CS 6120/4120: Natural Language Processing

Developing NLP measures for social science applications

Sarah Shugars

Northeastern University

shugars.s@northeastern.edu

she/her

What does “doing NLP” actually look like?

🤔 🤔 🤔

NLP in Practice

Some Phenomena

+

Observed in Text

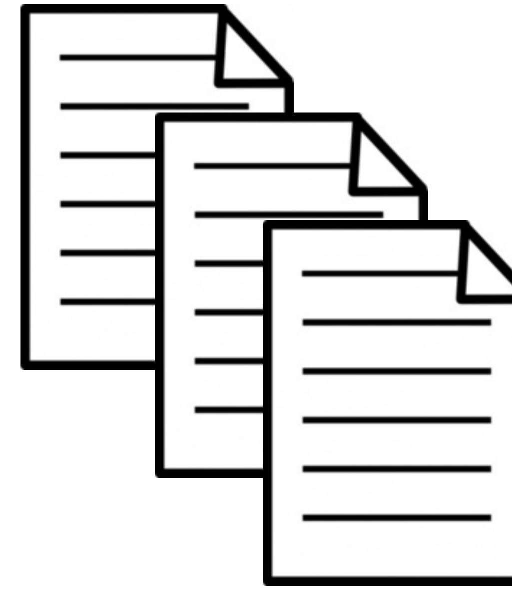

=

Interesting insights

🧐

NLP Applications

- (How) Are people of different genders described differently?
- What discursive moves influence debate outcomes?
- How do people talk about their political opinions?

Common NLP Approaches

(How) Are people of different genders described differently?

Common NLP Approaches

(How) Are people of different genders described differently?

Descriptive

- Top words

Managers Use More Positive Words to Describe Men in Performance Reviews and More Negative Ones to Describe Women

IN DESCENDING ORDER OF RELATIVE FREQUENCY

Words used to describe men

POSITIVE	NEGATIVE
Analytical	
Competent	
Athletic	
Dependable	Arrogant
Confident	
Versatile	
Articulate	
Level-headed	
	Irresponsible
Logical	
Practical	

Words used to describe women

POSITIVE	NEGATIVE
Compassionate	
	Inept
Enthusiastic	Selfish
Energetic	Frivolous
	Passive
Organized	Scattered
	Opportunistic
	Gossip
	Excitable
	Vain
	Panicky
	Temperamental
	Indecisive

SOURCE AN ANALYSIS OF 81,000 PERFORMANCE EVALUATIONS, DAVID G. SMITH ET AL., 2018

Common NLP Approaches

(How) Are people of different genders described differently?

Descriptive

- Top words
- Word distributions

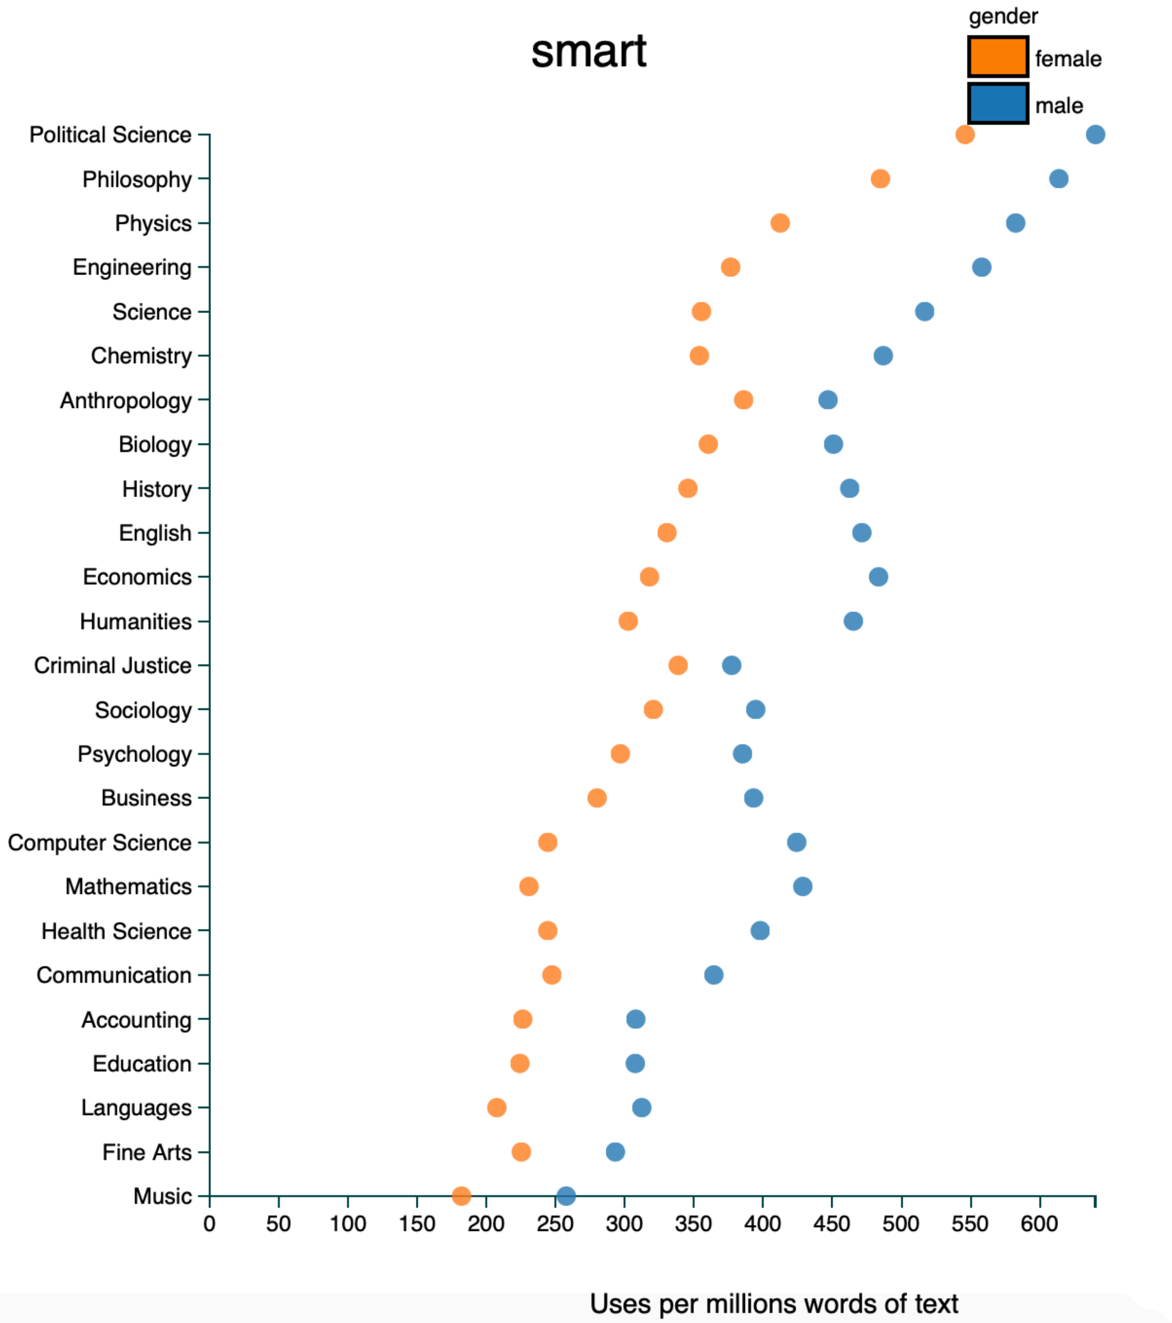

http://benschmidt.org/profGender

Common NLP Approaches

(How) Are people of different genders described differently?

Descriptive

- Top words
- Word distributions

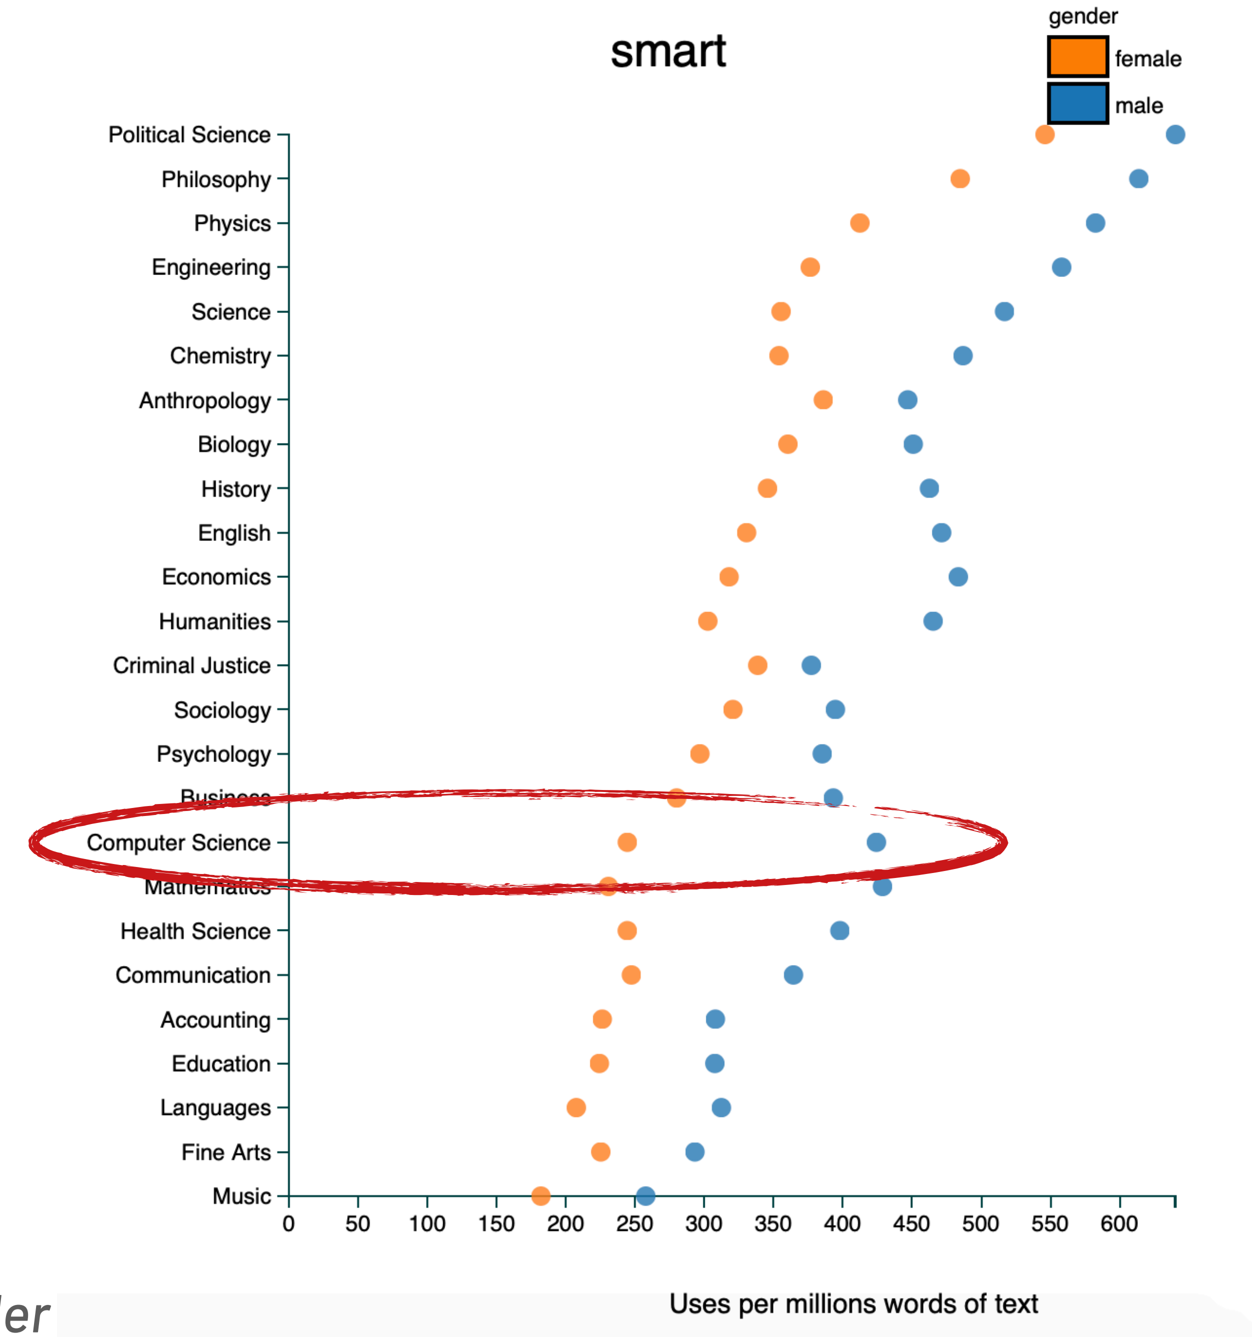

http://benschmidt.org/profGender

Common NLP Approaches

(How) Are people of different genders described differently?

Predictive

- Given that we see the word 'smart' what is the probability the sentence describes a woman?

Common NLP Approaches

(How) Are people of different genders described differently?

Predictive

- Given that we see the word 'smart' what is the probability the sentence describes a woman?

  Or, more generally:

- Given that we see features **X**, can we classify this text as relating to gender $\mathcal{G}$?

Common NLP Approaches

Given that we see features **X**, can we classify this text as relating to gender $\mathcal{G}$?

$$P(\mathcal{G} = 1) = \Phi\left(\beta\,\mathbf{X}\right)$$

Common NLP Approaches

Given that we see features **X**, can we classify this text as relating to gender $\mathcal{G}$?

$$P(\mathcal{G} = 1) \; = \; \Phi\left(\beta \mathbf{X} + \eta_f + \epsilon\right)$$

Common NLP Approaches

Given that we see features **X**, can we classify this text as relating to gender $\mathcal{G}$?

$$\underbrace{P(\mathcal{G} = 1)}_{\text{Probability gender = 1}} = \Phi\Big(\underbrace{\beta\,\mathbf{X}}_{\text{Features }\mathbf{X}\text{ with weights }\beta} + \eta_f + \epsilon\Big)$$

Error

Fixed effect for academic field

Common NLP Approaches

Given that we see features **X**, can we classify this text as relating to gender $\mathcal{G}$?

NLP goes here

$$\underbrace{P(\mathcal{G}=1)}_{\substack{\text{Probability} \\ \text{gender = 1}}} = \Phi\Big(\underbrace{\beta\,\mathbf{X}}_{\substack{\text{Features } \mathbf{X} \\ \text{with weights } \beta}} + \eta_f + \epsilon\Big)$$

Error

Fixed effect for academic field

Common NLP Approaches

What is $\beta \mathbf{X}$?

$$\beta \mathbf{X} = \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n$$

Common NLP Approaches

What is $\beta \mathbf{X}$?

$$\beta \mathbf{X} = \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n$$

Features

Common NLP Approaches

What is $\beta \mathbf{X}$?

$$\beta \mathbf{X} = \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n$$

Features

Weights

The Challenge:

How do we turn

into

$$x_1, x_2, \ldots, x_n$$

?

Case Study 1

Why Keep Arguing? Predicting Engagement in Political Conversations Online

Sarah Shugars, Network Science Institute

Nick Beauchamp, Department of Political Science

Why do people bother arguing online?

Online engagement is driven by individual, conversation, and content features

Online engagement is driven by individual, conversation, and content features

$$P(\underbrace{T_{ijt} = 1)}_{\text{Probability of engagement}} \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Individual features

Probability
of engagement

Some function of...

Online engagement is driven by individual, conversation, and content features

$$P(T_{ijt} = 1) \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Probability of engagement

Some function of...

Conversation features

Content features

Online engagement is driven by individual, conversation, and content features

$$\underbrace{P(T_{ijt} = 1)}_{\text{Probability of engagement}} \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Probability of engagement

Some function of...

