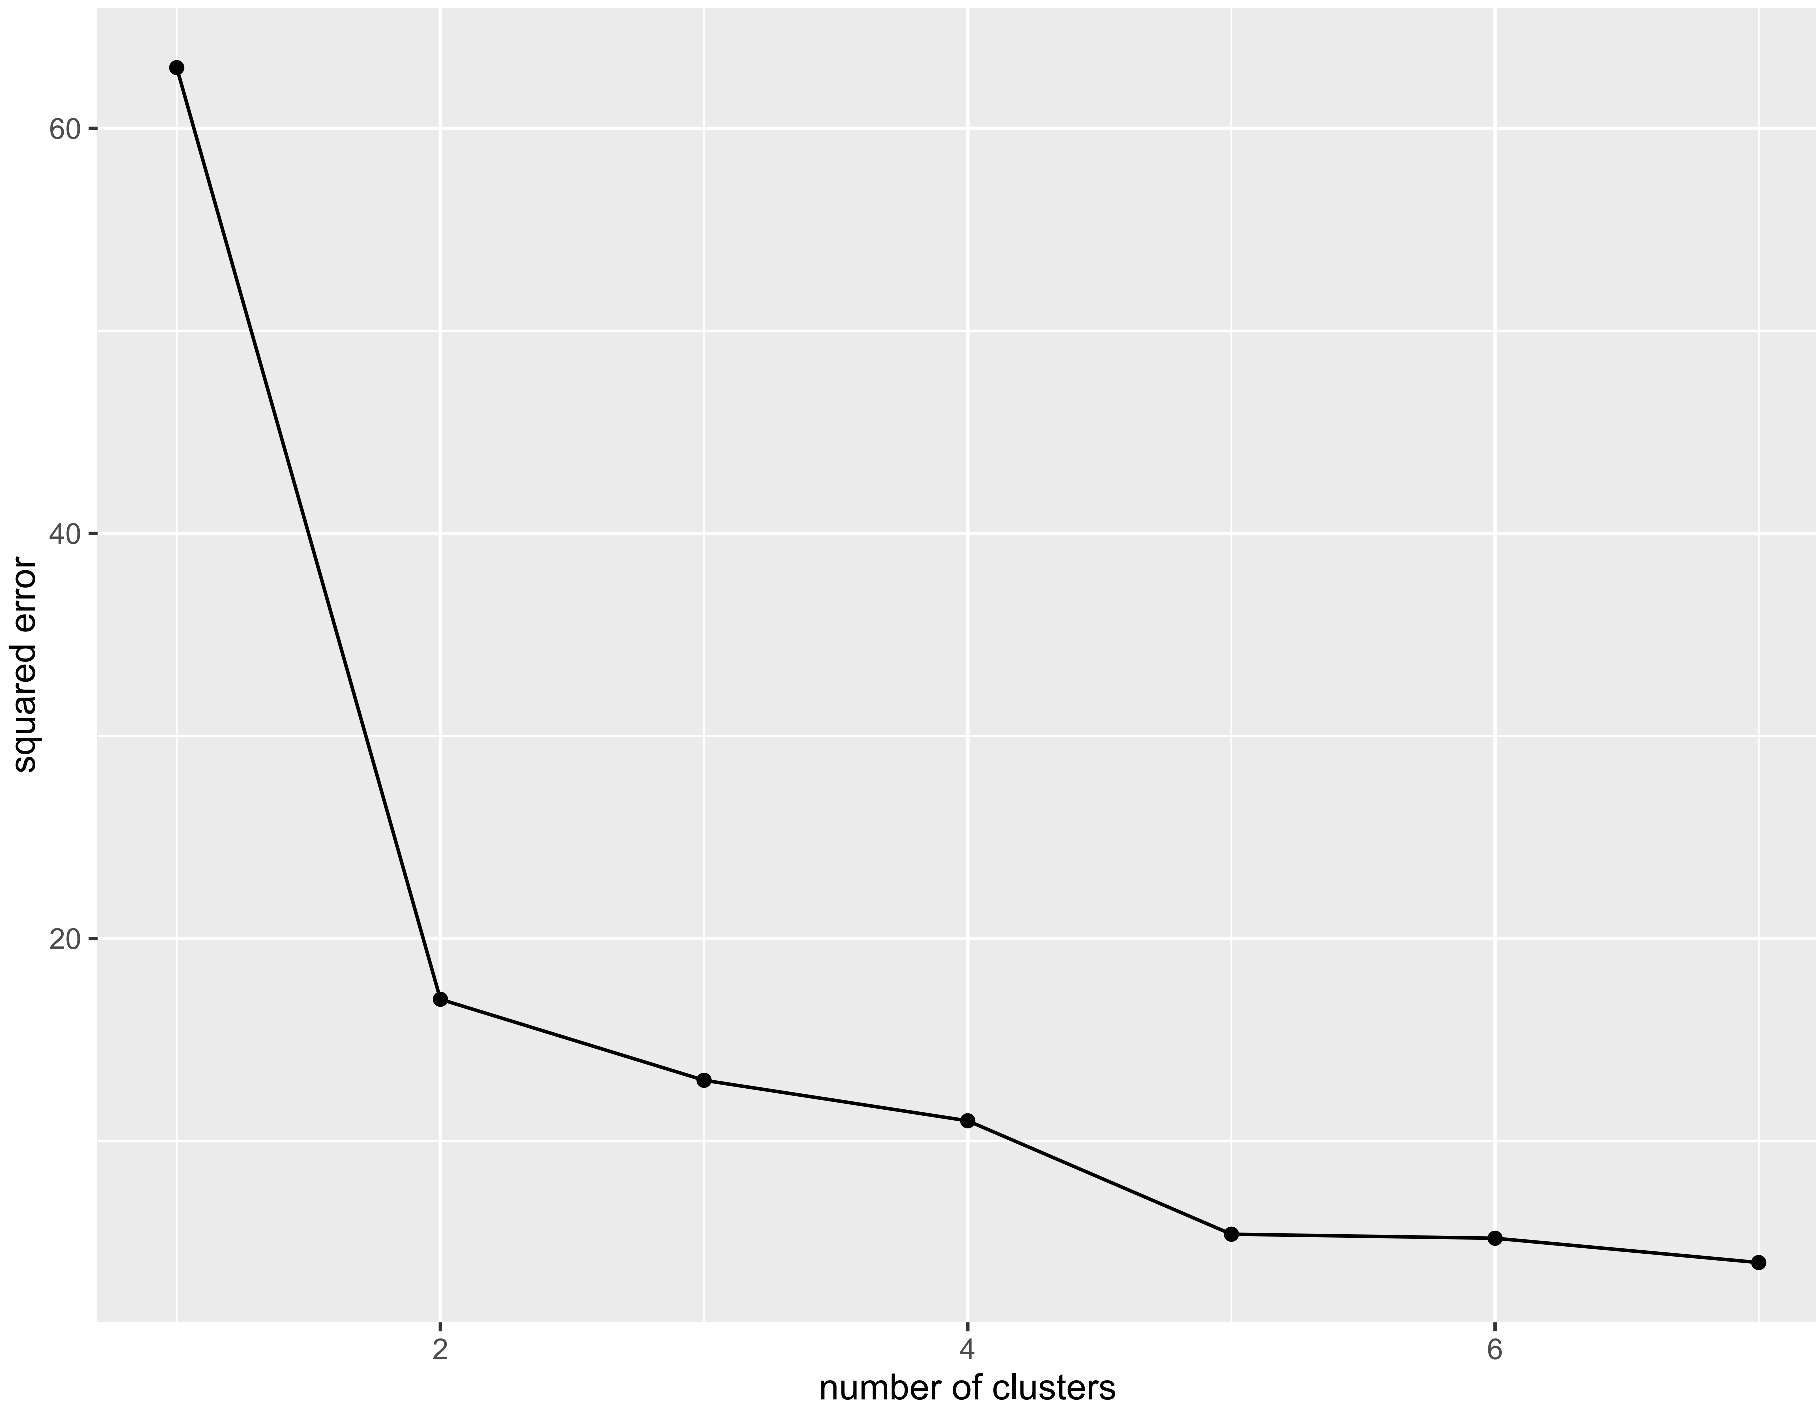

Gap statistic

- How much variance should we expect to see for a given number of clusters?
- Choose number of clusters that maximizes the “gap” between the observed variance and the expected variance for a given K.

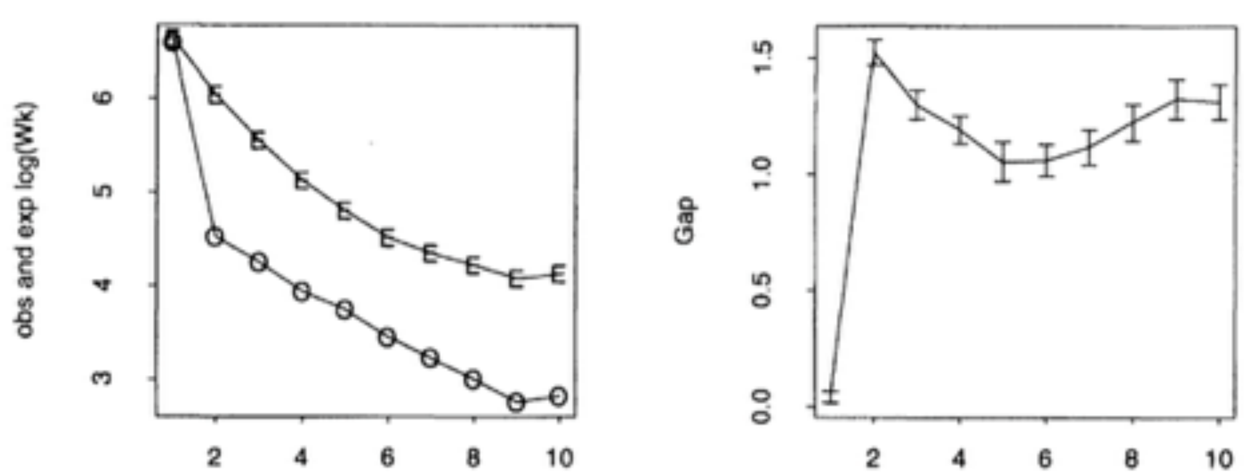

Tibshirani et al., “Estimating the number of clusters in a data set via the gap statistic”
http://web.stanford.edu/~hastie/Papers/gap.pdf