Content features

Hypotheses

Individual

- Baseline activity
- Popularity

$$P(T_{ijt} = 1) \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Conversations

- Popularity
- Recent engagement

$$P(T_{ijt} = 1) \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Hypotheses

Content

- Emotionally extreme users more likely to re-engage
- Emotionally extreme tweets more likely to receive a response
- Topic effects

$$P(T_{ijt} = 1) \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Data

- 7053 conversations
- 63,671 tweets
- Keyword "Trump"
- October 2017

$$P(T_{ijt} = 1) \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Data

- 7053 conversations
- 63,671 tweets
- 1,016,49 observations

$$P(T_{ijt} = 1) \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

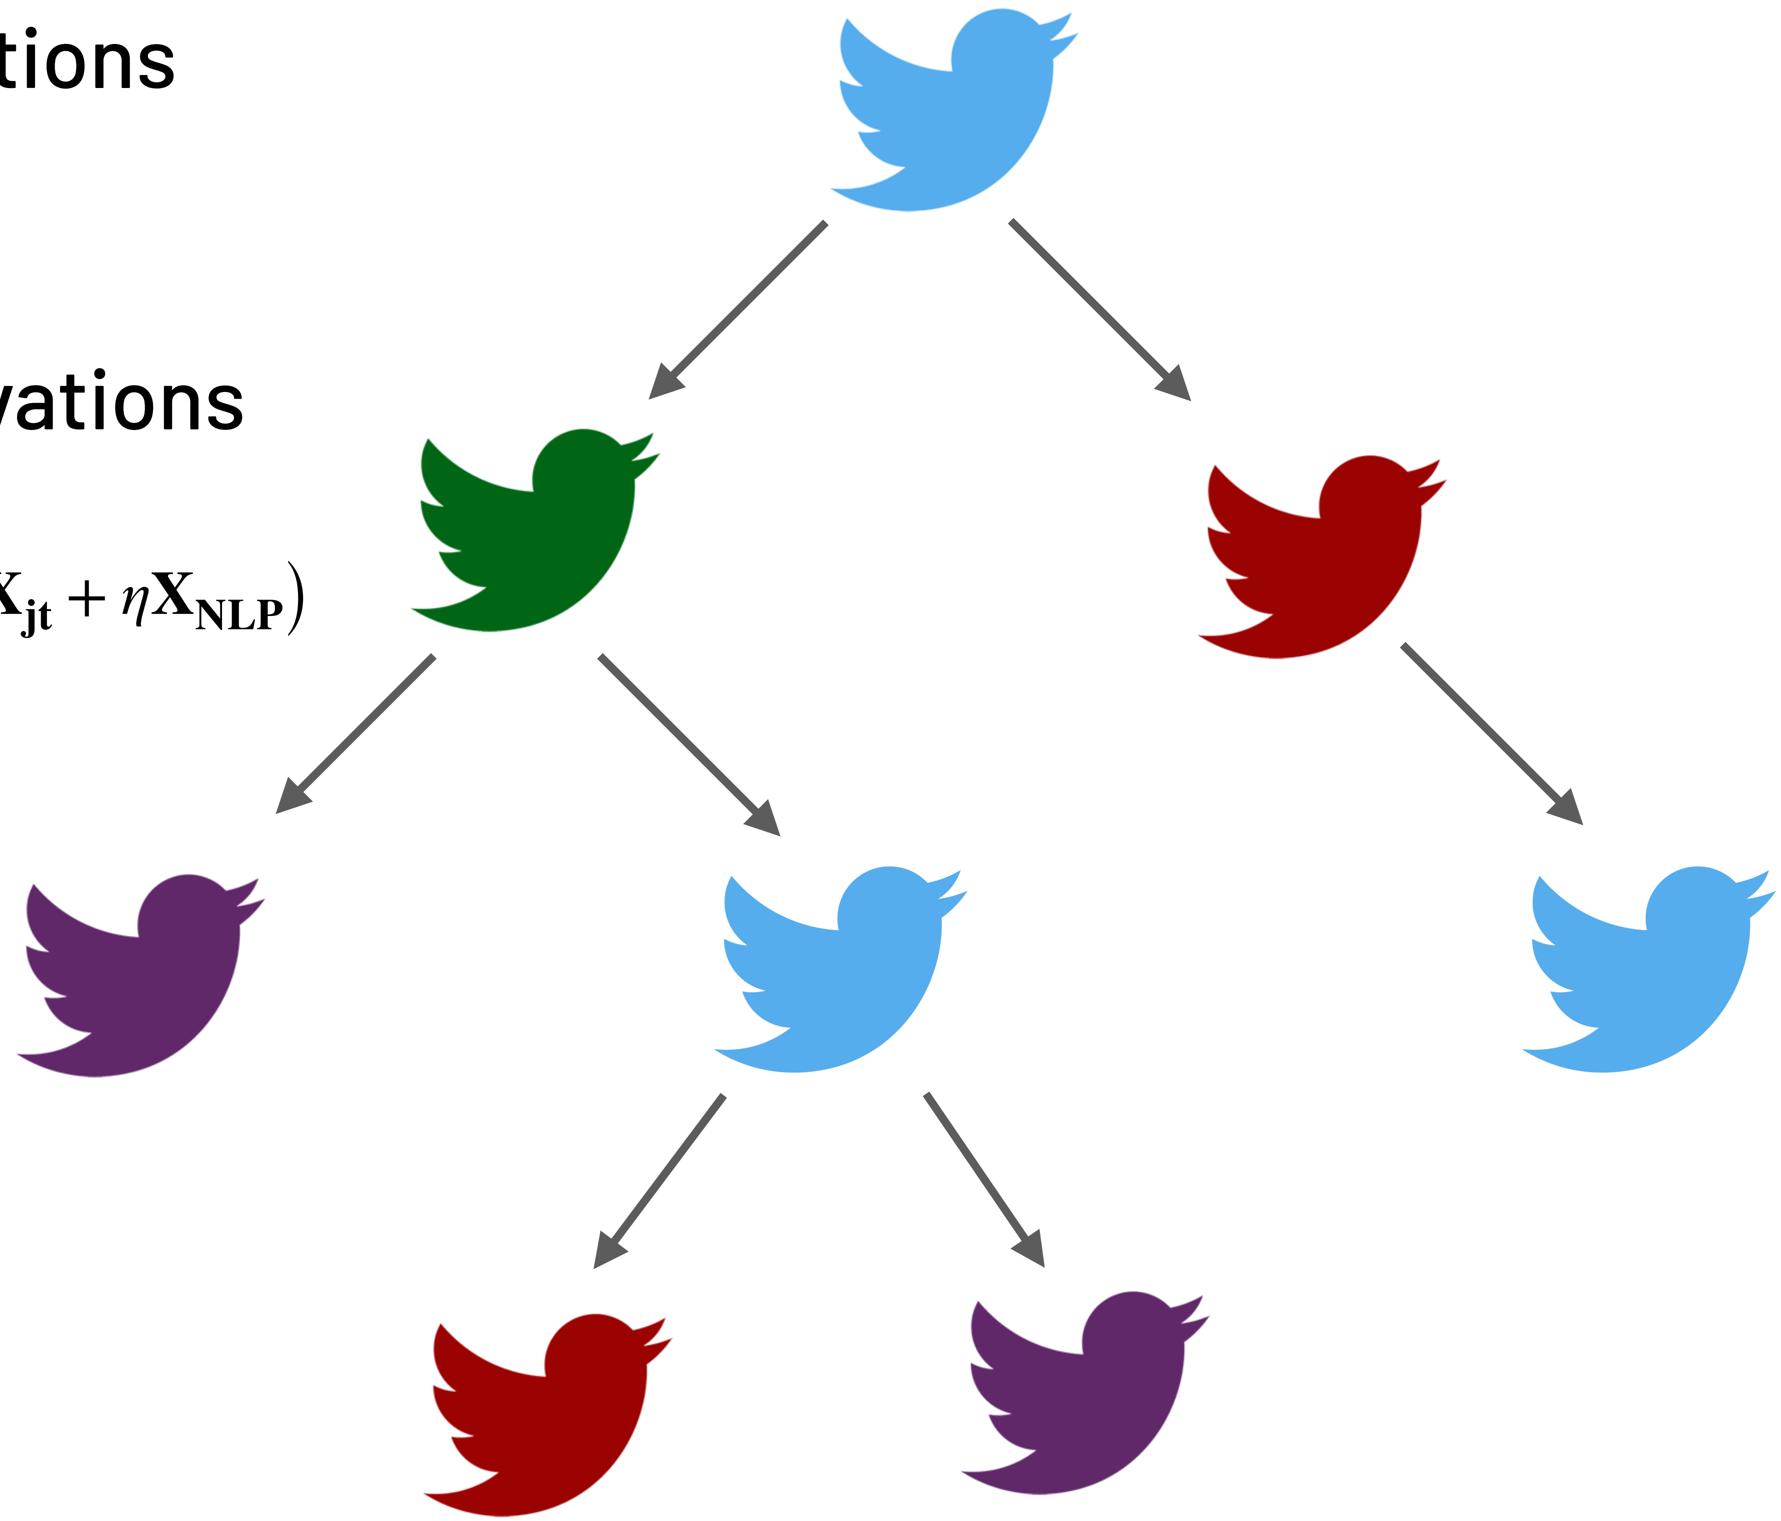

Conversational data

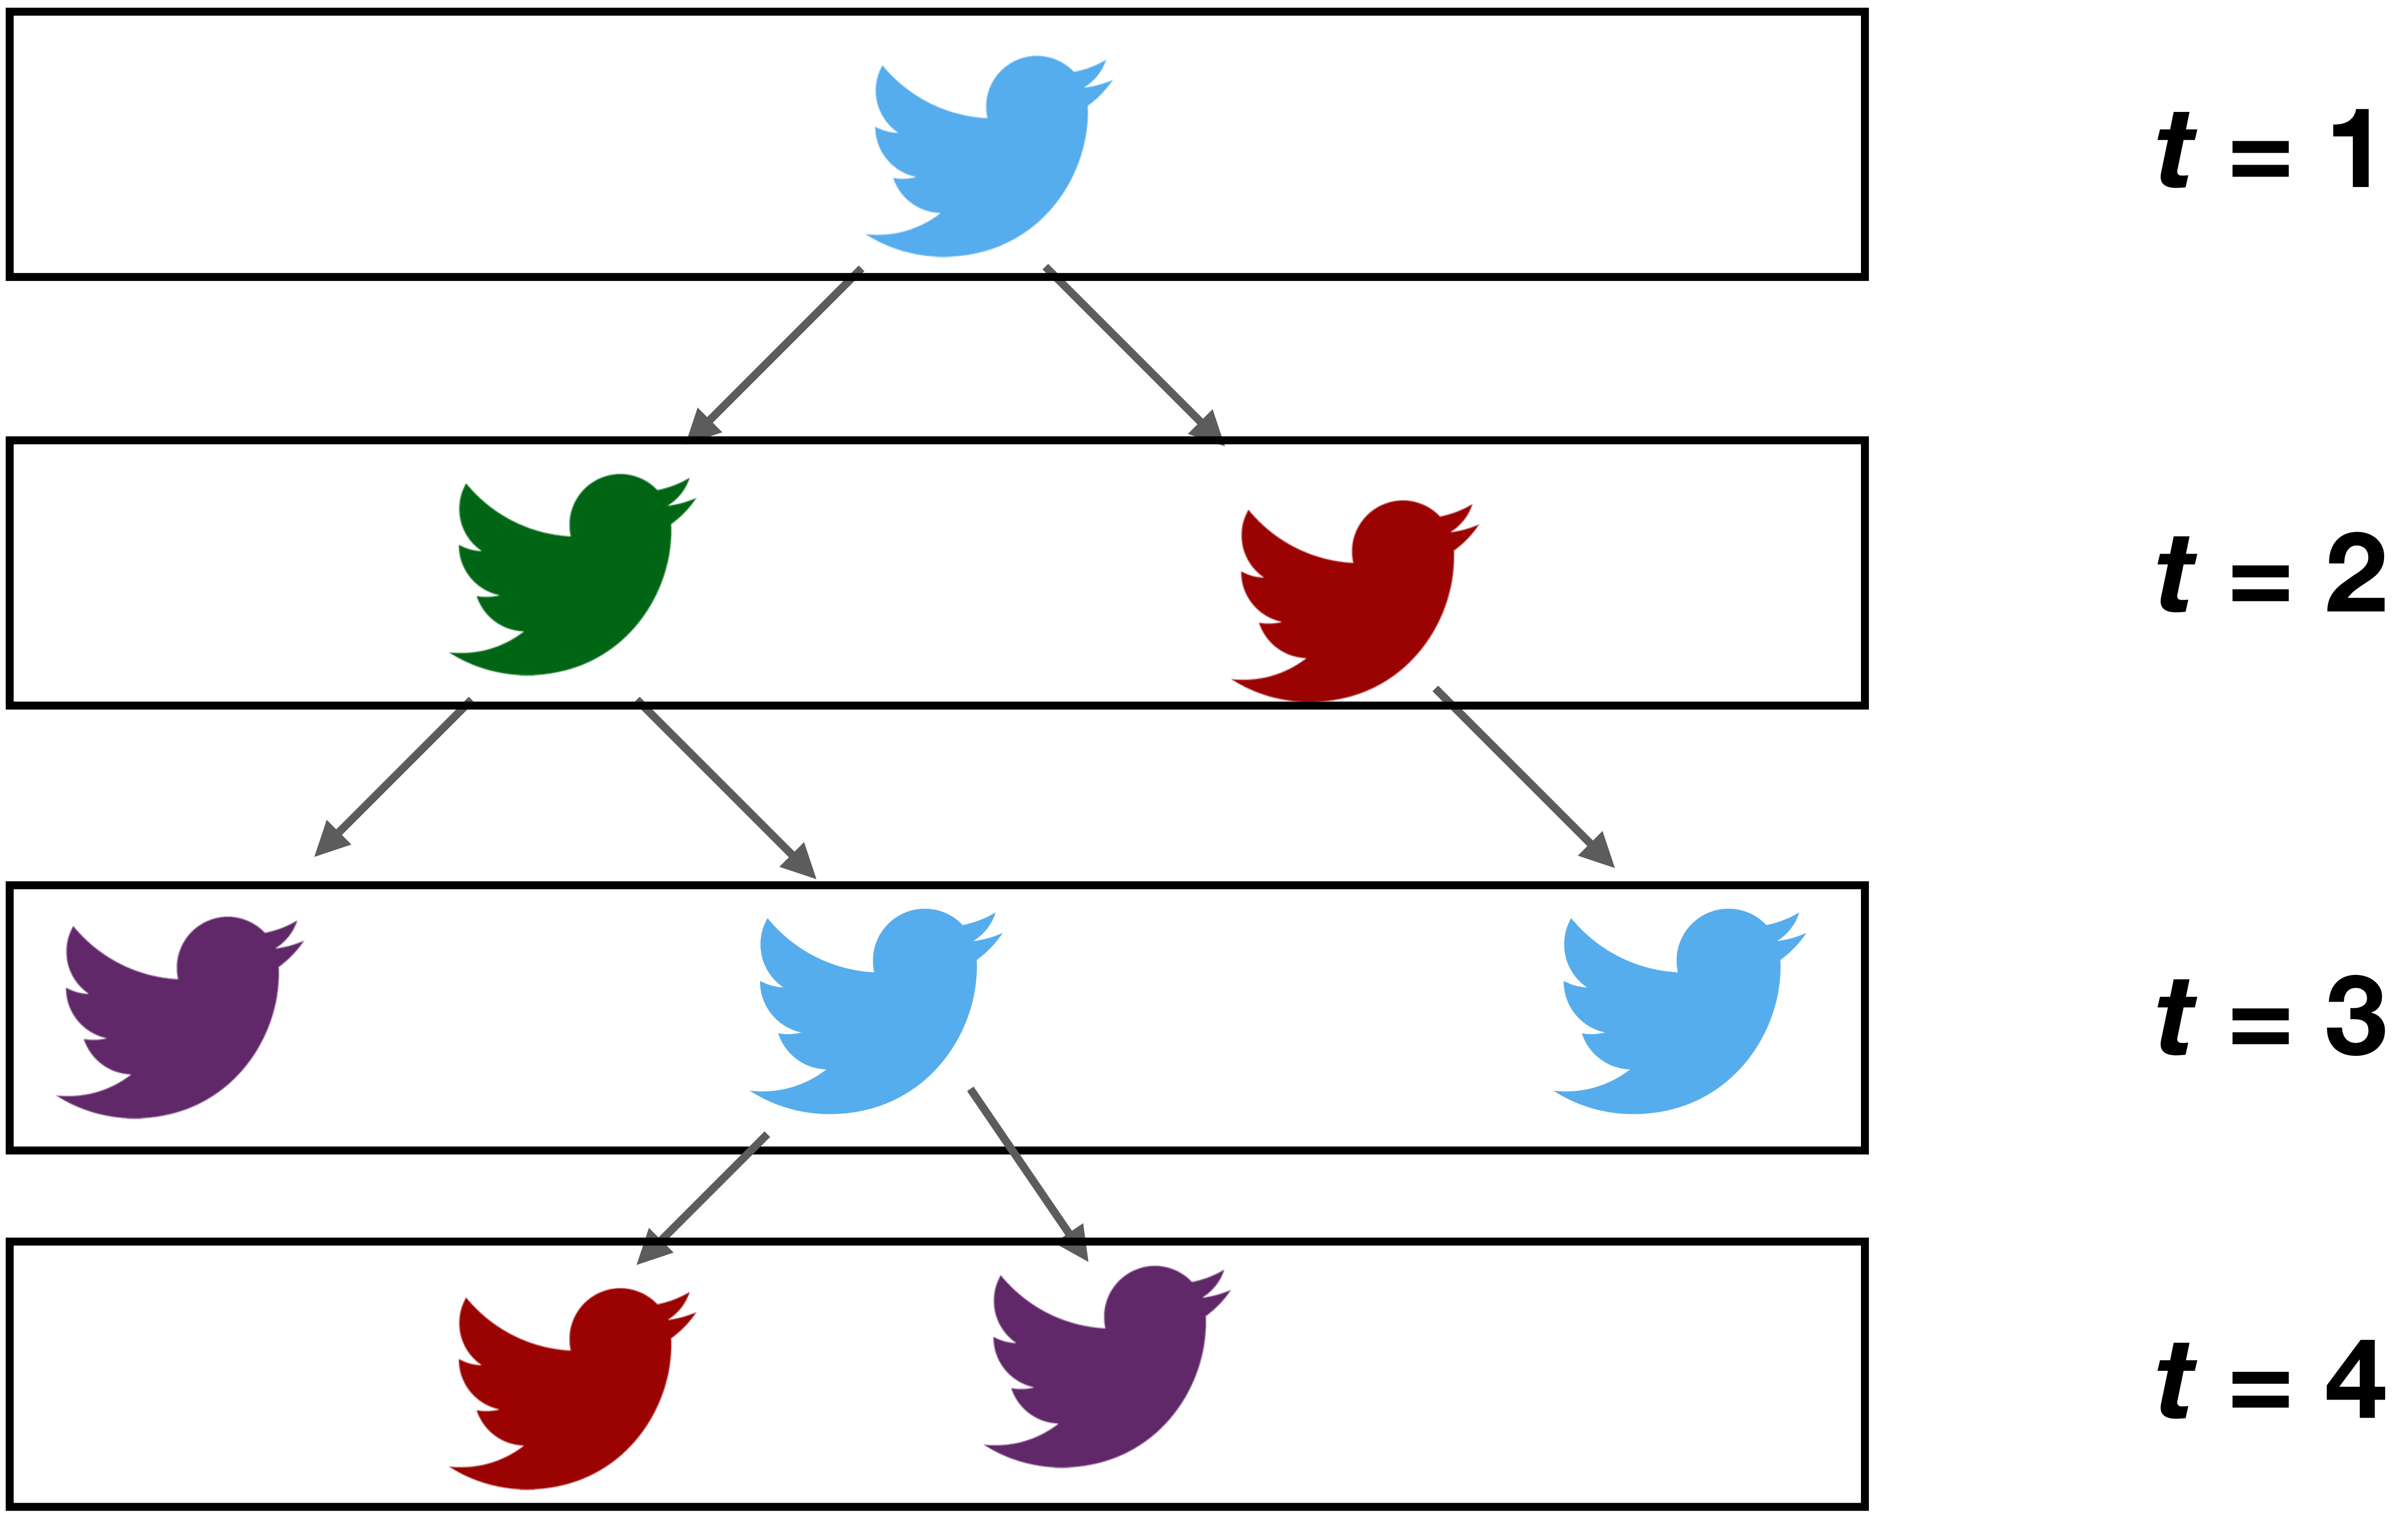

Conversational data

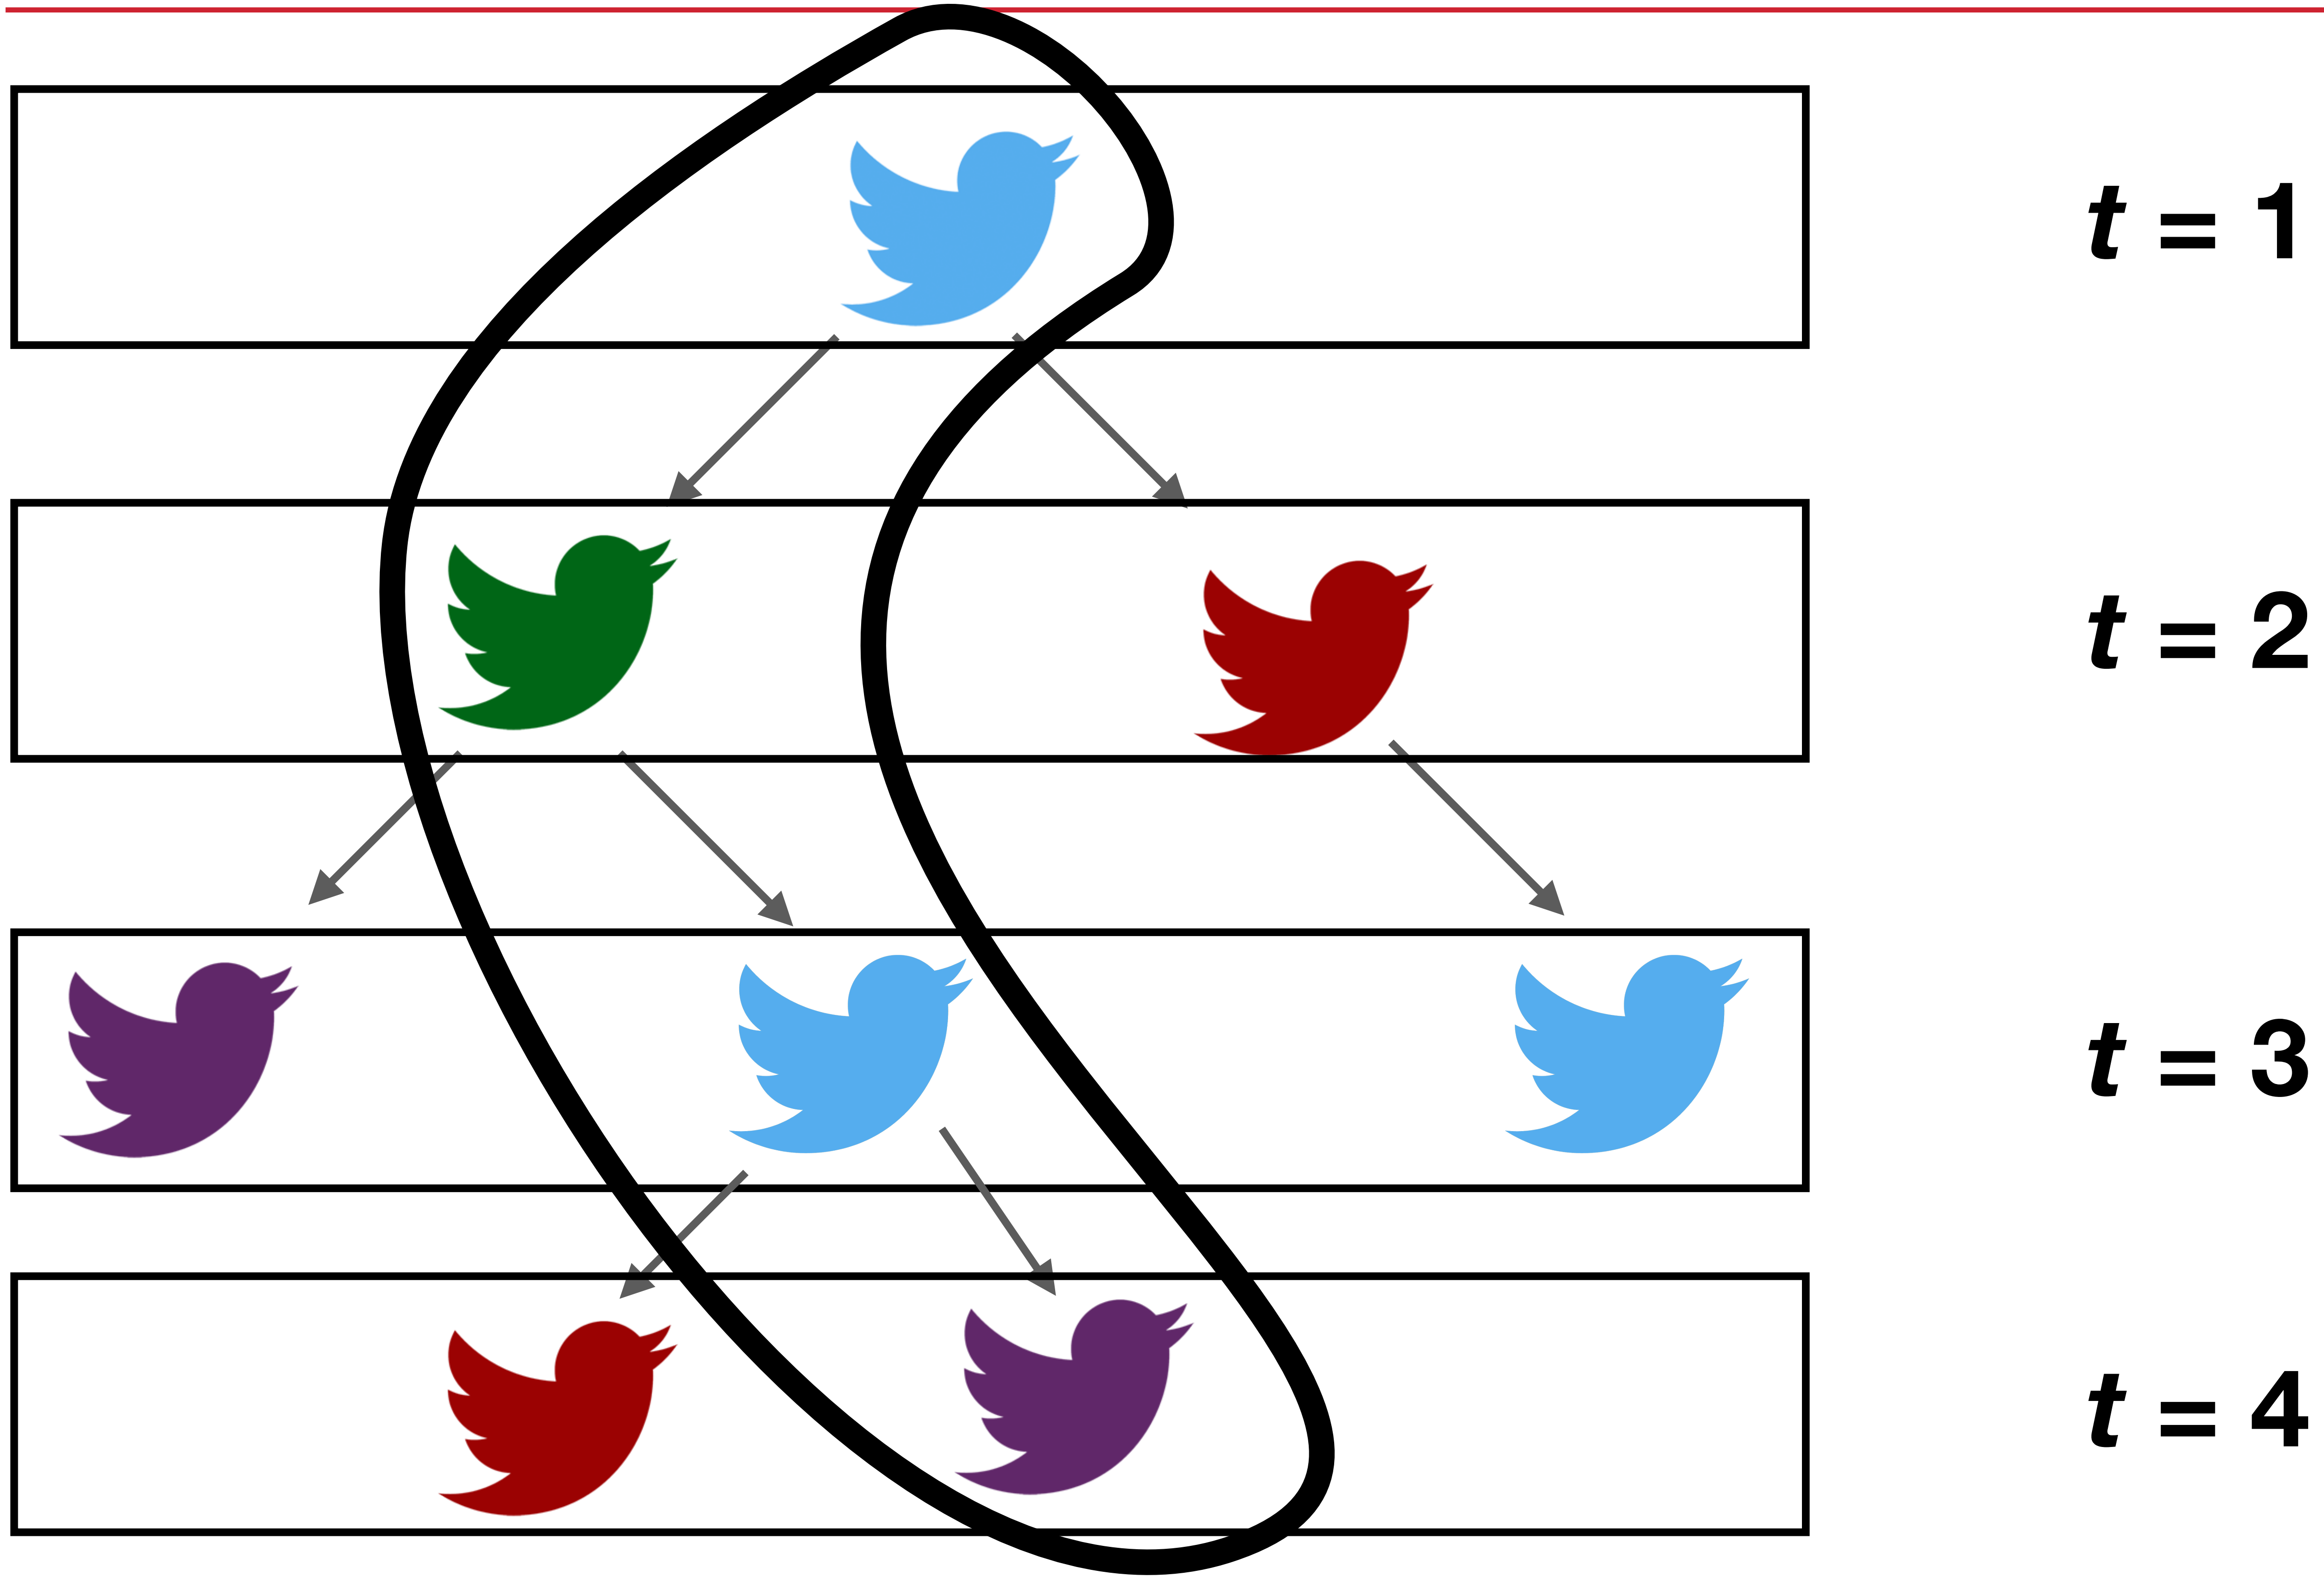

Conversational data

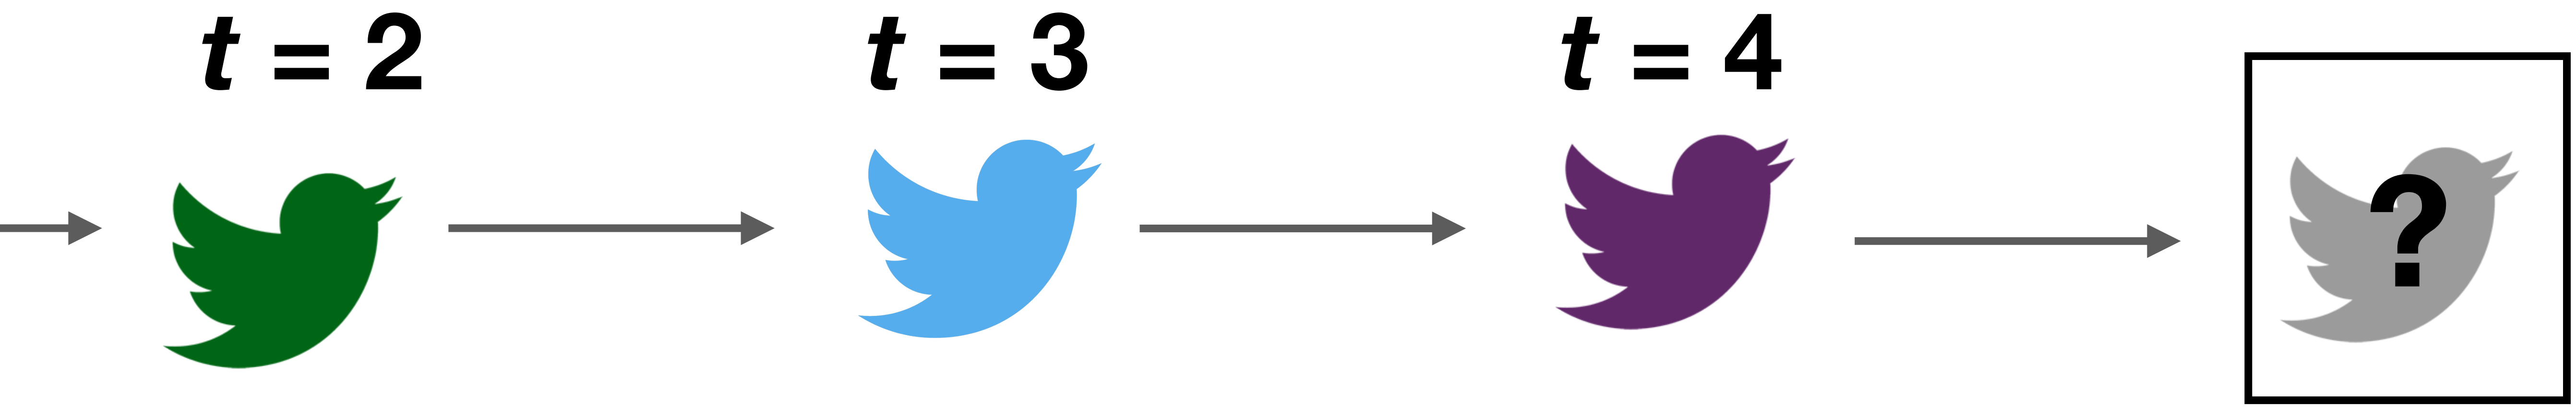

For time $t > 2$, how does ideology influence who will **remain active** in a conversation?

Conversational data

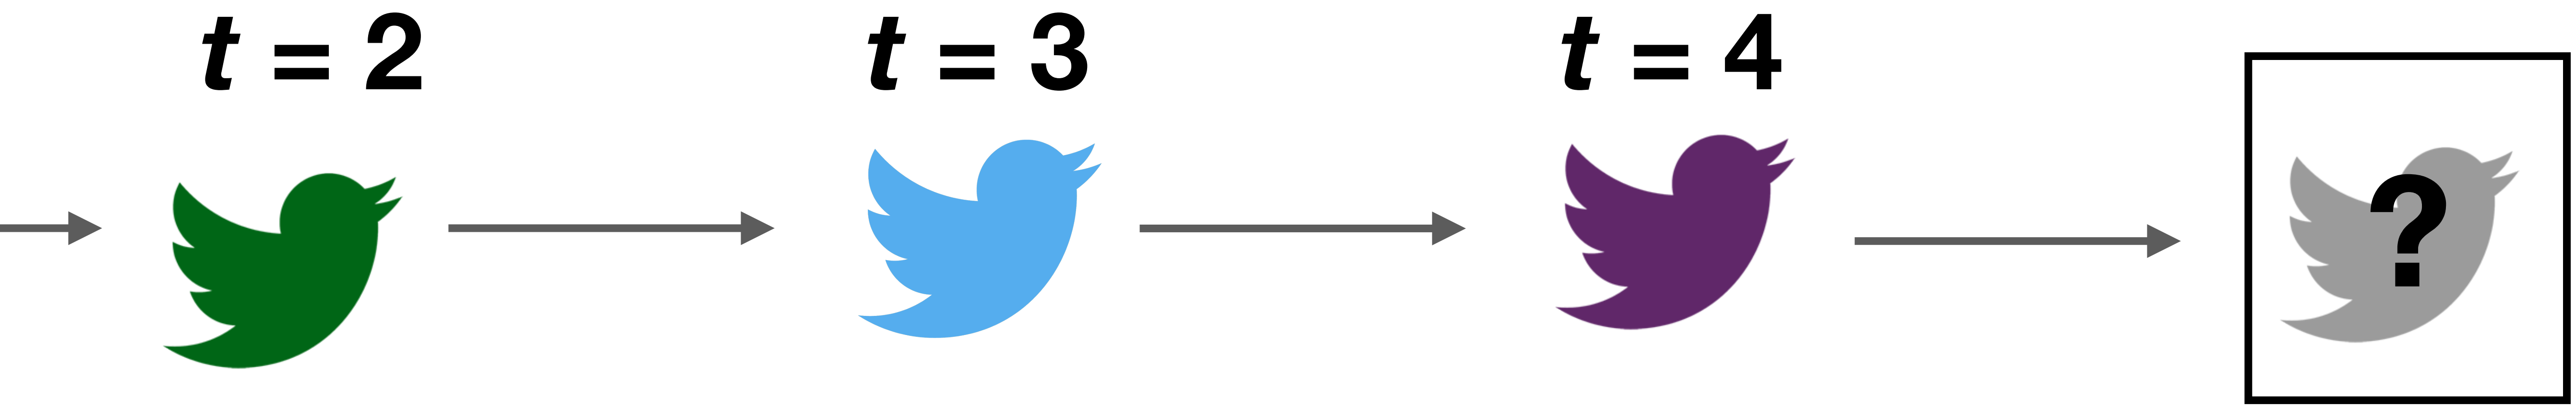

Candidates for re-entry:

Conversational data

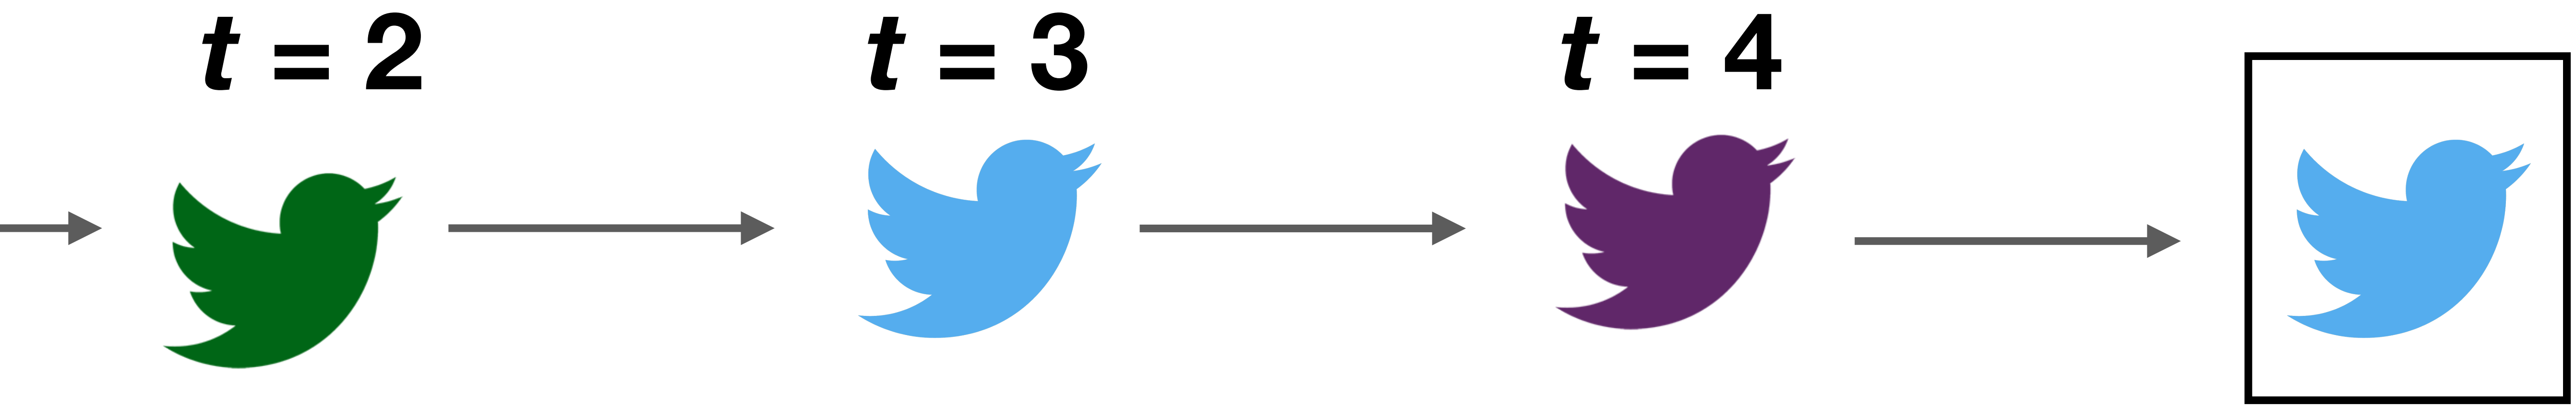

Candidates for re-entry:

Observed outcome:

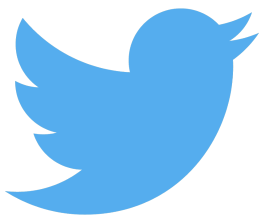

1

0

Conversational data

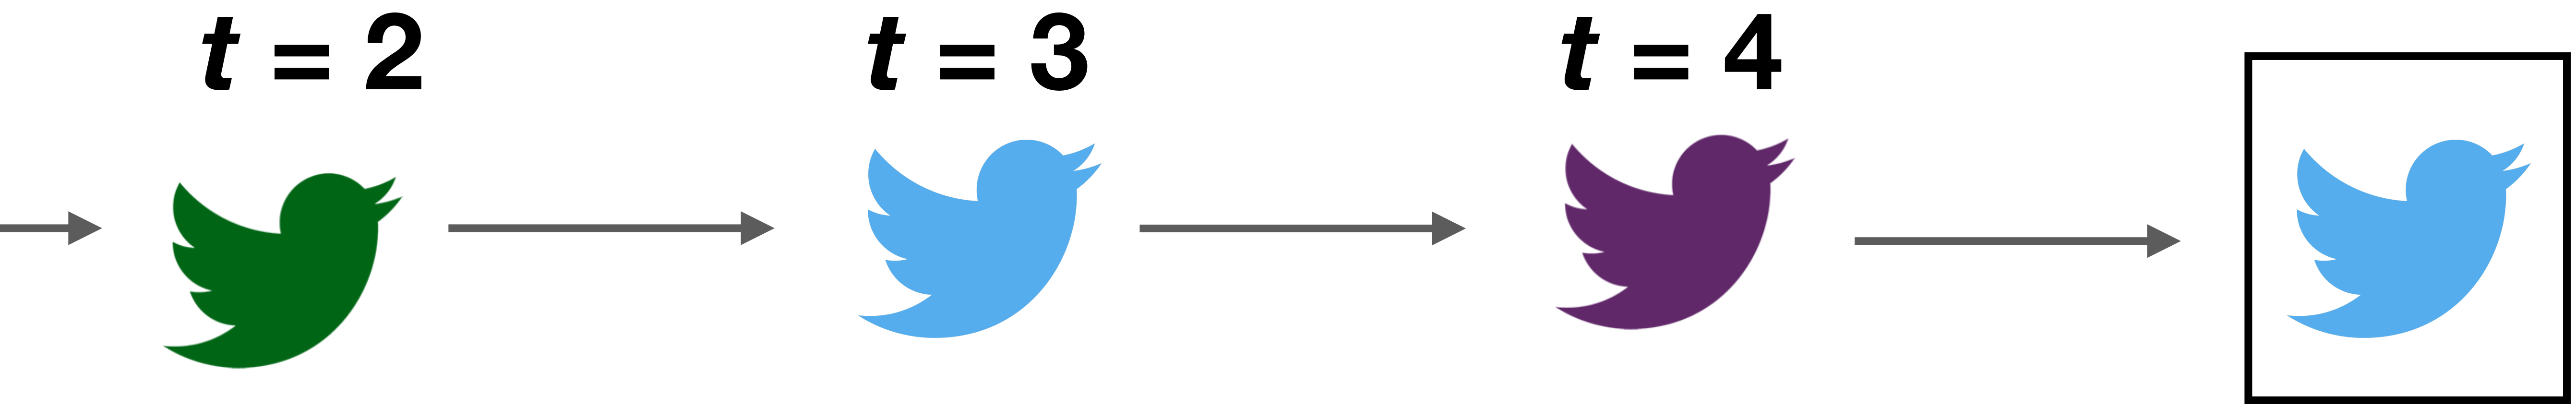

Candidates for re-entry:	Observed outcome:
	**1**
	**0**

Function of: individual, conversation, and content features

Conversational data

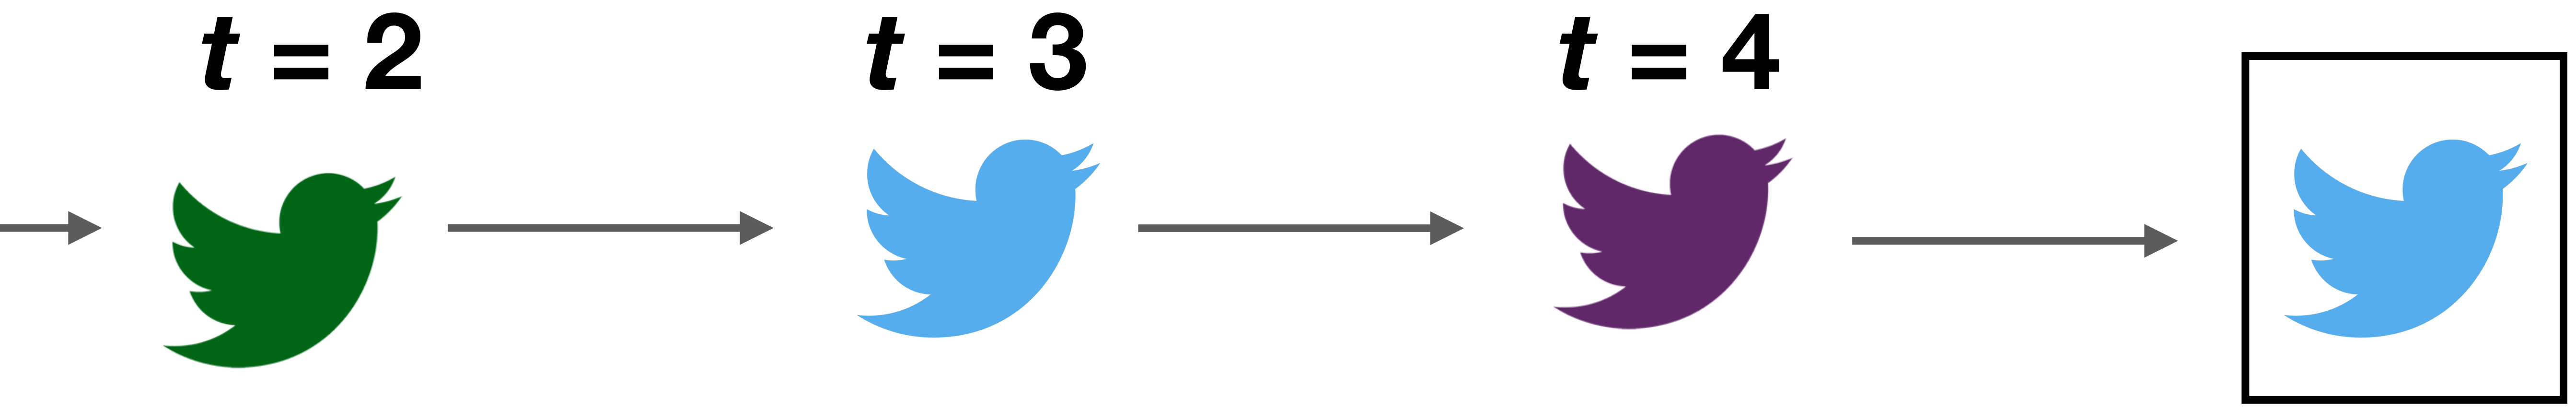

$$P(T_{ijt} = 1) \sim \Phi\left(\beta \mathbf{X_i} + \gamma \mathbf{X_{jt}} + \eta \mathbf{X_{NLP}}\right)$$

Findings

- 110,035 observations
- 89% of observations are 0 (non-response)
- Achieve 94% accuracy with logistic regression
- 98% accuracy with SVM