Hierarchical clustering

Core idea: build a binary tree of a set of data points by repeatedly merging the two most similar elements

Hierarchical clustering

Algorithm 1 Hierarchical agglomerative clustering

Data: N training data points $x \in \mathbb{R}^F$
Let X denote a set of objects x
Given some linkage function $d(X, X') \to \mathbb{R}$
Initialize clusters $\mathcal{C} = \{C_1, \ldots, C_N\}$ to singleton data points
while data points not in one cluster **do**
    Identify X, Y as clusters with smallest linkage function among clusters in $\mathcal{C}$
    Create new cluster $Z = X \cup Y$
    remove X, Y from $\mathcal{C}$
    add Z to $\mathcal{C}$
end while

Hierarchical clustering

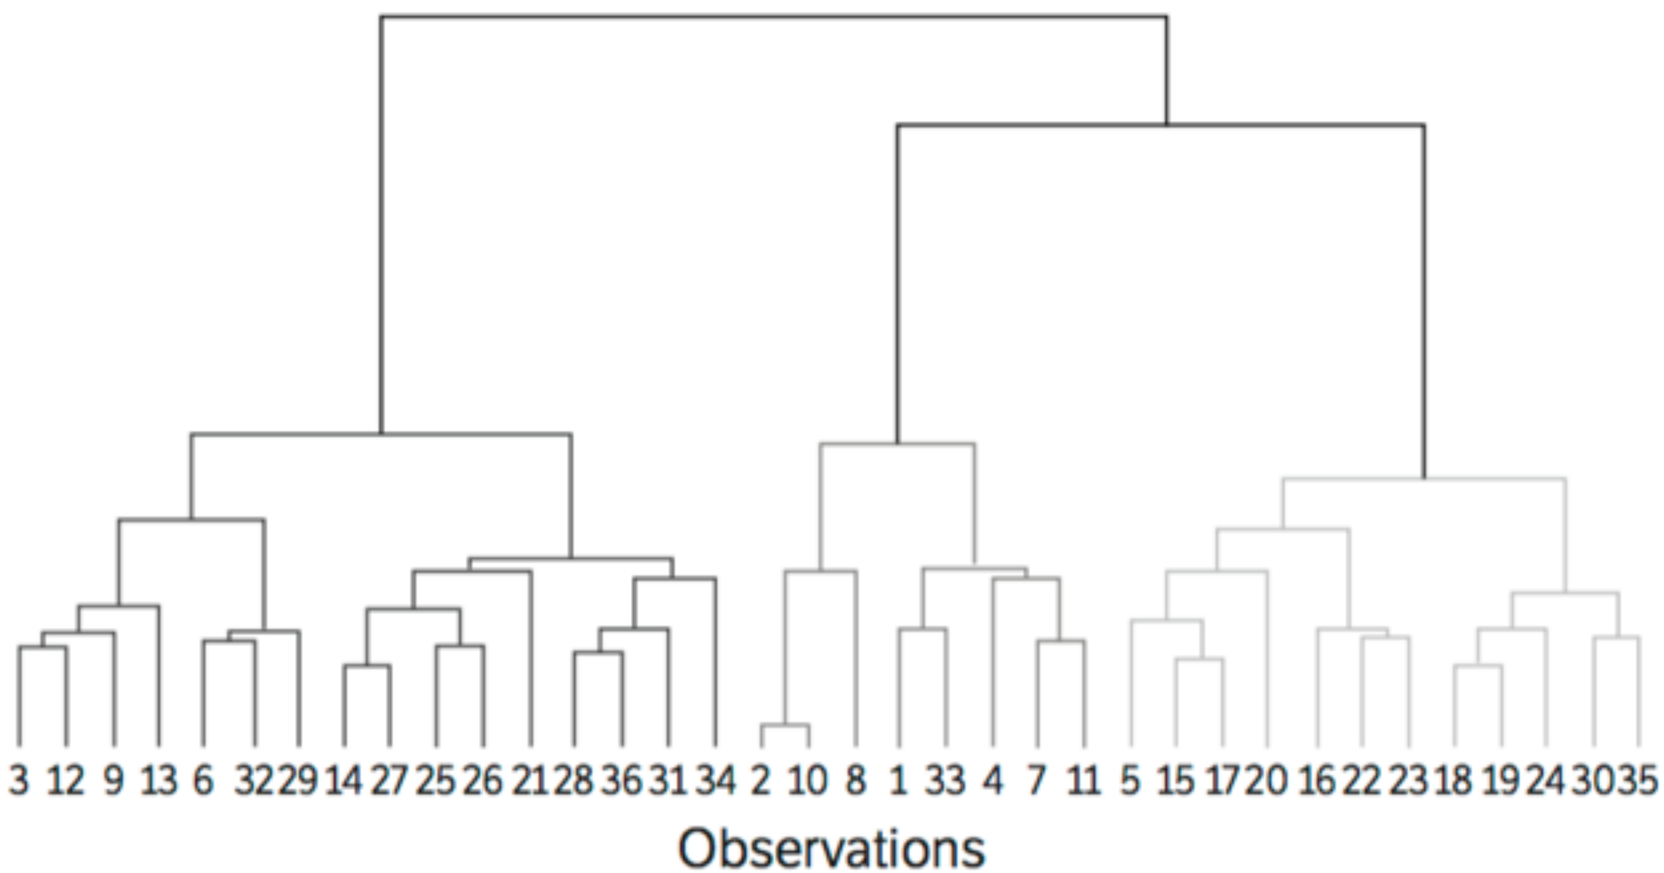

A Midsummer Night's Dream (3)
Twelfth Night (12)
Much Ado About Nothing (9)
Two Gentlemen (13)
Measure for Measure (6)
Othello (32)
Julius Caesar (29)

The Winter's Tale (14)
Cymbeline (27)
Antony and Cleopatra (25)
Coriolanus (26)
Henry VIII (21)
Hamlet (28)
Troilus and Cressida (36)
Macbeth (31)
Timon of Athens (34)

All's Well That Ends Well (2)
Taming of the Shrew (10)
Merry Wives of Windsor (8)
A Midsummer Night's Dream (1)
Romeo and Juliet (33)
Comedy of Errors (4)
Merchant of Venice (7)
The Tempest (11)

Love's Labours' Lost (5)
1 Henry IV (15)
2 Henry IV (17)
Henry V (20)
1 Henry VI (16)
King John (22)
Richard II (23)