Response predictors: Individual

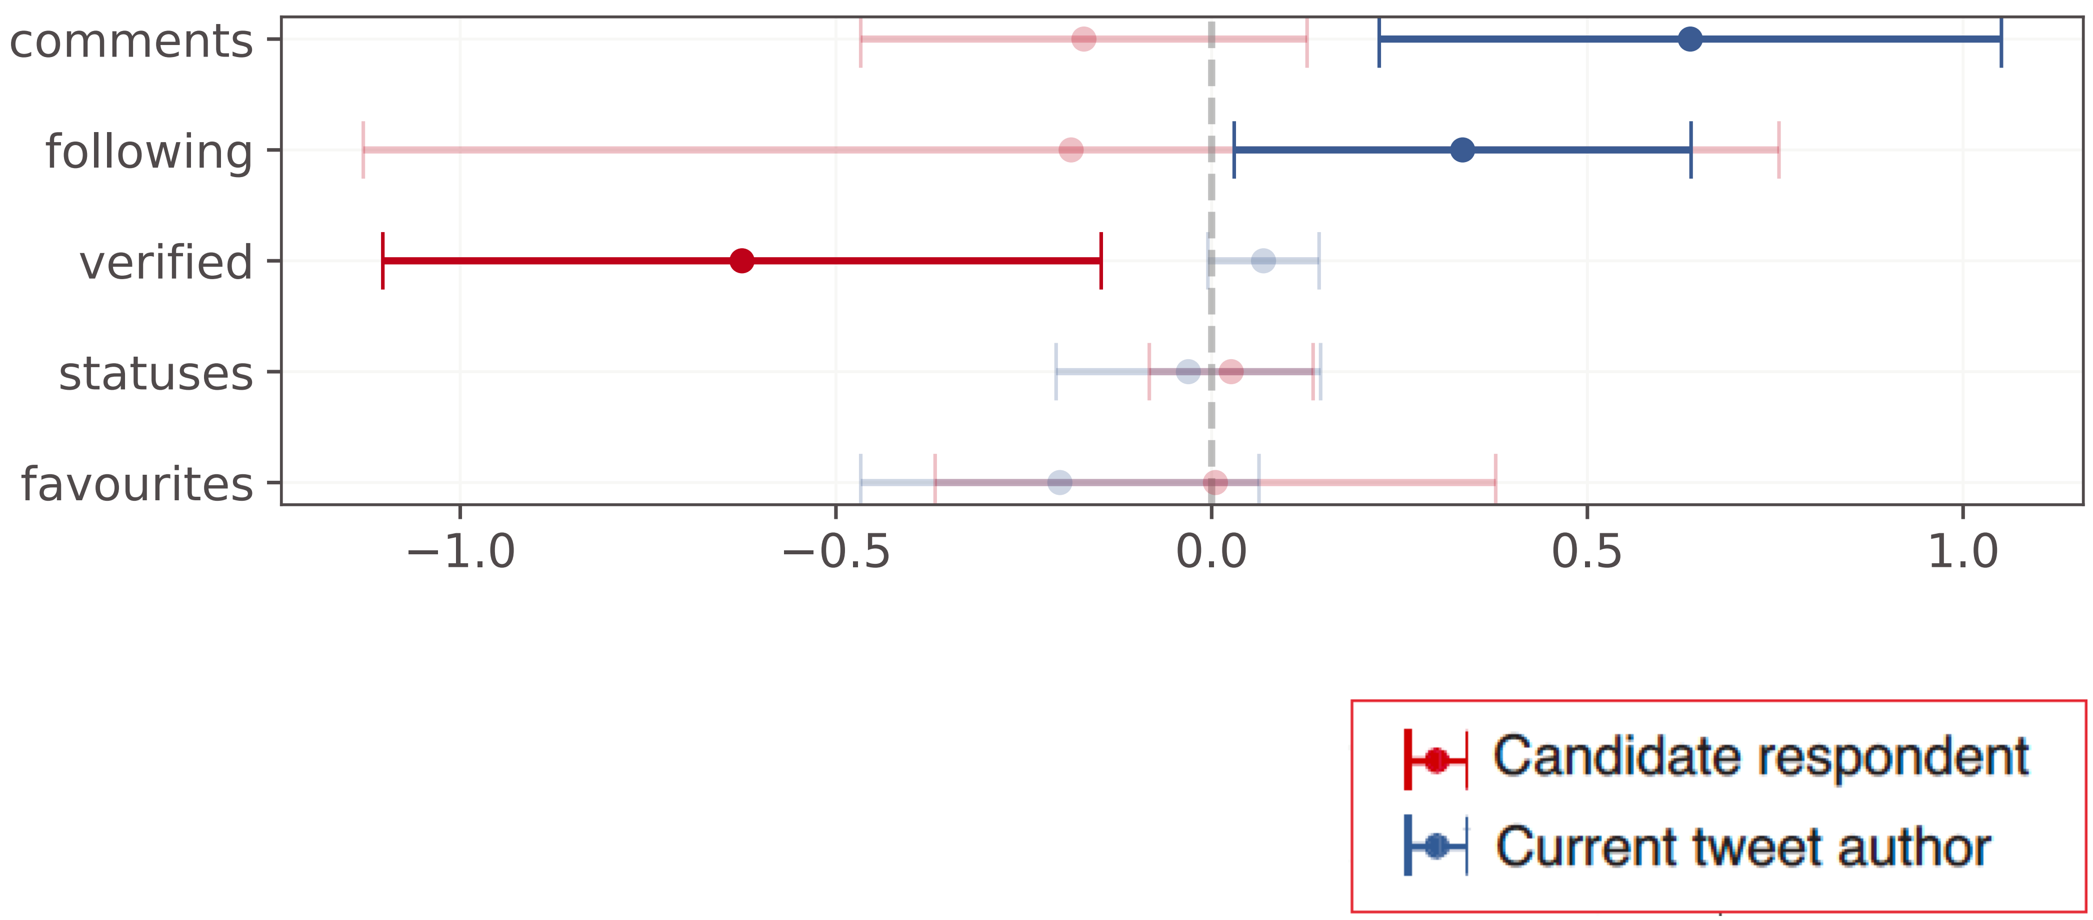

Response predictors: Conversation

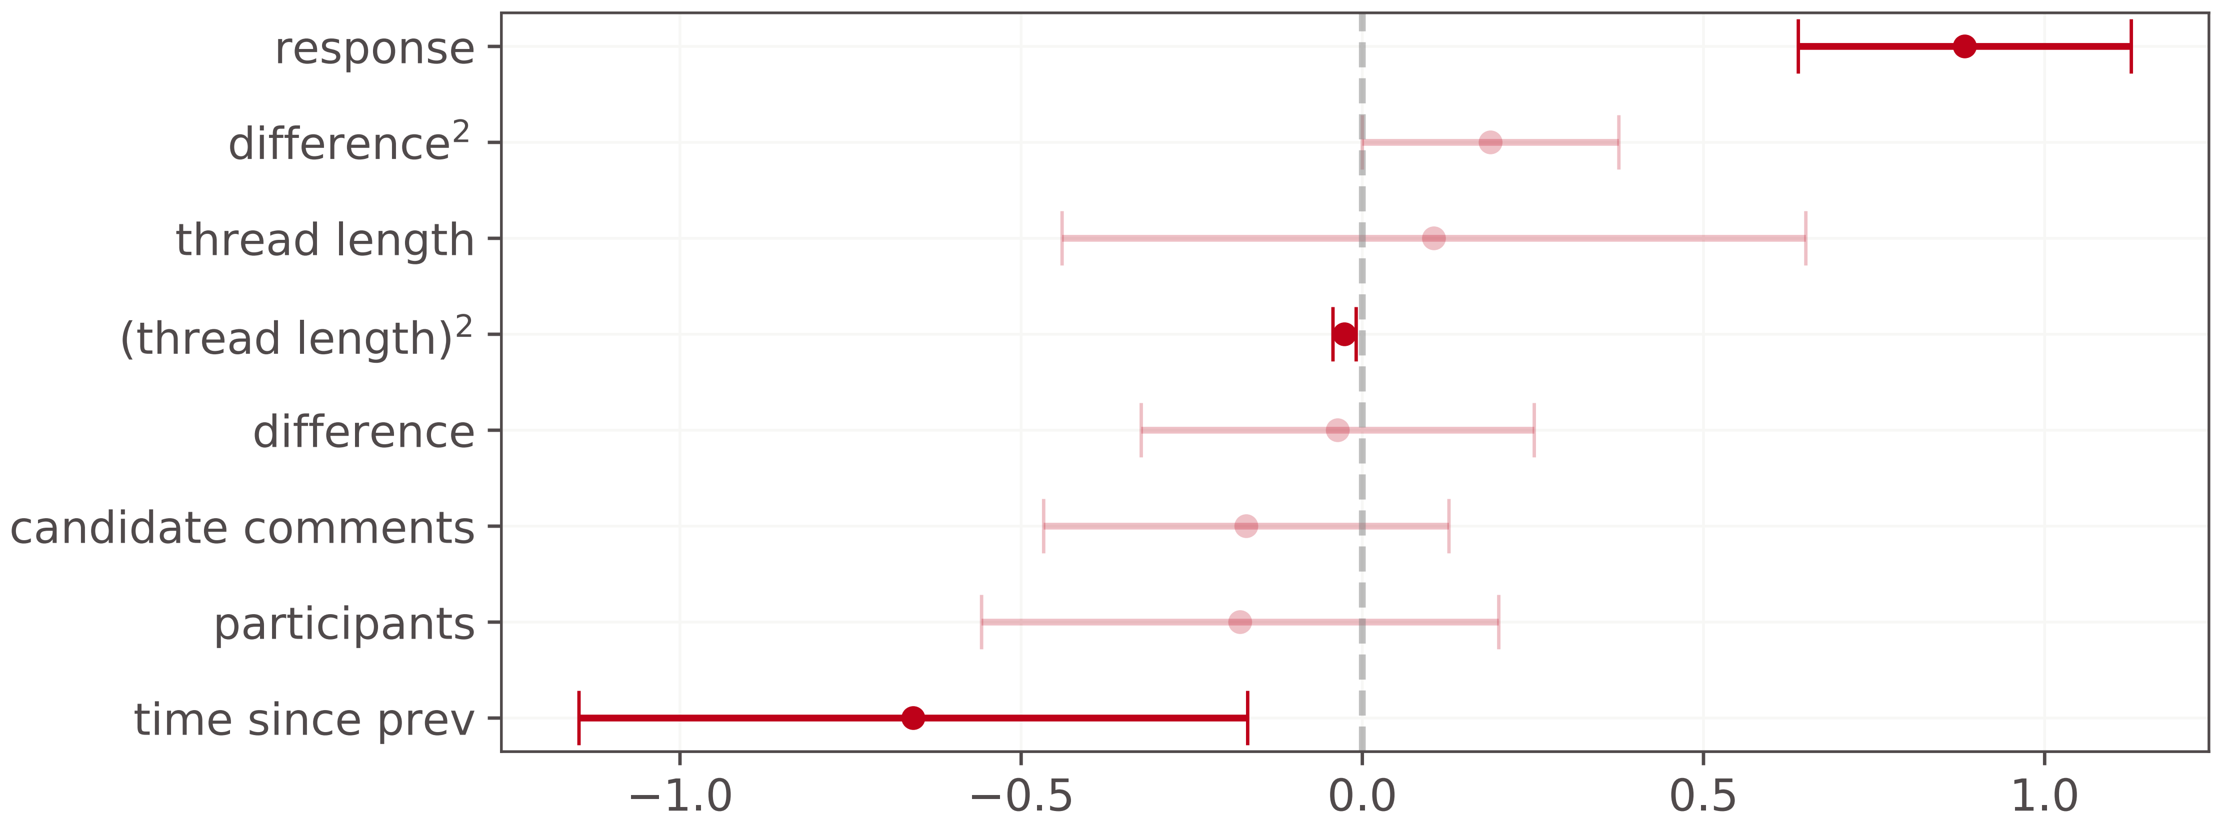

Response predictors: Content

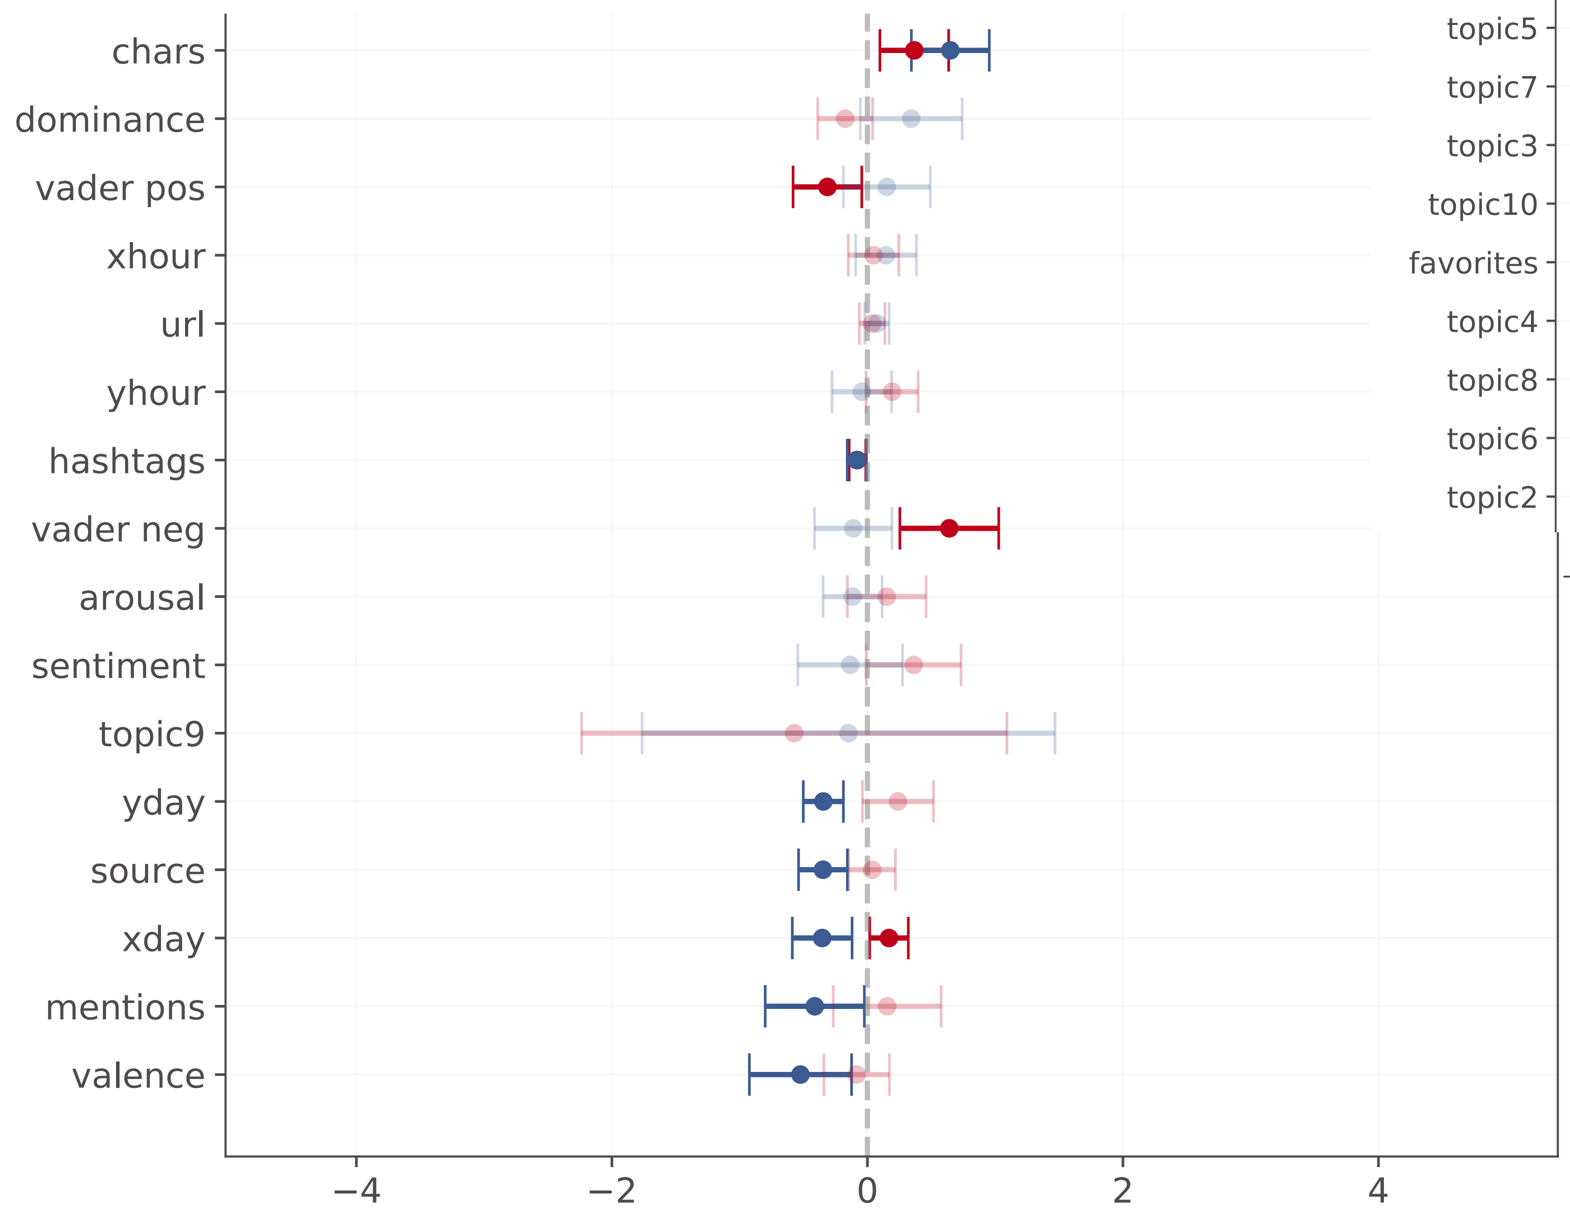

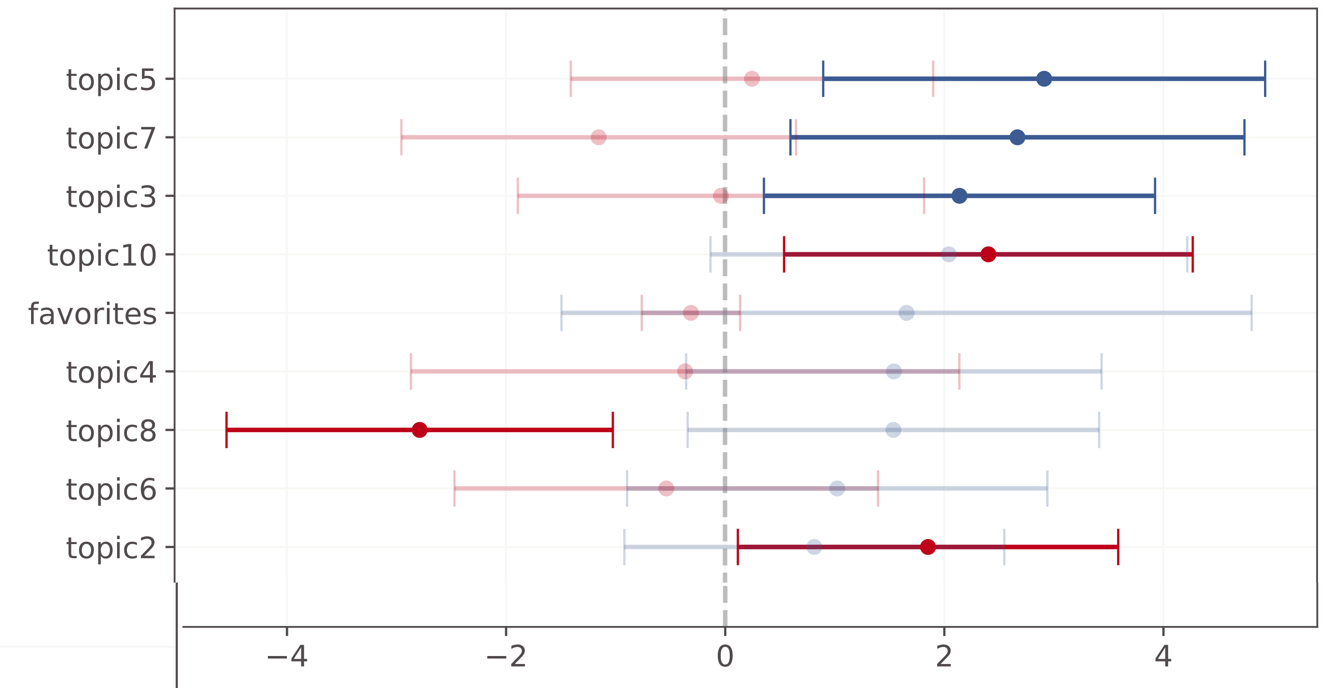

Candidate respondent
Current tweet author

10-topic LDA

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10
hope	love	people	good	people	thought	puerto	true	news	live
hillary	sad	pr	mayor	lol	evidence	rico	wrong	fake	usa
bot	big	money	day	black	funny	years	obama	real	war
agree	yeah	power	god	white	russian	lies	people	time	matter
cnn	people	dying	work	racist	means	people	president	flag	country
happen	dont	water	supplies	point	food	understand	vote	protest	marathi
states	person	tax	great	hate	read	party	thing	stand	tweeting
liar	blame	hurricane	job	guy	act	white	shit	talking	class
argument	wow	taking	san	bad	facts	rich	care	watch	leader
clinton	pr	days	juan	problem	helping	world	donald	anthem	place

Blei et al, 2003

10-topic LDA

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10
hope	love	people	good	people	thought	puerto	true	news	live
hillary	sad	pr	mayor	lol	evidence	rico	wrong	fake	usa
bot	big	money	day	black	funny	years	obama	real	war
agree	yeah	power	god	white	russian	lies	people	time	matter
cnn	people	dying	work	racist	means	people	president	flag	country
happen	dont	water	supplies	point	food	understand	vote	protest	marathi
states	person	tax	great	hate	read	party	thing	stand	tweeting
liar	blame	hurricane	job	guy	act	white	shit	talking	class
argument	wow	taking	san	bad	facts	rich	care	watch	leader
clinton	pr	days	juan	problem	helping	world	donald	anthem	place

Blei et al, 2003

10-topic LDA

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10
hope	love	people	good	people	thought	puerto	true	news	live
hillary	sad	pr	mayor	lol	evidence	rico	wrong	fake	usa
bot	big	money	day	black	funny	years	obama	real	war
agree	yeah	power	god	white	russian	lies	people	time	matter
cnn	people	dying	work	racist	means	people	president	flag	country
happen	dont	water	supplies	point	food	understand	vote	protest	marathi
states	person	tax	great	hate	read	party	thing	stand	tweeting
liar	blame	hurricane	job	guy	act	white	shit	talking	class
argument	wow	taking	san	bad	facts	rich	care	watch	leader
clinton	pr	days	juan	problem	helping	world	donald	anthem	place

Blei et al, 2003

10-topic LDA

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10
hope	love	people	good	people	thought	puerto	true	news	live
hillary	sad	pr	mayor	lol	evidence	rico	wrong	fake	usa
bot	big	money	day	black	funny	years	obama	real	war
agree	yeah	power	god	white	russian	lies	people	time	matter
cnn	people	dying	work	racist	means	people	president	flag	country
happen	dont	water	supplies	point	food	understand	vote	protest	marathi
states	person	tax	great	hate	read	party	thing	stand	tweeting
liar	blame	hurricane	job	guy	act	white	shit	talking	class
argument	wow	taking	san	bad	facts	rich	care	watch	leader
clinton	pr	days	juan	problem	helping	world	donald	anthem	place

Blei et al, 2003

Side note about LDA

Figure 1. The intuitions behind latent Dirichlet allocation. We assume that some number of "topics," which are distributions over words, exist for the whole collection (far left). Each document is assumed to be generated as follows. First choose a distribution over the topics (the histogram at right); then, for each word, choose a topic assignment (the colored coins) and choose the word from the corresponding topic. The topics and topic assignments in this figure are illustrative—they are not fit from real data. See Figure 2 for topics fit from data.