2 Henry VI (18)
2 Henry VI (19)
Richard III (24)
King Lear (30)
Titus Andronicus (35)

Allison et al. 2009

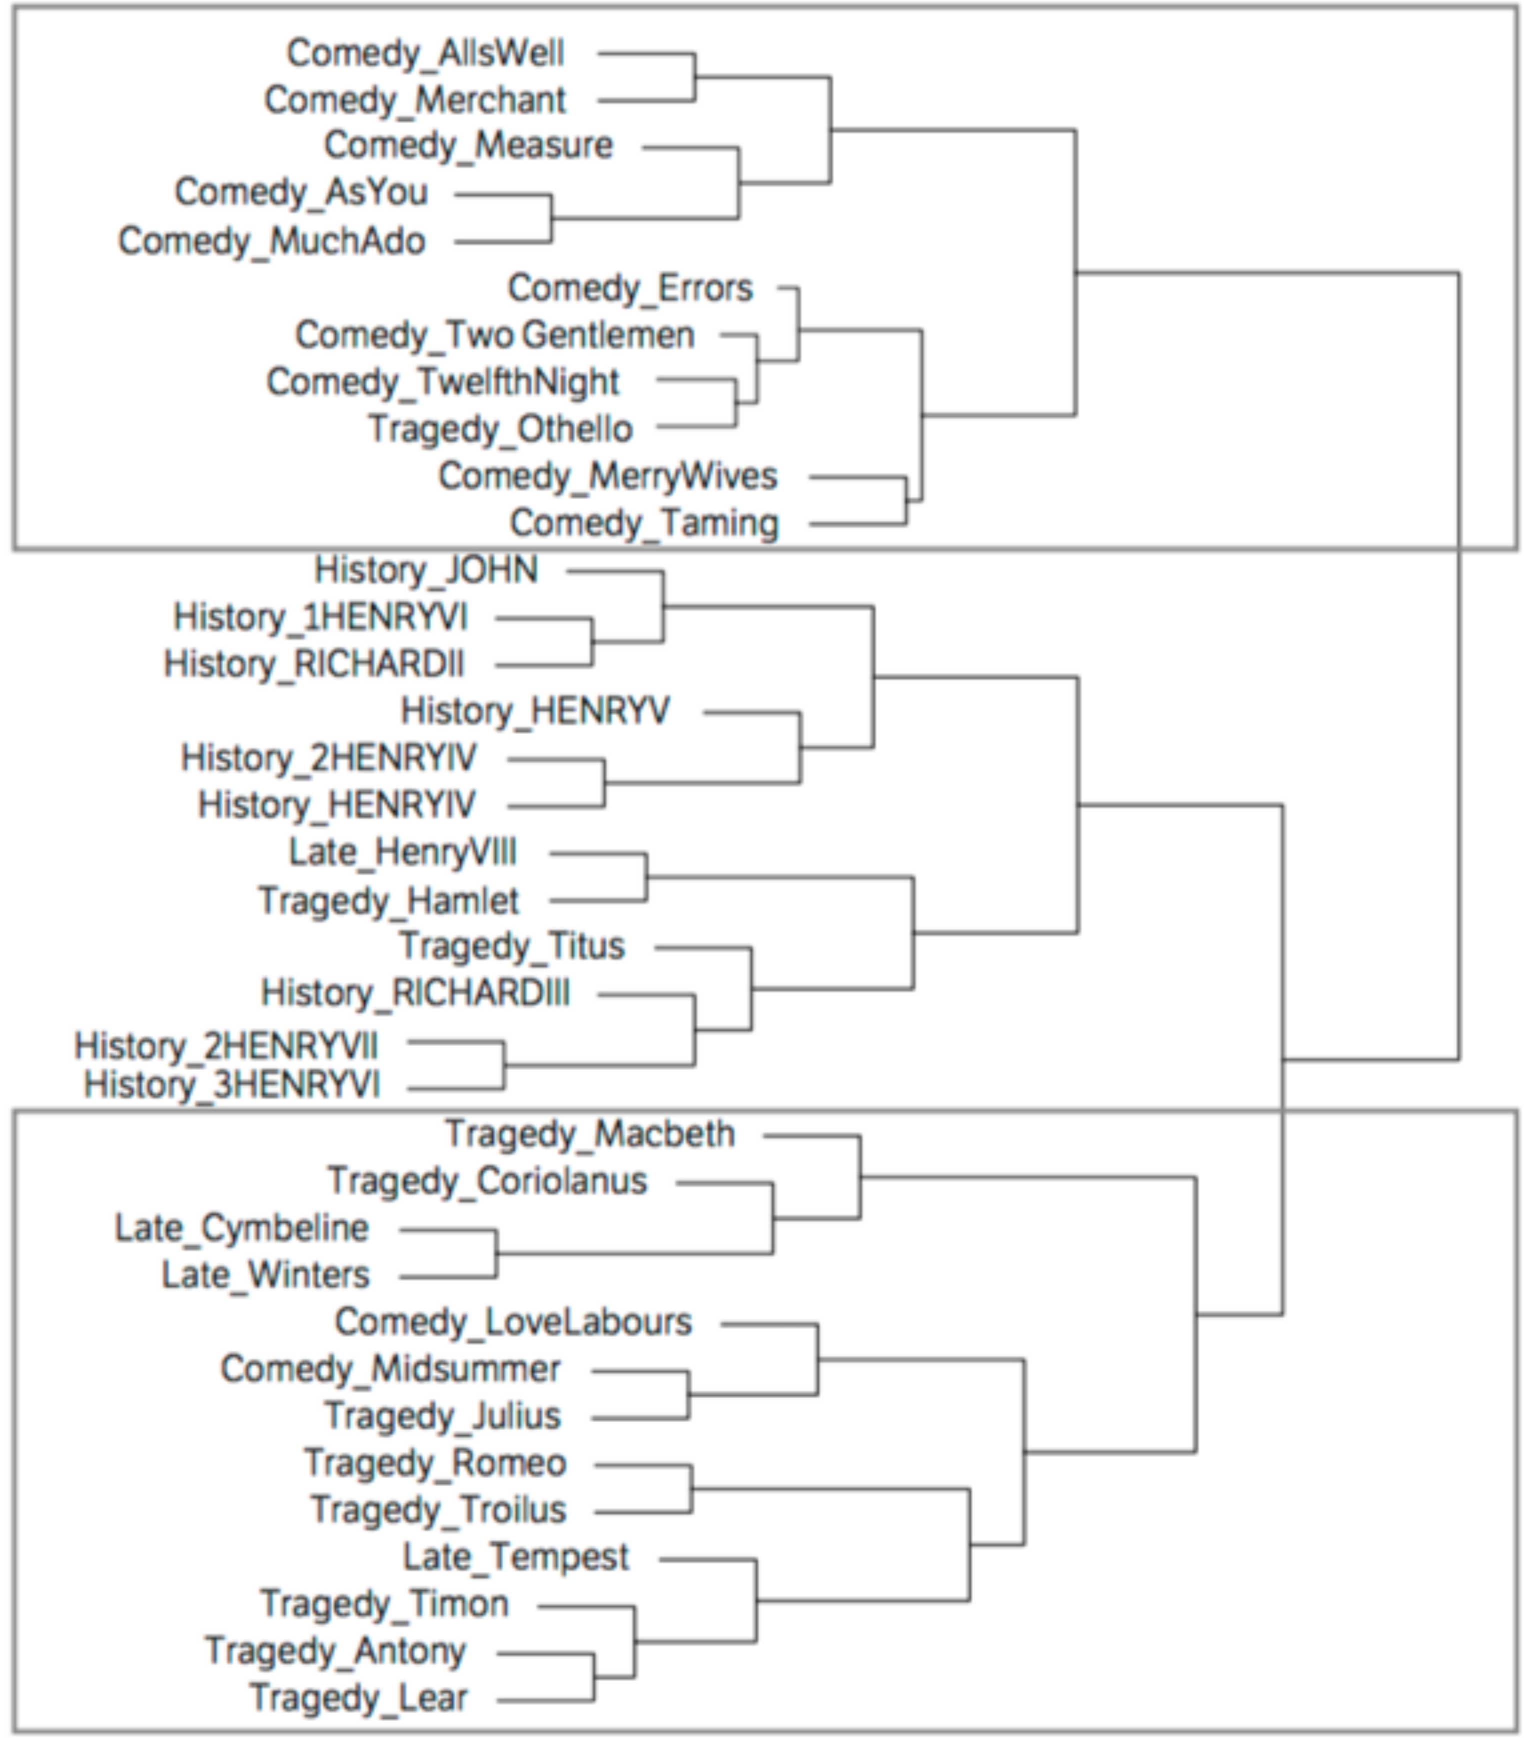

Allison et al. 2009

Hierarchical clustering

We know how to compare data points with distance metrics.

How do we compare sets of data points?

Single linkage

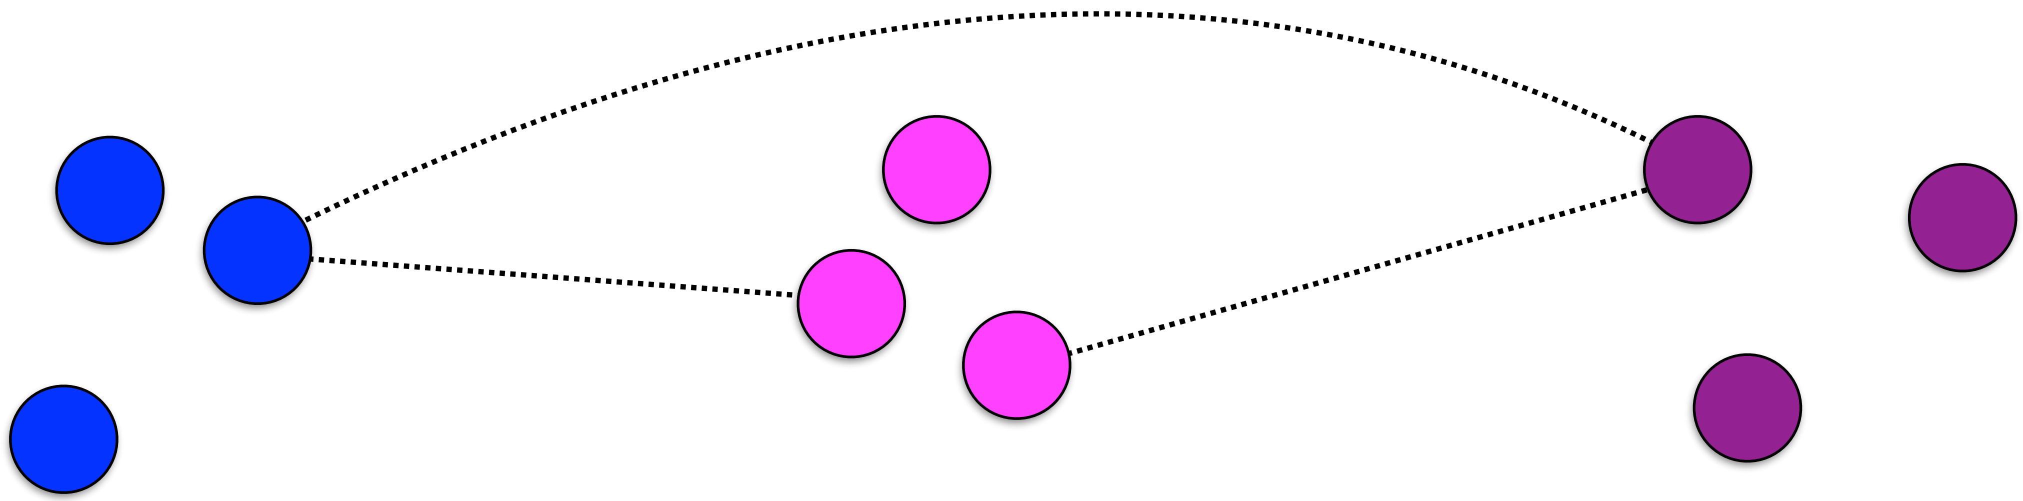

$$\min_{x \in A,\, y \in B} \mathrm{Dis}(x, y)$$

Complete linkage

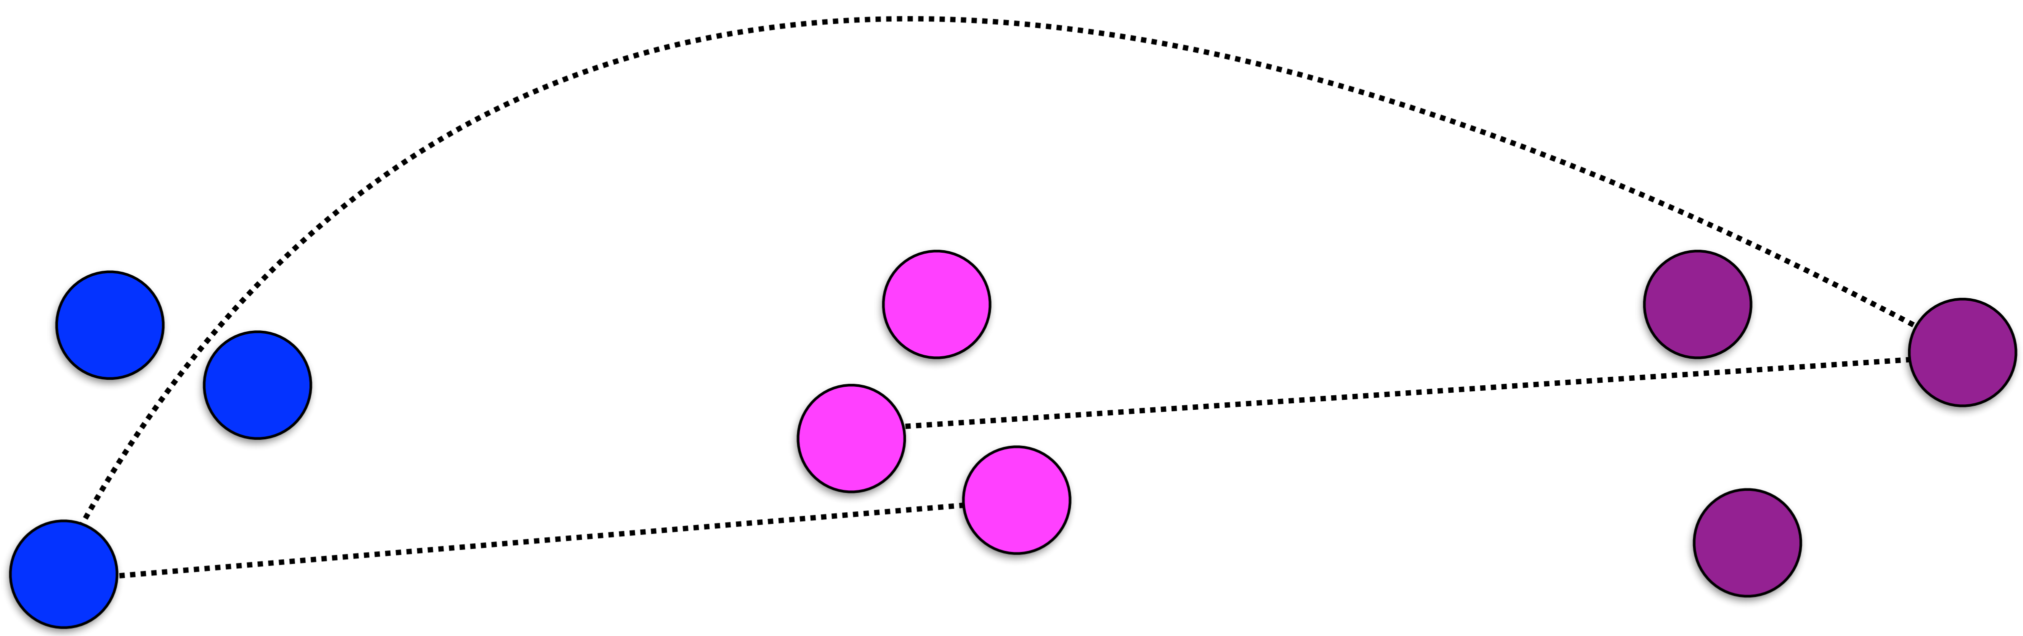

$$\max_{x \in A,\, y \in B} \mathrm{Dis}(x, y)$$

Average linkage

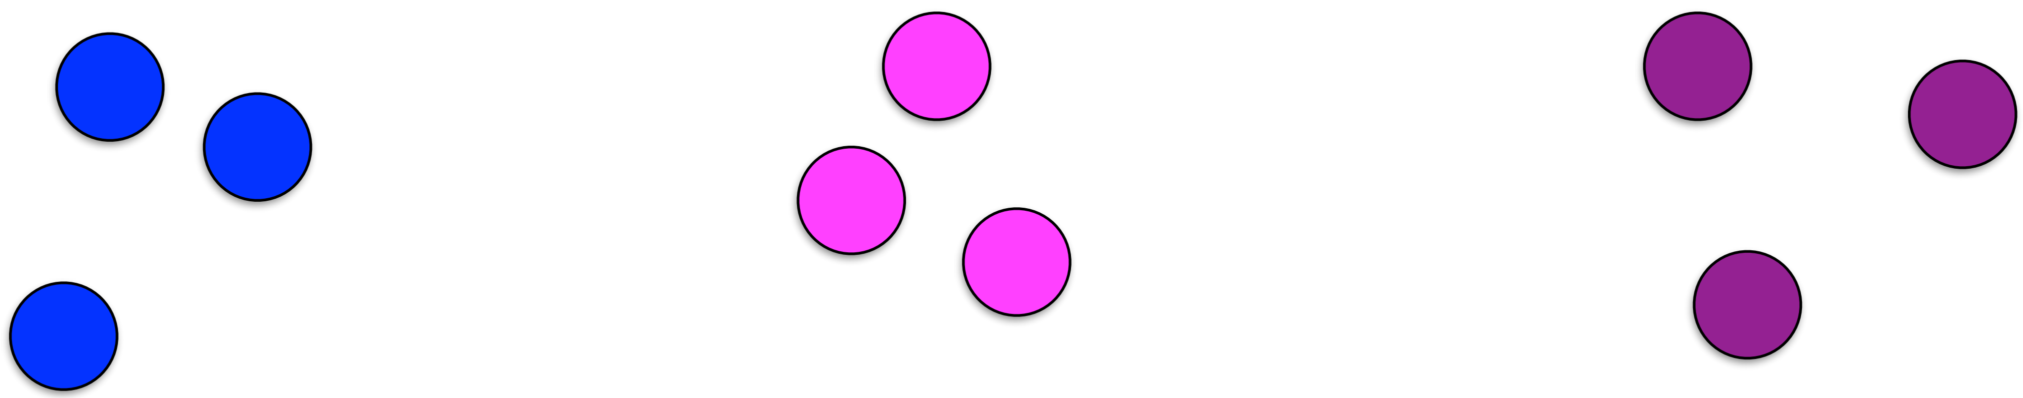

$$\frac{\sum_{x \in A,\, y \in B} \mathrm{Dis}(x, y)}{|A| \times |B|}$$

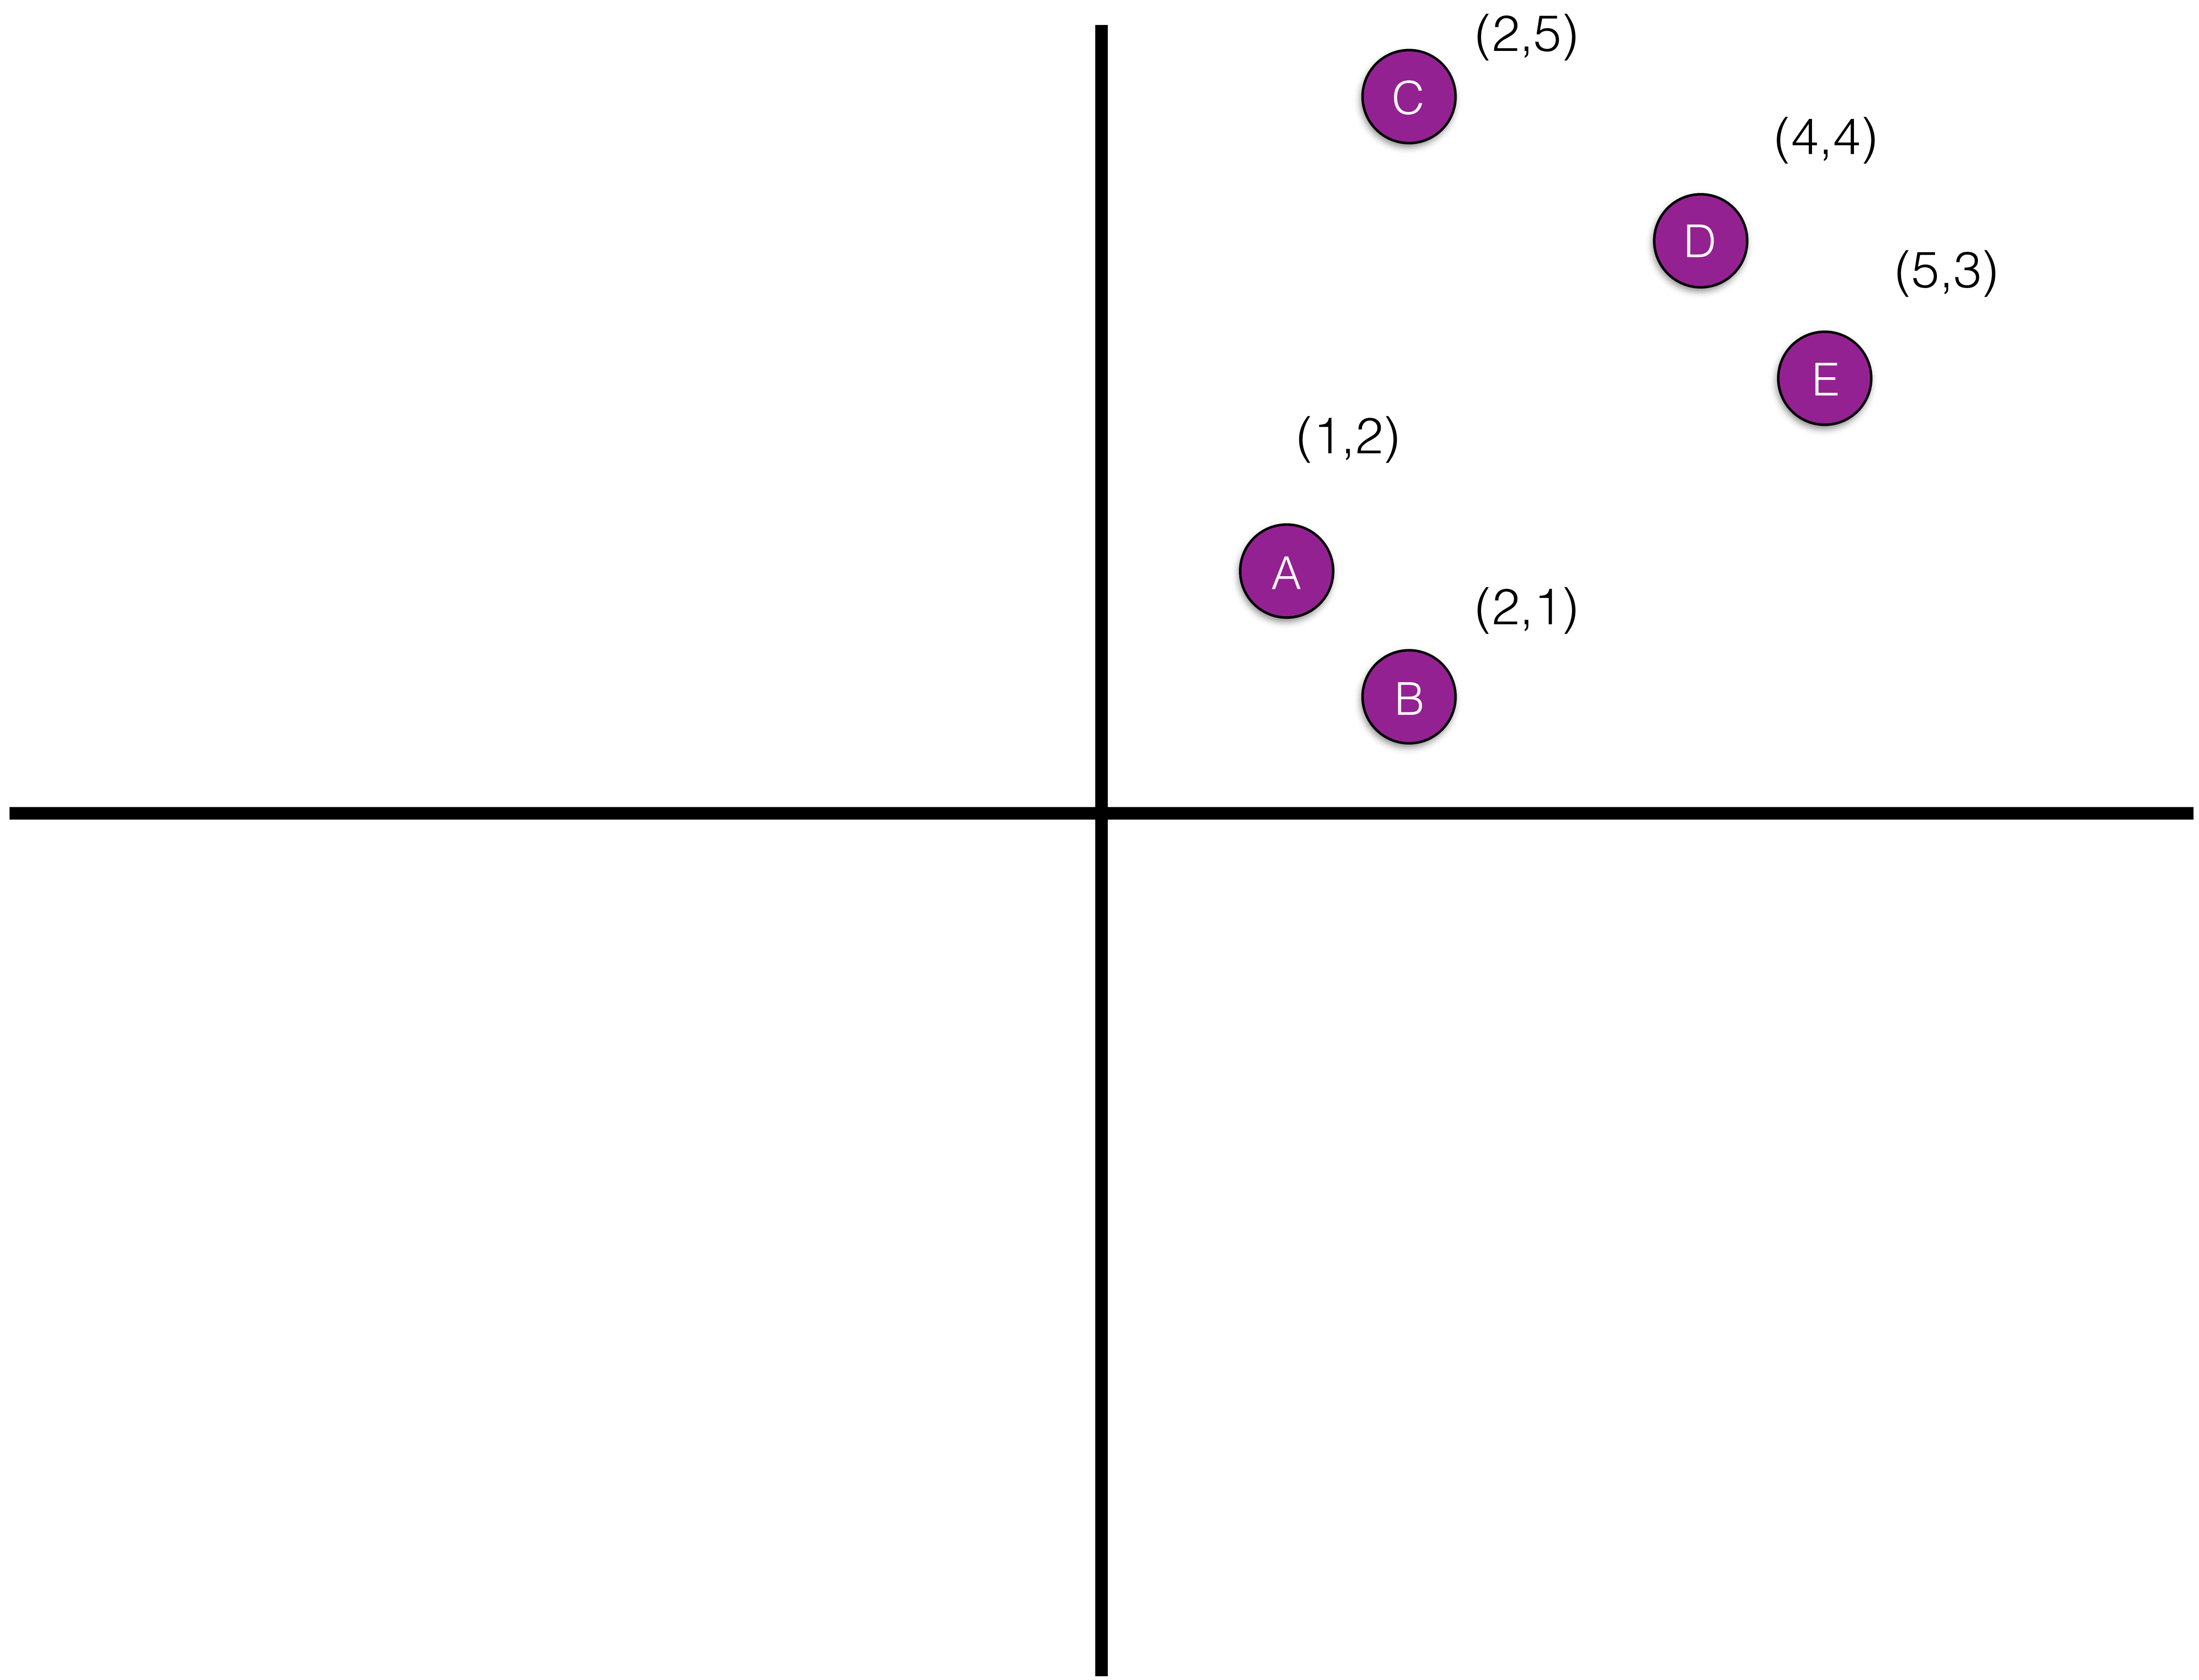

(2,5)
C
(4,4)
D
(5,3)
E
(1,2)
A
(2,1)
B

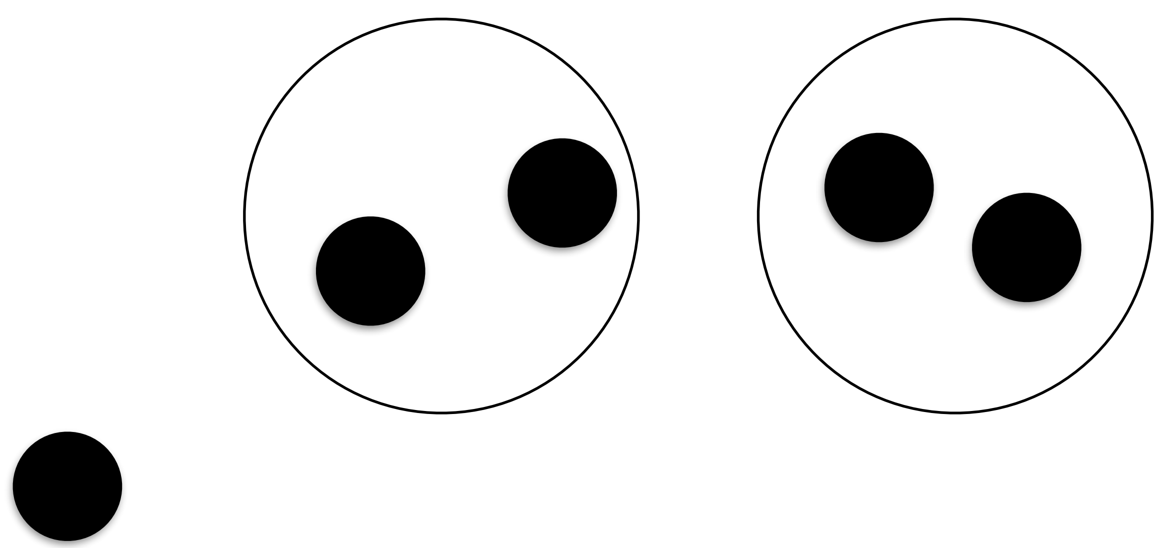

Single linkage may link bigger clusters together
before outliers

Complete linkage

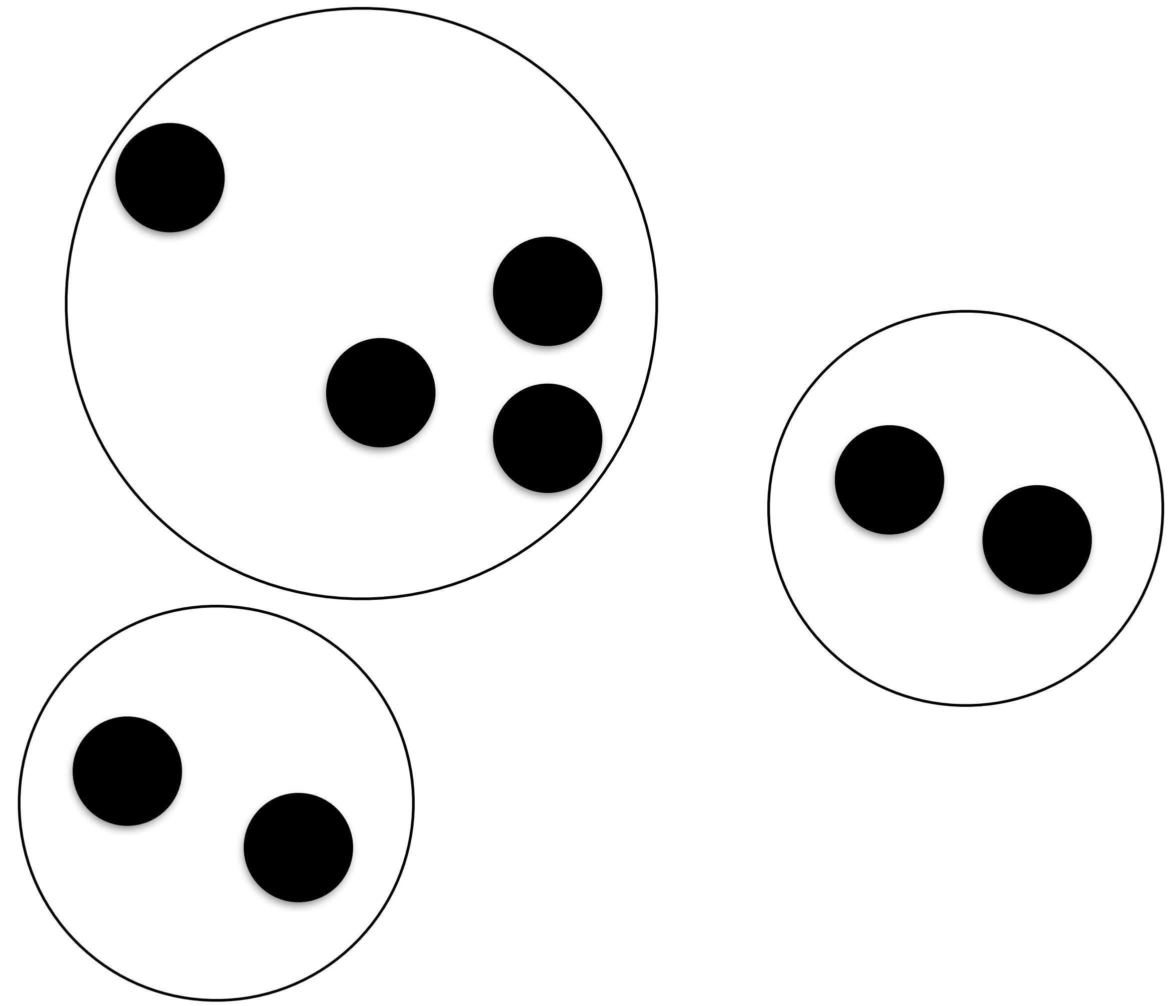

Complete linkage may *not* link close clusters together because of outliers

- Allison et al., “Quantitative Formalism: an Experiment”

DocuScope