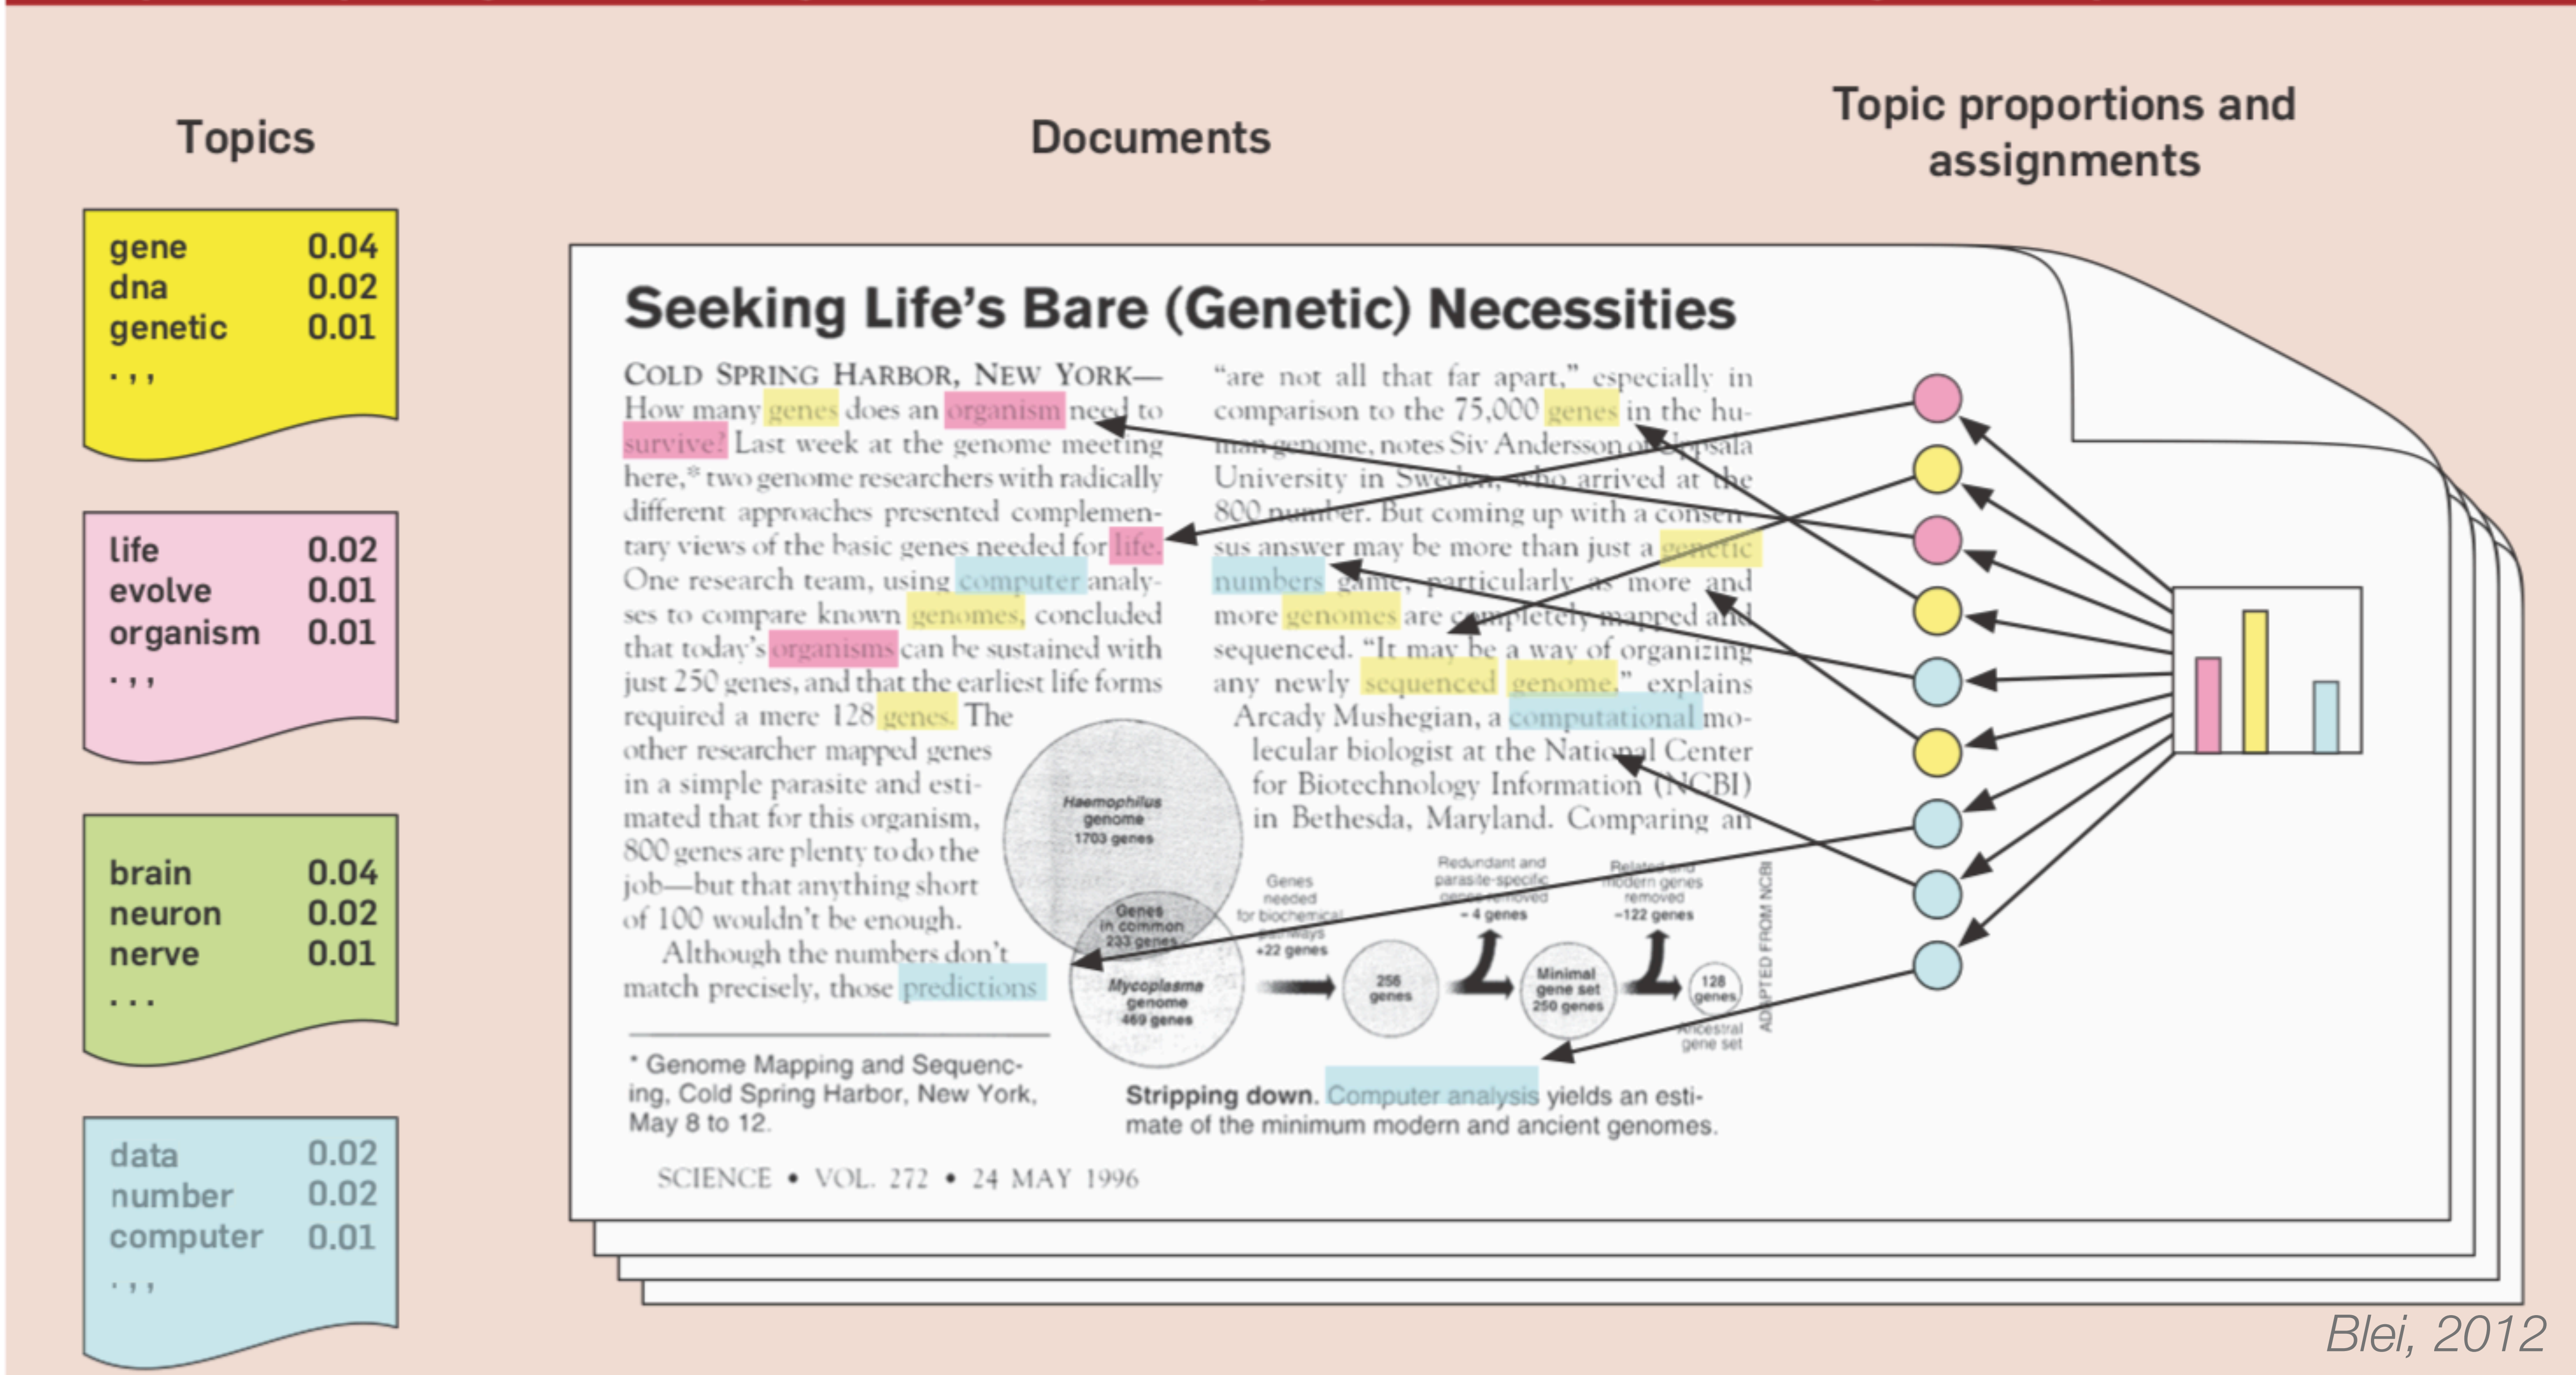

Blei, 2012

Side note about LDA

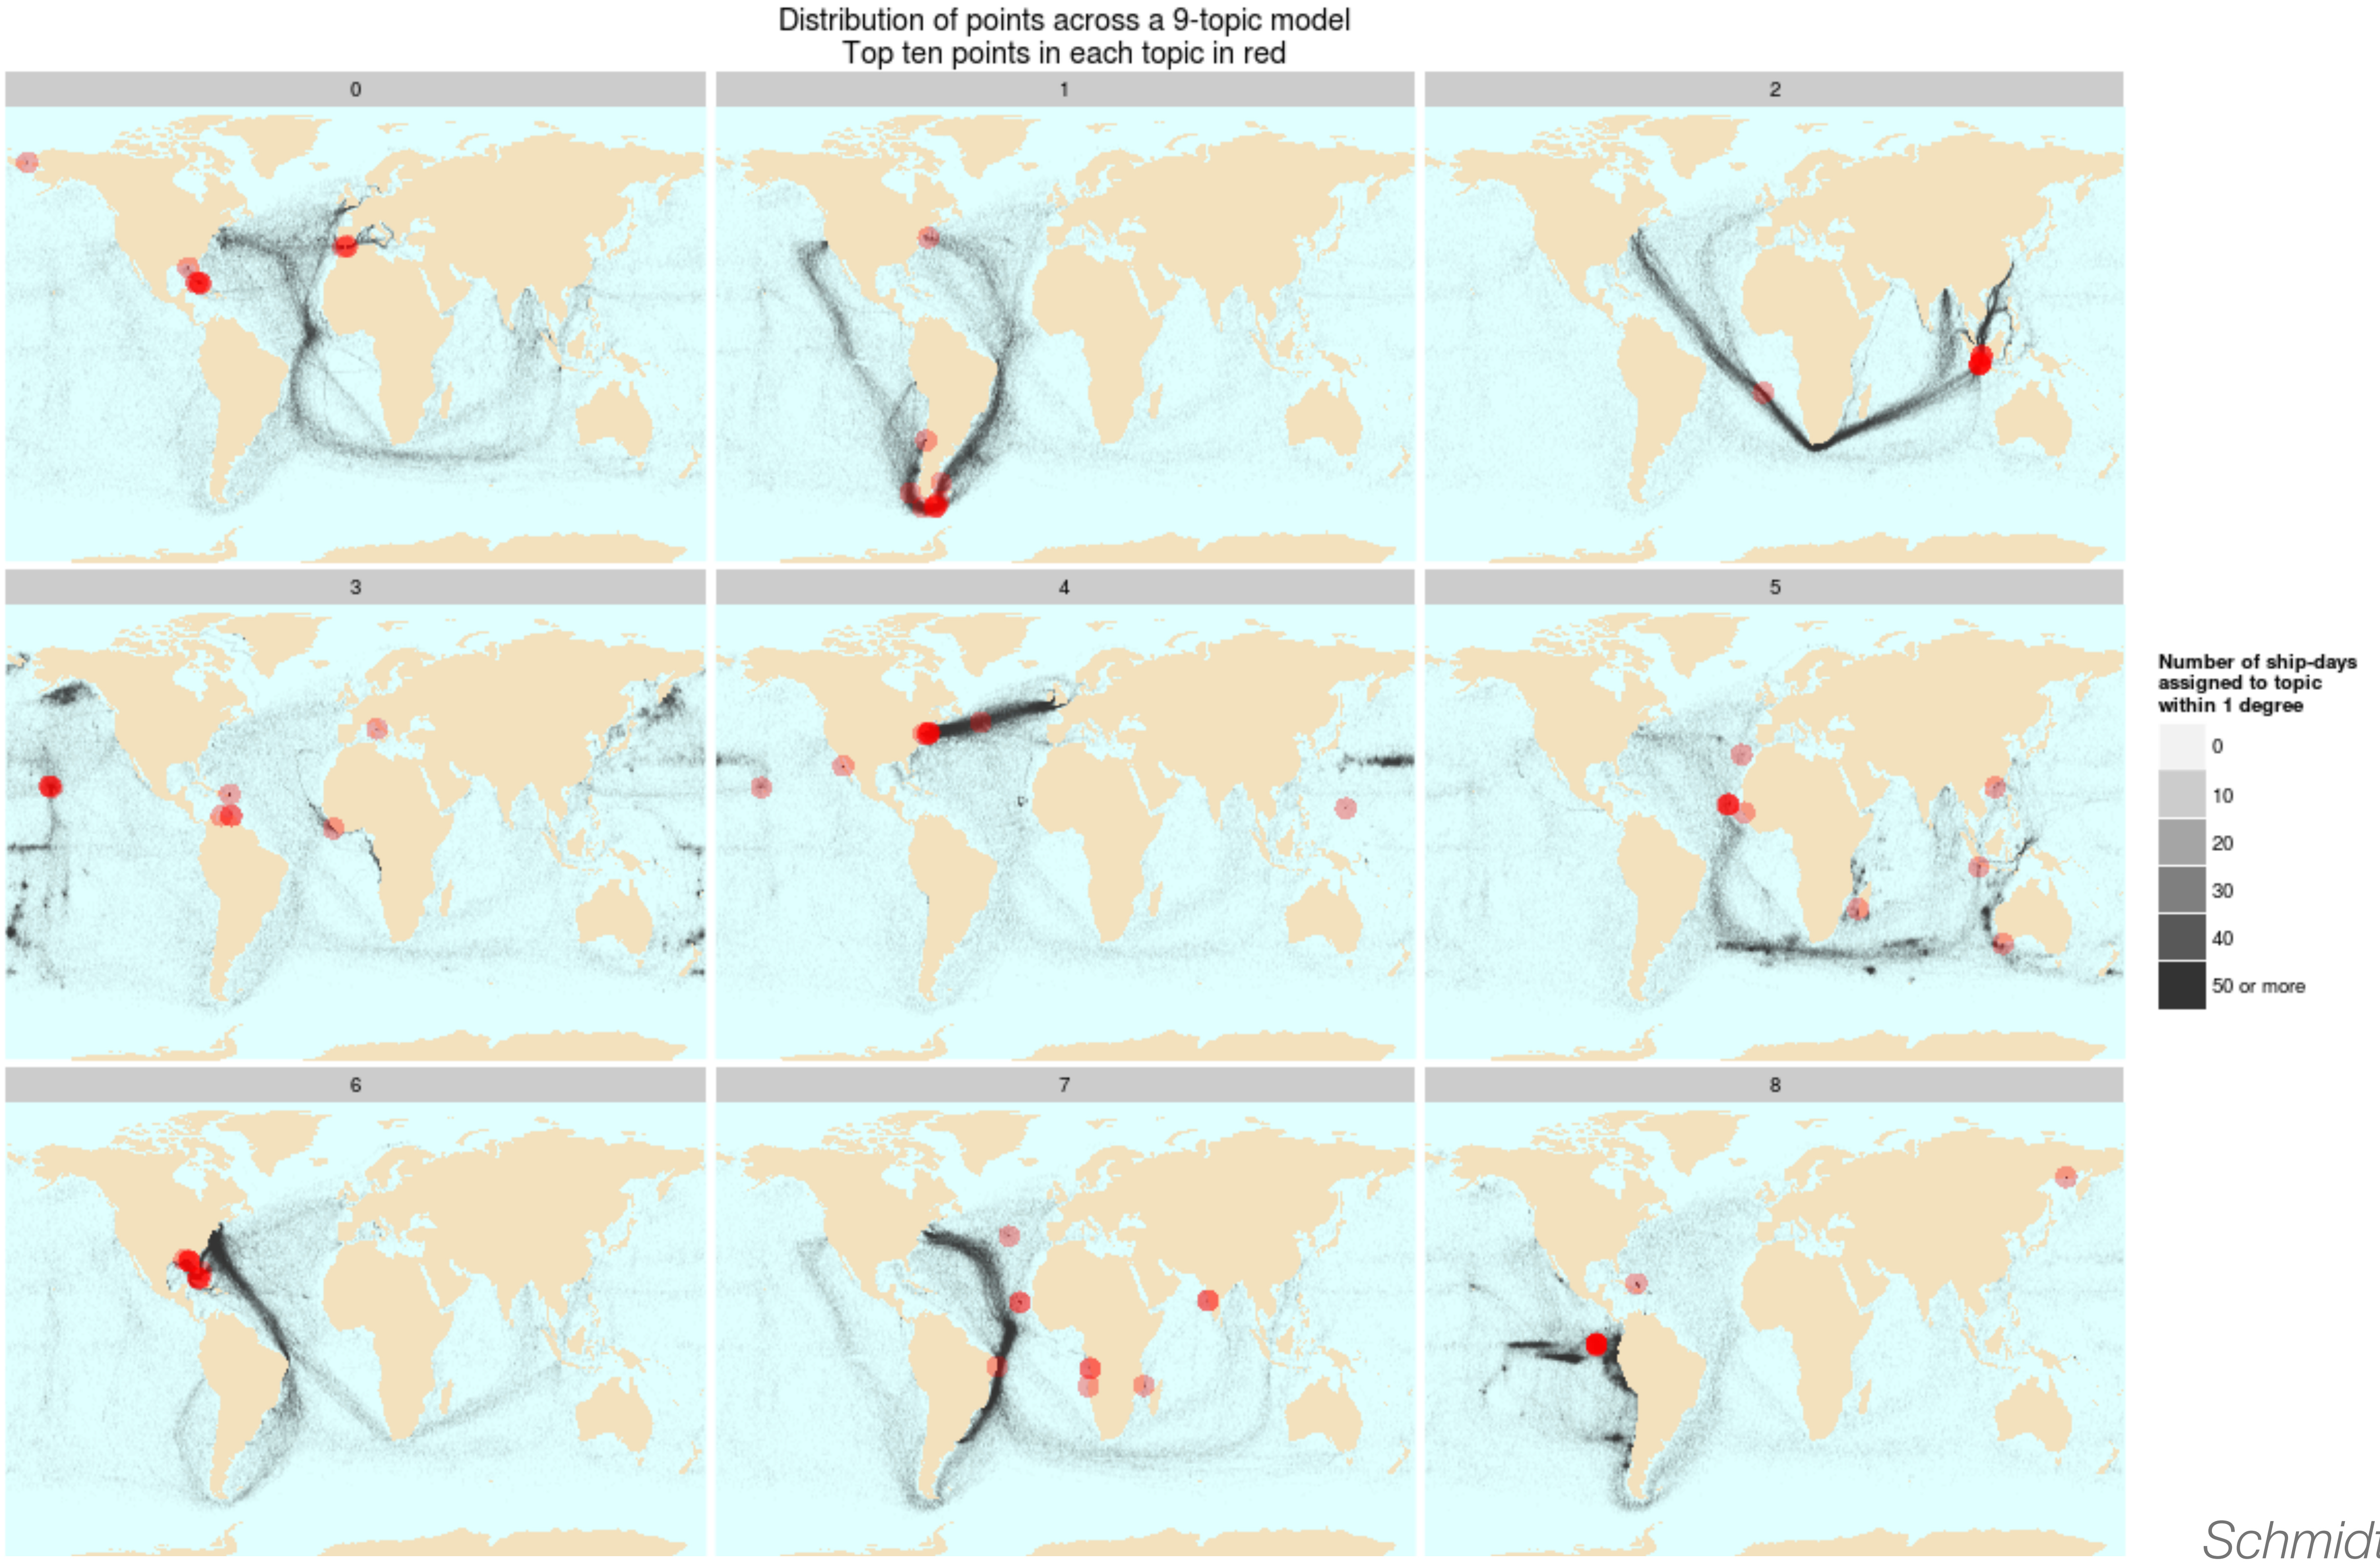

Schmidt, 2012

Sentiment and valence

- Positive tweets often received a positive response
- Tweets often crossed topics and ideological divides