Dictionary mapping ngrams to classes

First Person	Numbers	Positivity
about me	six-wheeled	perpetual adorations
about my	275 degrees	mated with
am	three-card loo	hugging yourself
I	695	striking responsive cord
I'd	four-ply	wassailing
I'll	half-way	plucked up your spirits
I'm	three parts	offers ourselves
I for one	eight-member	promotive of
ich	third-world	enshrining
ich dien	3,5	devotes yourself
me	half-and-half measures	music lover
mea	8,3	delectated
meum	half-reclining	recharging my batteries
mine	26	recommends you for
my	634	shadow of your smile
myself	five-rater	regaining our composure

MFW

Only unigrams with relative frequency > 0.03

a	not
all	of
and	on
as	p_apos
at	p_comma
be	p_exlam
but	p_hyphen
by	p_period
for	p_ques
from	p_quote
had	p_semi
have	said
he	she
her	so
him	that
his	the
i	this
in	to
is	was
it	which
me	with
my	you

Hierarchical clustering

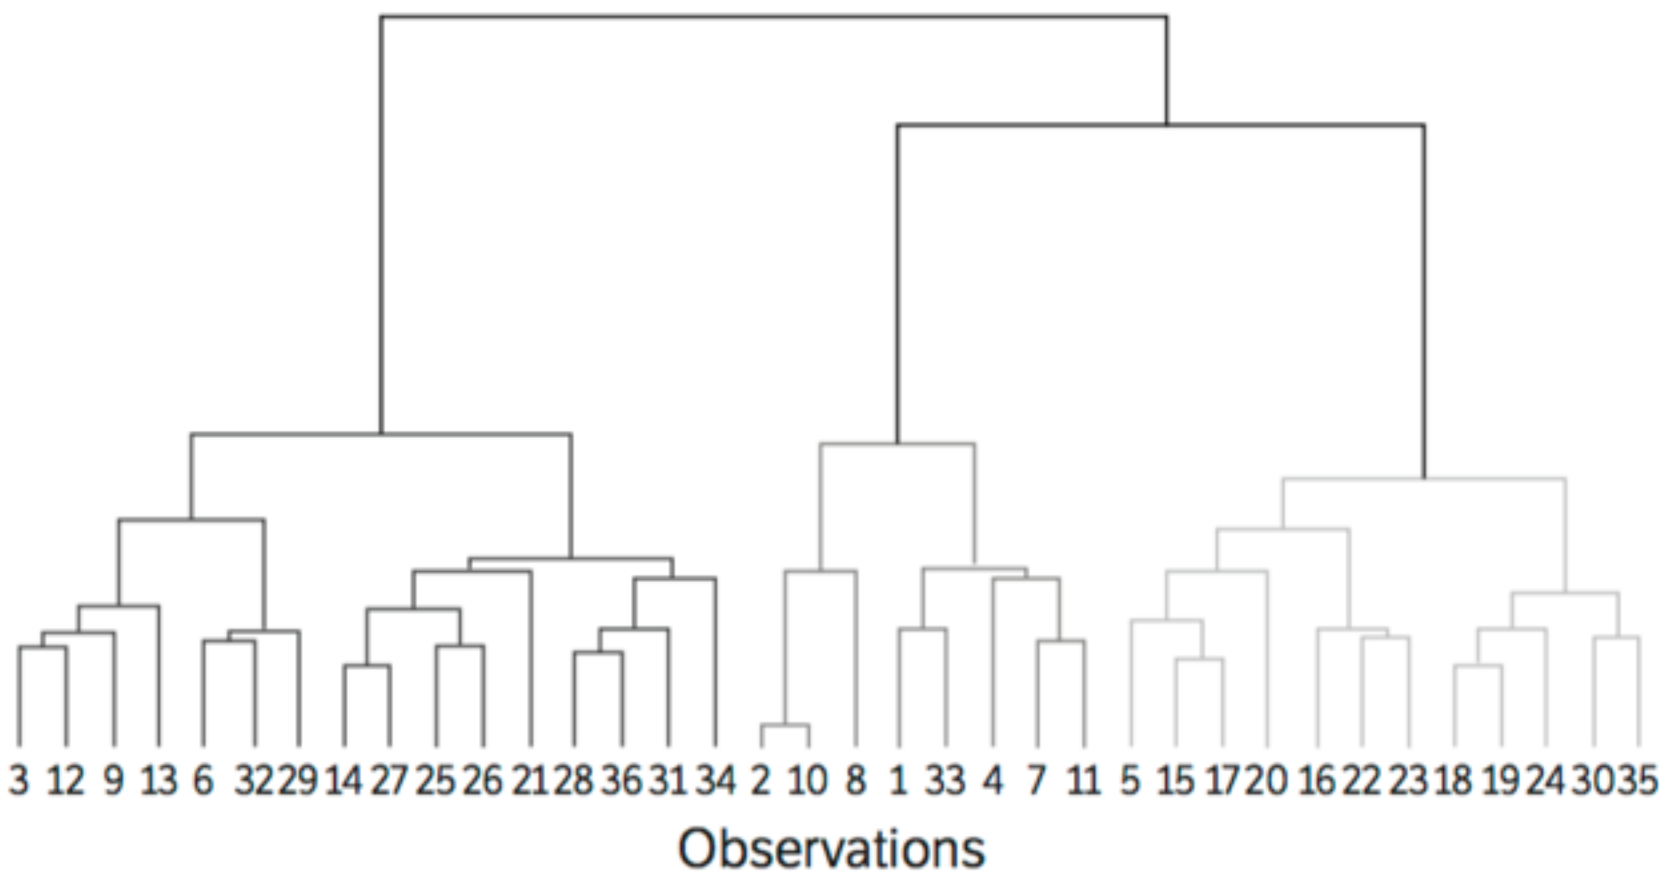

A Midsummer Night's Dream (3)
Twelfth Night (12)
Much Ado About Nothing (9)
Two Gentlemen (13)
Measure for Measure (6)
Othello (32)
Julius Caesar (29)

The Winter's Tale (14)
Cymbeline (27)
Antony and Cleopatra (25)
Coriolanus (26)
Henry VIII (21)
Hamlet (28)
Troilus and Cressida (36)
Macbeth (31)
Timon of Athens (34)

All's Well That Ends Well (2)
Taming of the Shrew (10)
Merry Wives of Windsor (8)
A Midsummer Night's Dream (1)
Romeo and Juliet (33)
Comedy of Errors (4)
Merchant of Venice (7)
The Tempest (11)

Love's Labours' Lost (5)
1 Henry IV (15)
2 Henry IV (17)
Henry V (20)
1 Henry VI (16)
King John (22)
Richard II (23)

2 Henry VI (18)
2 Henry VI (19)
Richard III (24)
King Lear (30)
Titus Andronicus (35)

Allison et al. 2009

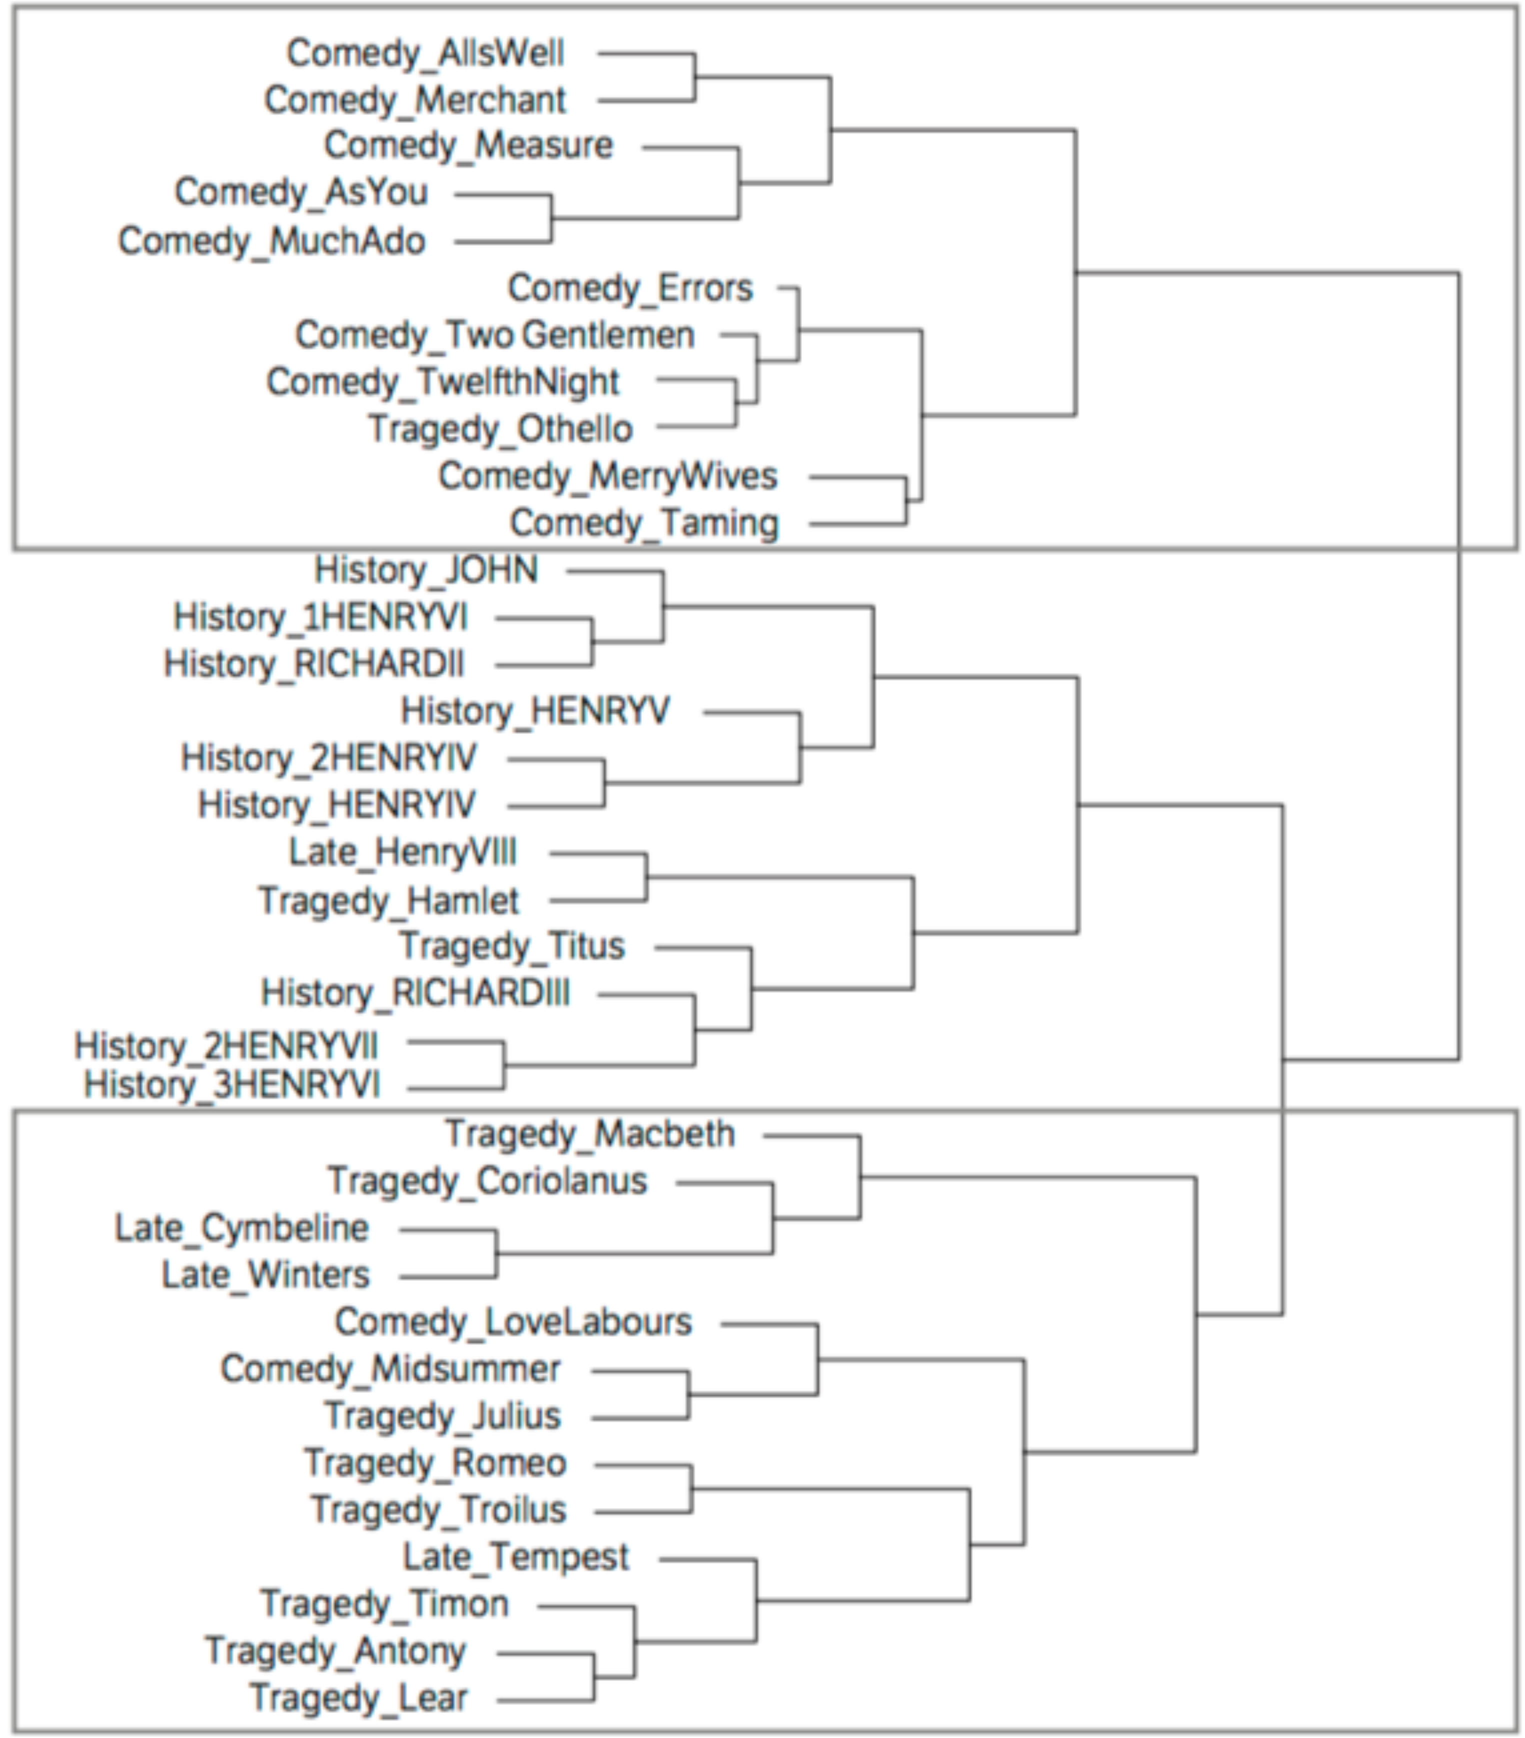

Allison et al. 2009

“But there is also a simpler explanation: namely, that these features which are so effective at differentiating genres, and so entwined with their overall texture – these features cannot offer new insights into structure, because they aren't independent traits, but mere consequences of higher-order choices. Do you want to write a story where each and every room may be full of surprises? Then locative prepositions, articles and verbs in the past tense are bound to follow. They are the effects of the chosen narrative structure.”

Project presentation

Tuesday April 25 (3) + Thursday April 27 (3)

12 min presentation +
5 min questions

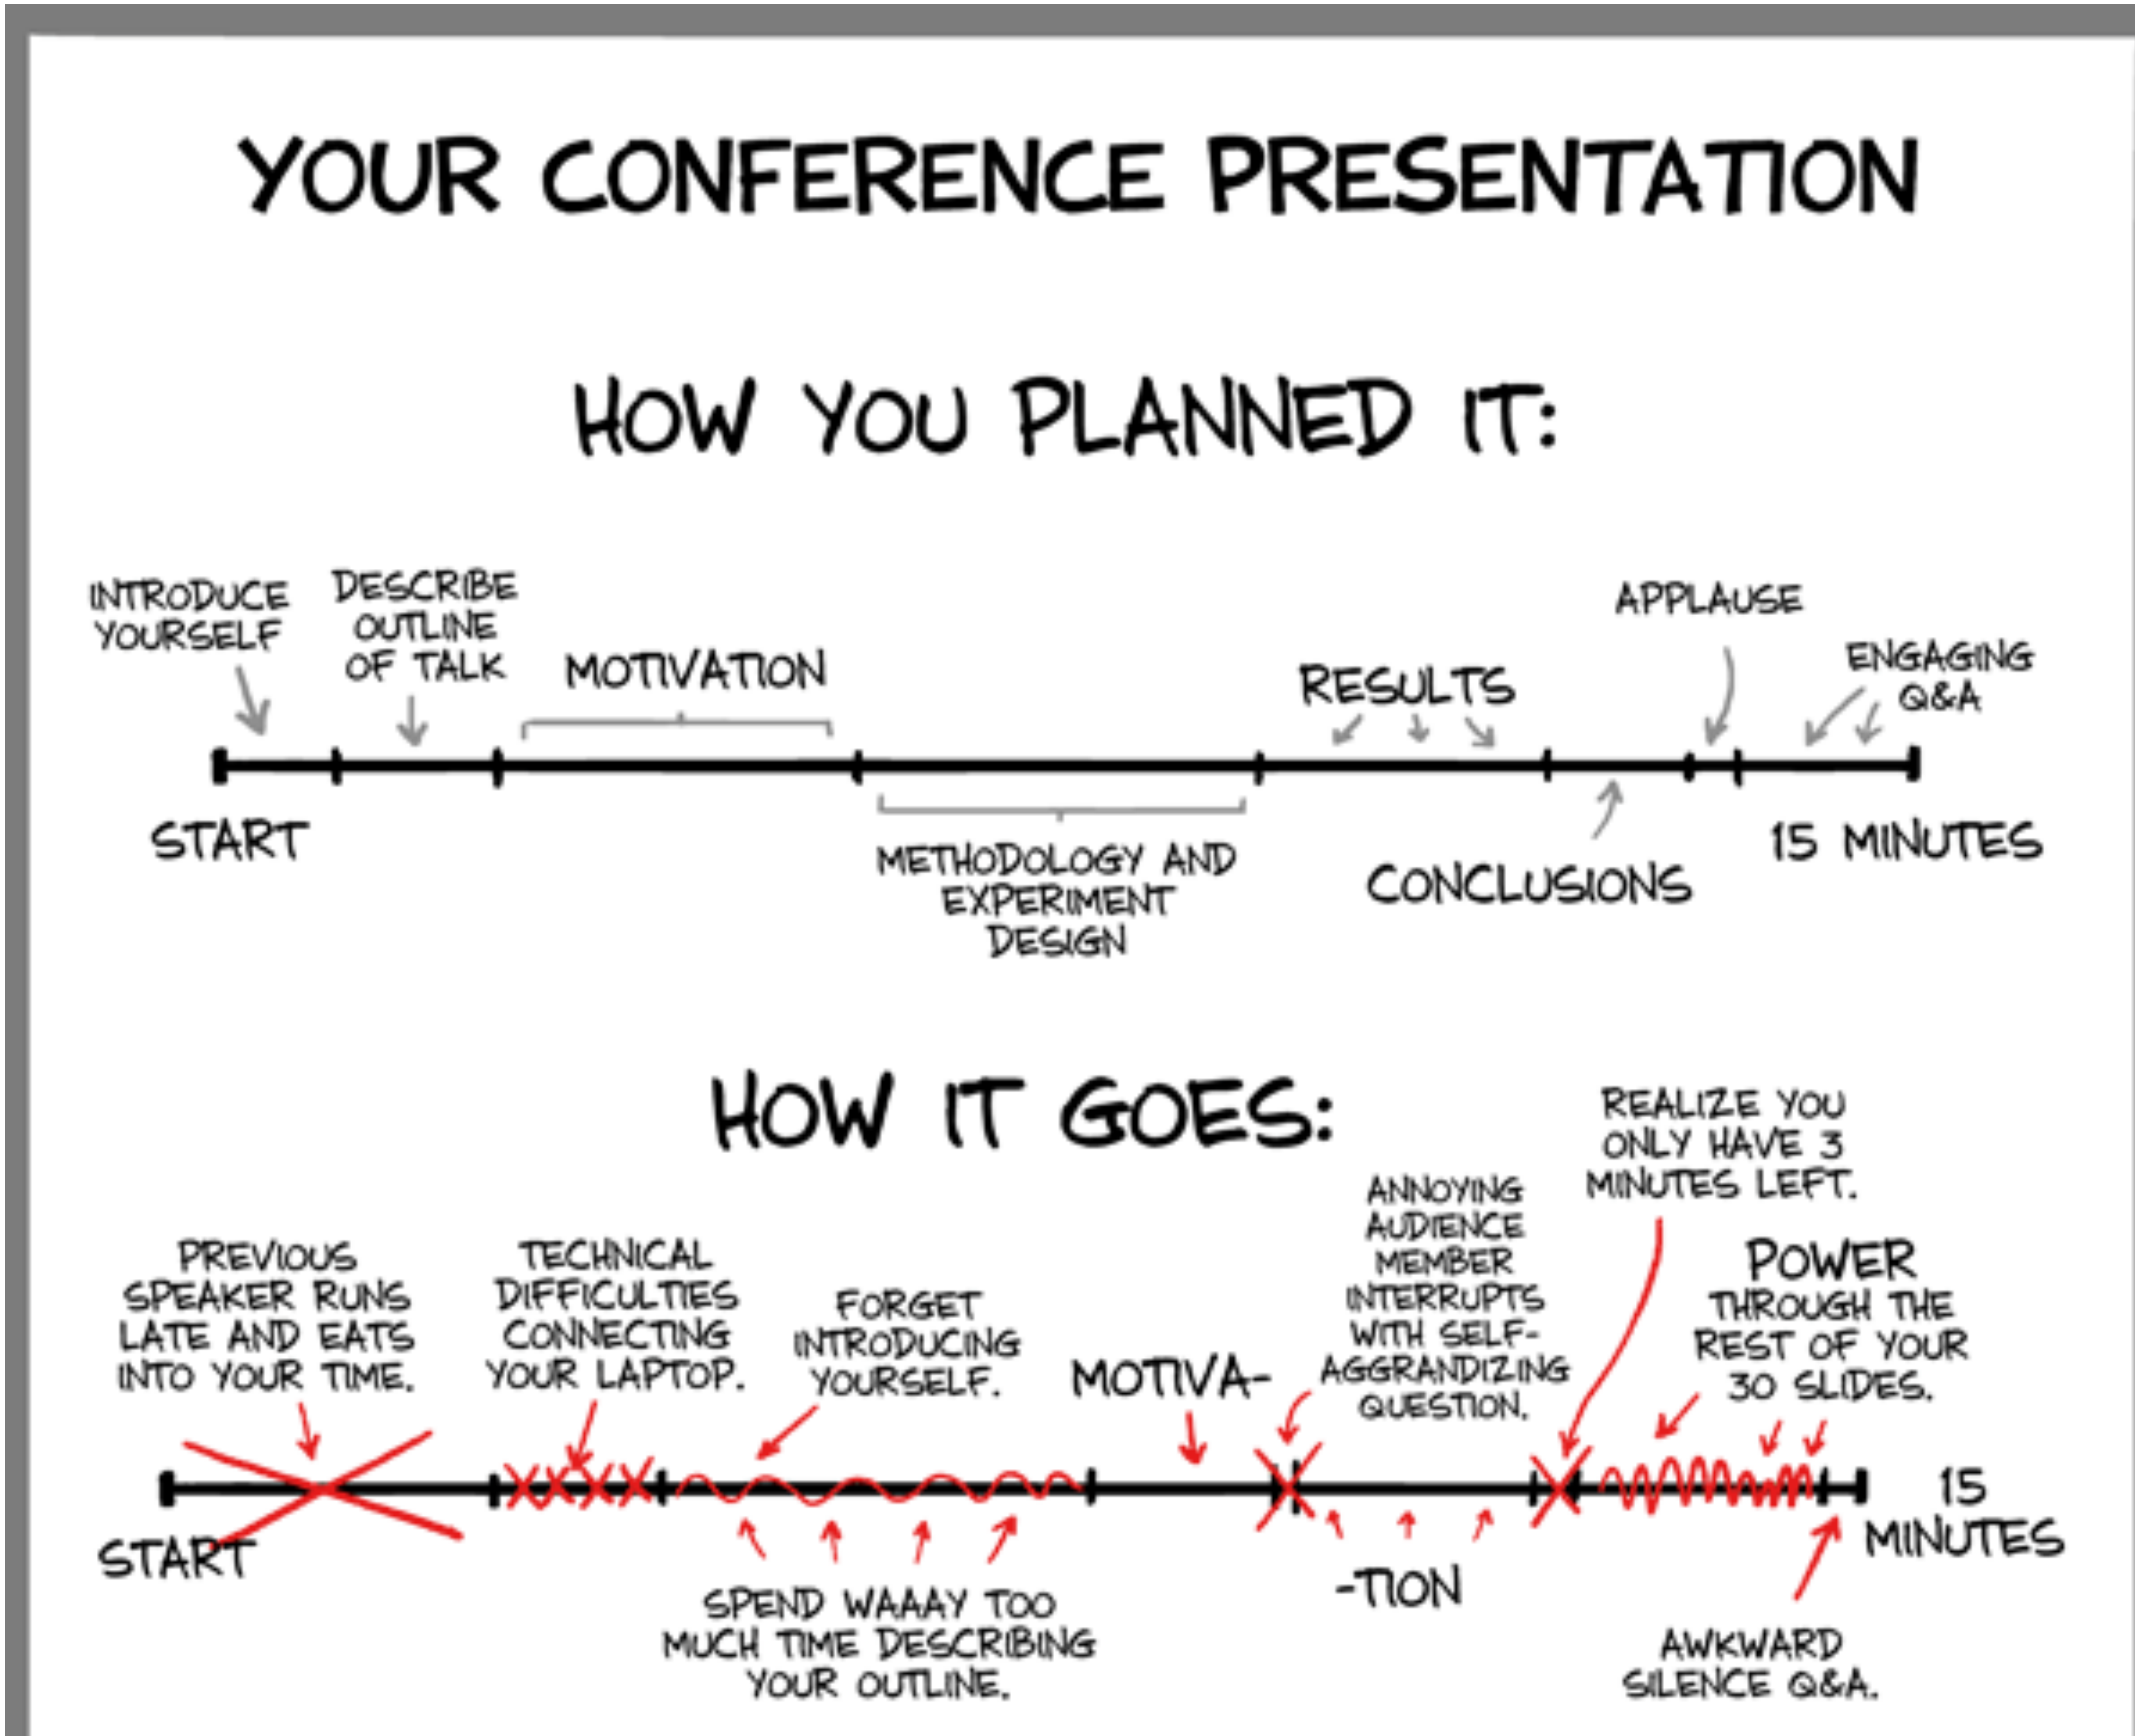

http://www.phdcomics.com/comics.php?f=1553

Final report

- 8 pages, single spaced.
- Complete description of work undertaken
    - Data collection
    - Methods
    - Experimental details
    - Comparison with past work
    - Analysis
- See many of the papers we've read this semester for examples.

Final report

- Clarity.   For the reasonably well-prepared reader, is it clear what was done and why? Is the paper well-written and well-structured?
- Originality. How original is the approach or problem presented in this paper? Does this paper break new ground in topic, methodology, or content? How exciting and innovative is the research it describes?
- Soundness. Is the technical approach sound and well-chosen? Second, can one trust the claims of the paper -- are they supported by proper experiments, proofs, or other argumentation?
- Substance. Does this paper have enough substance, or would it benefit from more ideas or results? Do the authors identify potential limitations of their work?
- Evaluation. To what extent has the application or tool been tested and evaluated? Does this paper present a compelling argument for
- Meaningful comparison. Do the authors make clear where the presented system sits with respect to existing literature? Are the references adequate? Are the benefits of the system/application well-supported and are the limitations identified?
- Impact. How significant is the work described? Will novel aspects of the system result in other researchers adopting the approach in their own work?