Topic correlations

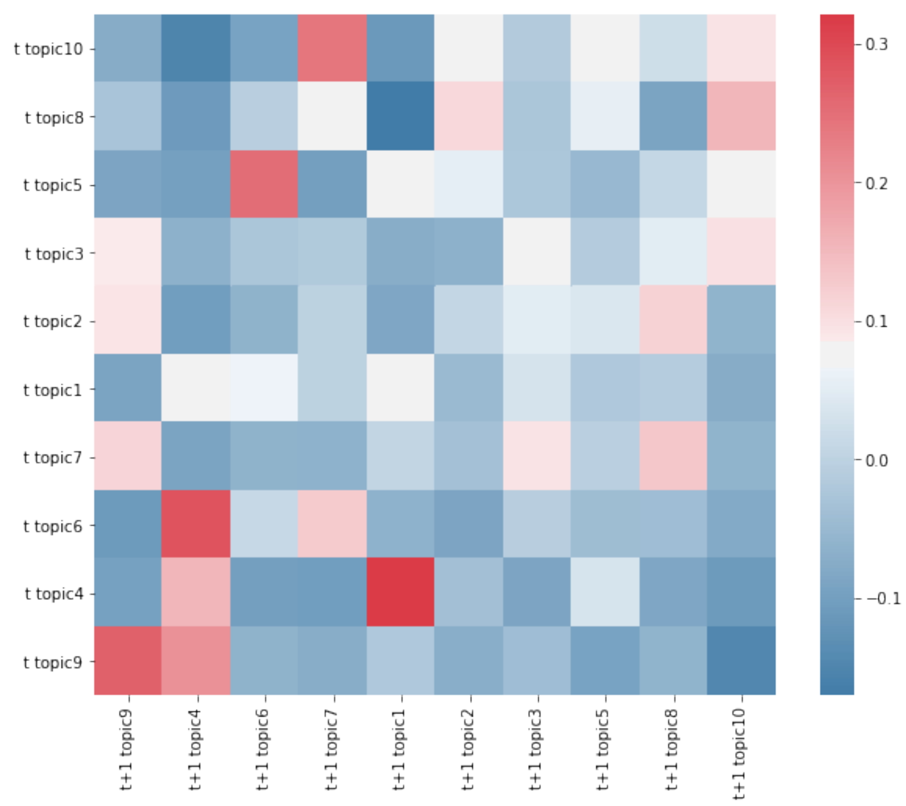

Emotional correlations

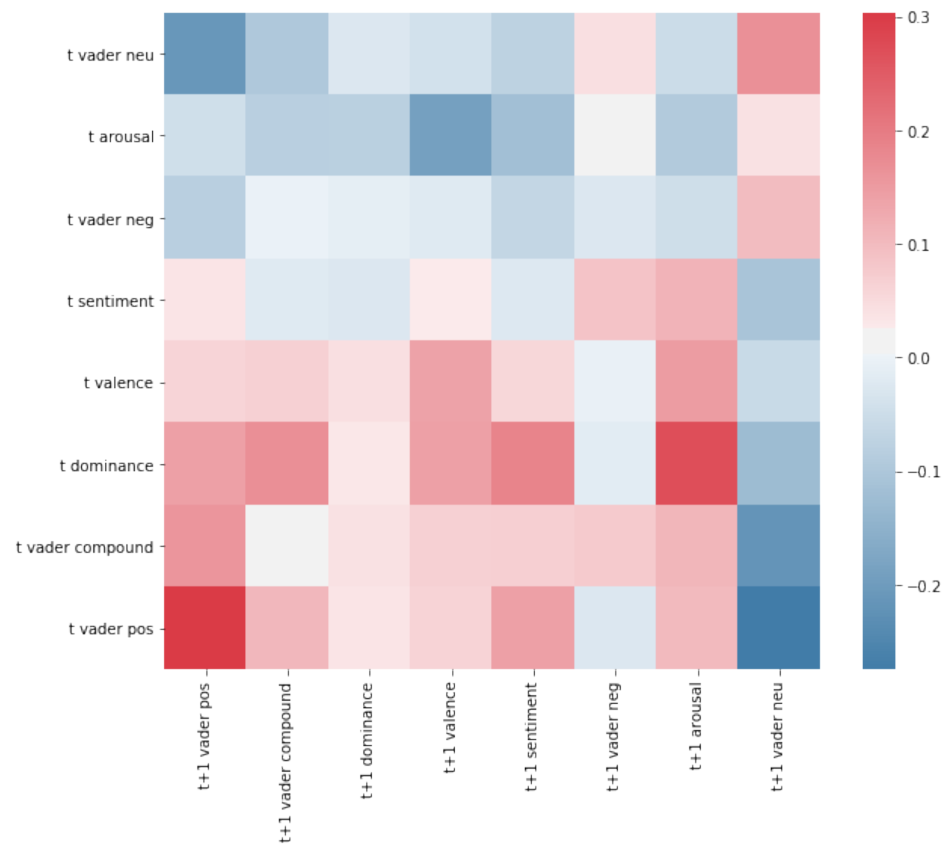

Implications

- This is promising because even single-comment interventions have shown to improve discourse quality

- Deliberative theory suggests that repeated interactions can have a greater positive impact on discourse quality

Friggeri et al., 2014; Munger, 2017; Bednar and Page, 2007; Habermas, 1984

Summary

- Negative tweets spark sustained conversation
- These conversations cross ideological divides
- Some twitter conversations remain civil
- There is hope for productive political conversations!

Case Study 2

The Structure of Reasoning: Measuring Expressions of Political Preference

Sarah Shugars, Network Science Institute

Case Study 2

Can we measure:

- Individual variation in how
- People structure their political expressions
- And do we really care anyway?

Case Study 2

Can we measure:

- Individual variation in how
- People structure their political expressions
- And do we really care anyway?

Spoiler Alert:

Case Study 2

Can we measure:

- Individual variation in how
- People structure their political expressions
- And do we really care anyway?

Spoiler Alert: Yes.

Case Study 2

Can we measure:

- Individual variation in how
- **People structure their political expressions**
- And do we really care anyway?

Spoiler Alert: Yes.

Case Study 2

Can we measure:

- **Individual variation** in how
- People structure their political expressions
- And do we really care anyway?

Spoiler Alert: Yes.

Case Study 2

Can we **measure**:

- Individual variation in how
- People structure their political expressions
- And do we really care anyway?

Spoiler Alert: Yes.

Case Study 2

Can we measure:

- ➡ Individual variation in how
- ➡ People structure their political expressions
- ➡ And do we really **care** anyway?

Spoiler Alert: Yes.

Case Study 2

Can we measure:

- ➡ Individual variation in how
- ➡ People structure their political expressions
- ➡ **Potential for behavioral insights**

Spoiler Alert: Yes.

New Tools for a Classic Problem

This idea is not new

- A classic element of public opinion scholarship
- Efforts used interviews or hand-coding of text
- Largely abandoned as too difficult / time consuming

Lane, 1962; Axelrod, 1976; Campbell, 1960

New Tools for a Classic Problem

This idea is not new

- A classic element of public opinion scholarship
- Efforts used interviews or hand-coding of text
- Largely abandoned as too difficult / time consuming

Lane, 1962; Axelrod, 1976; Campbell, 1960

➡ Modern computational tools make this task tractable

Roadmap

1. Elaborate on “structure” of political reasoning
2. Define approach for inferring and measuring structure
3. Demonstrate potential for behavioral insights
   – using two distinct datasets

1. Political Reasoning is Structured

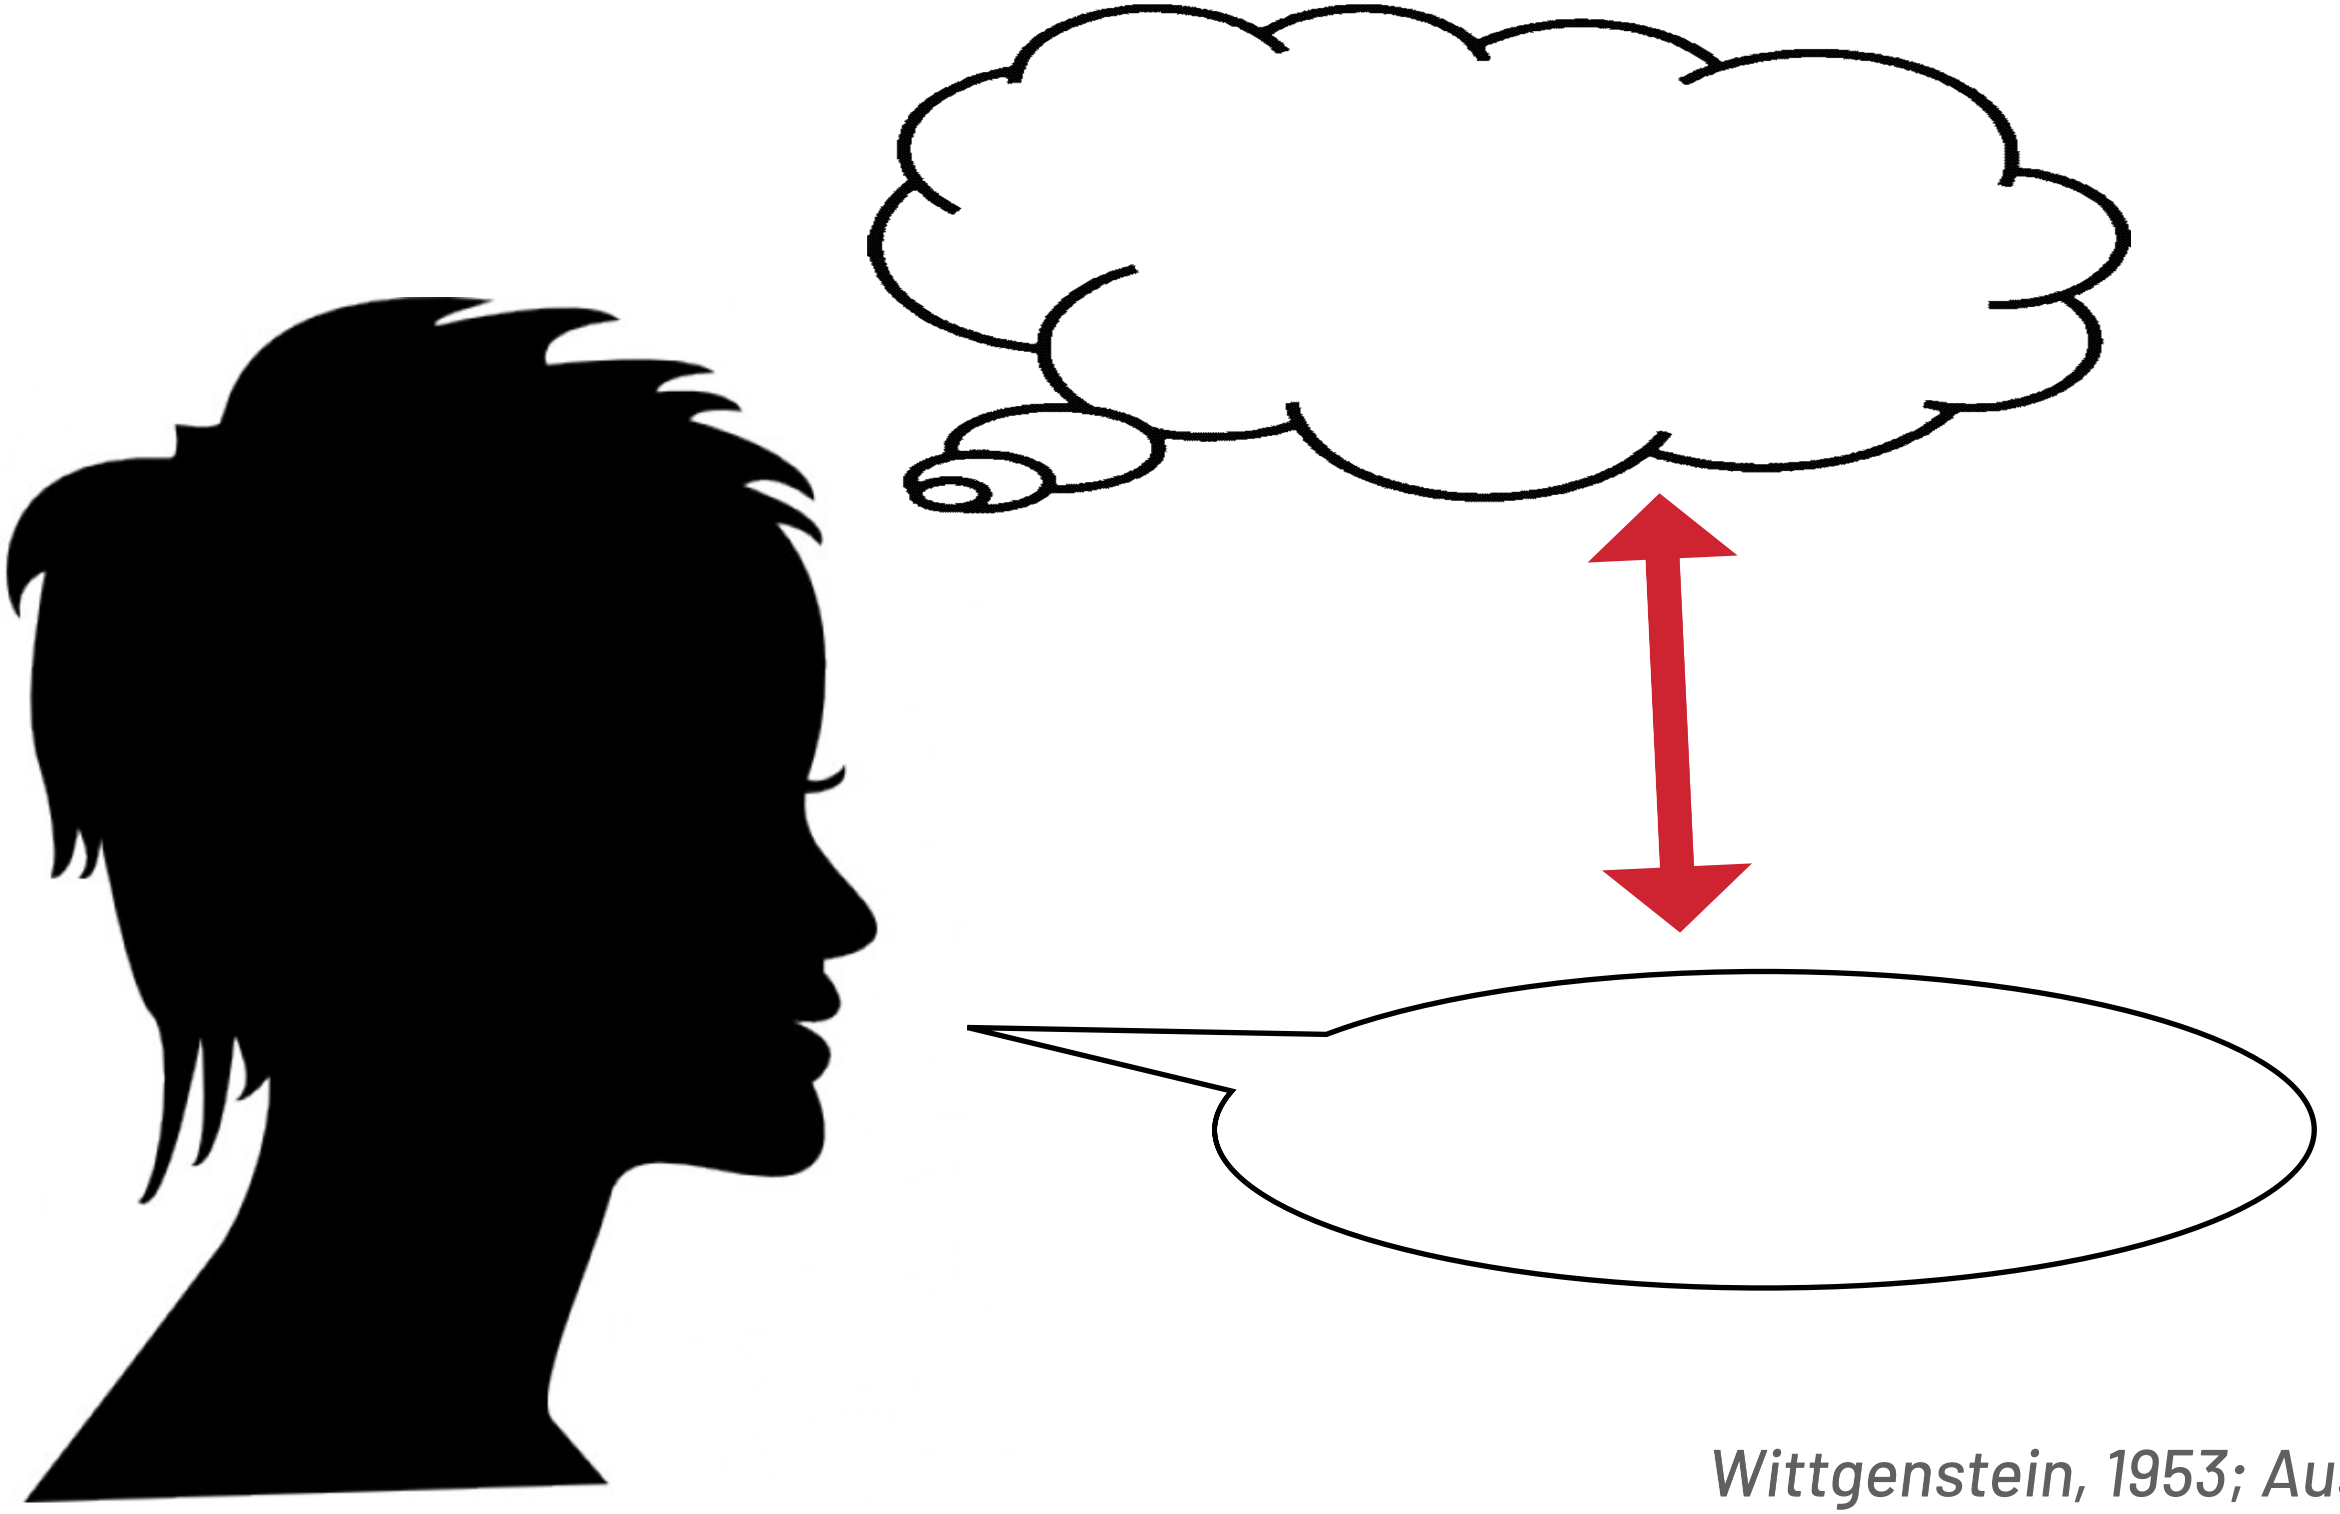

Wittgenstein, 1953; Austin, 1962

1. Political Reasoning is Structured

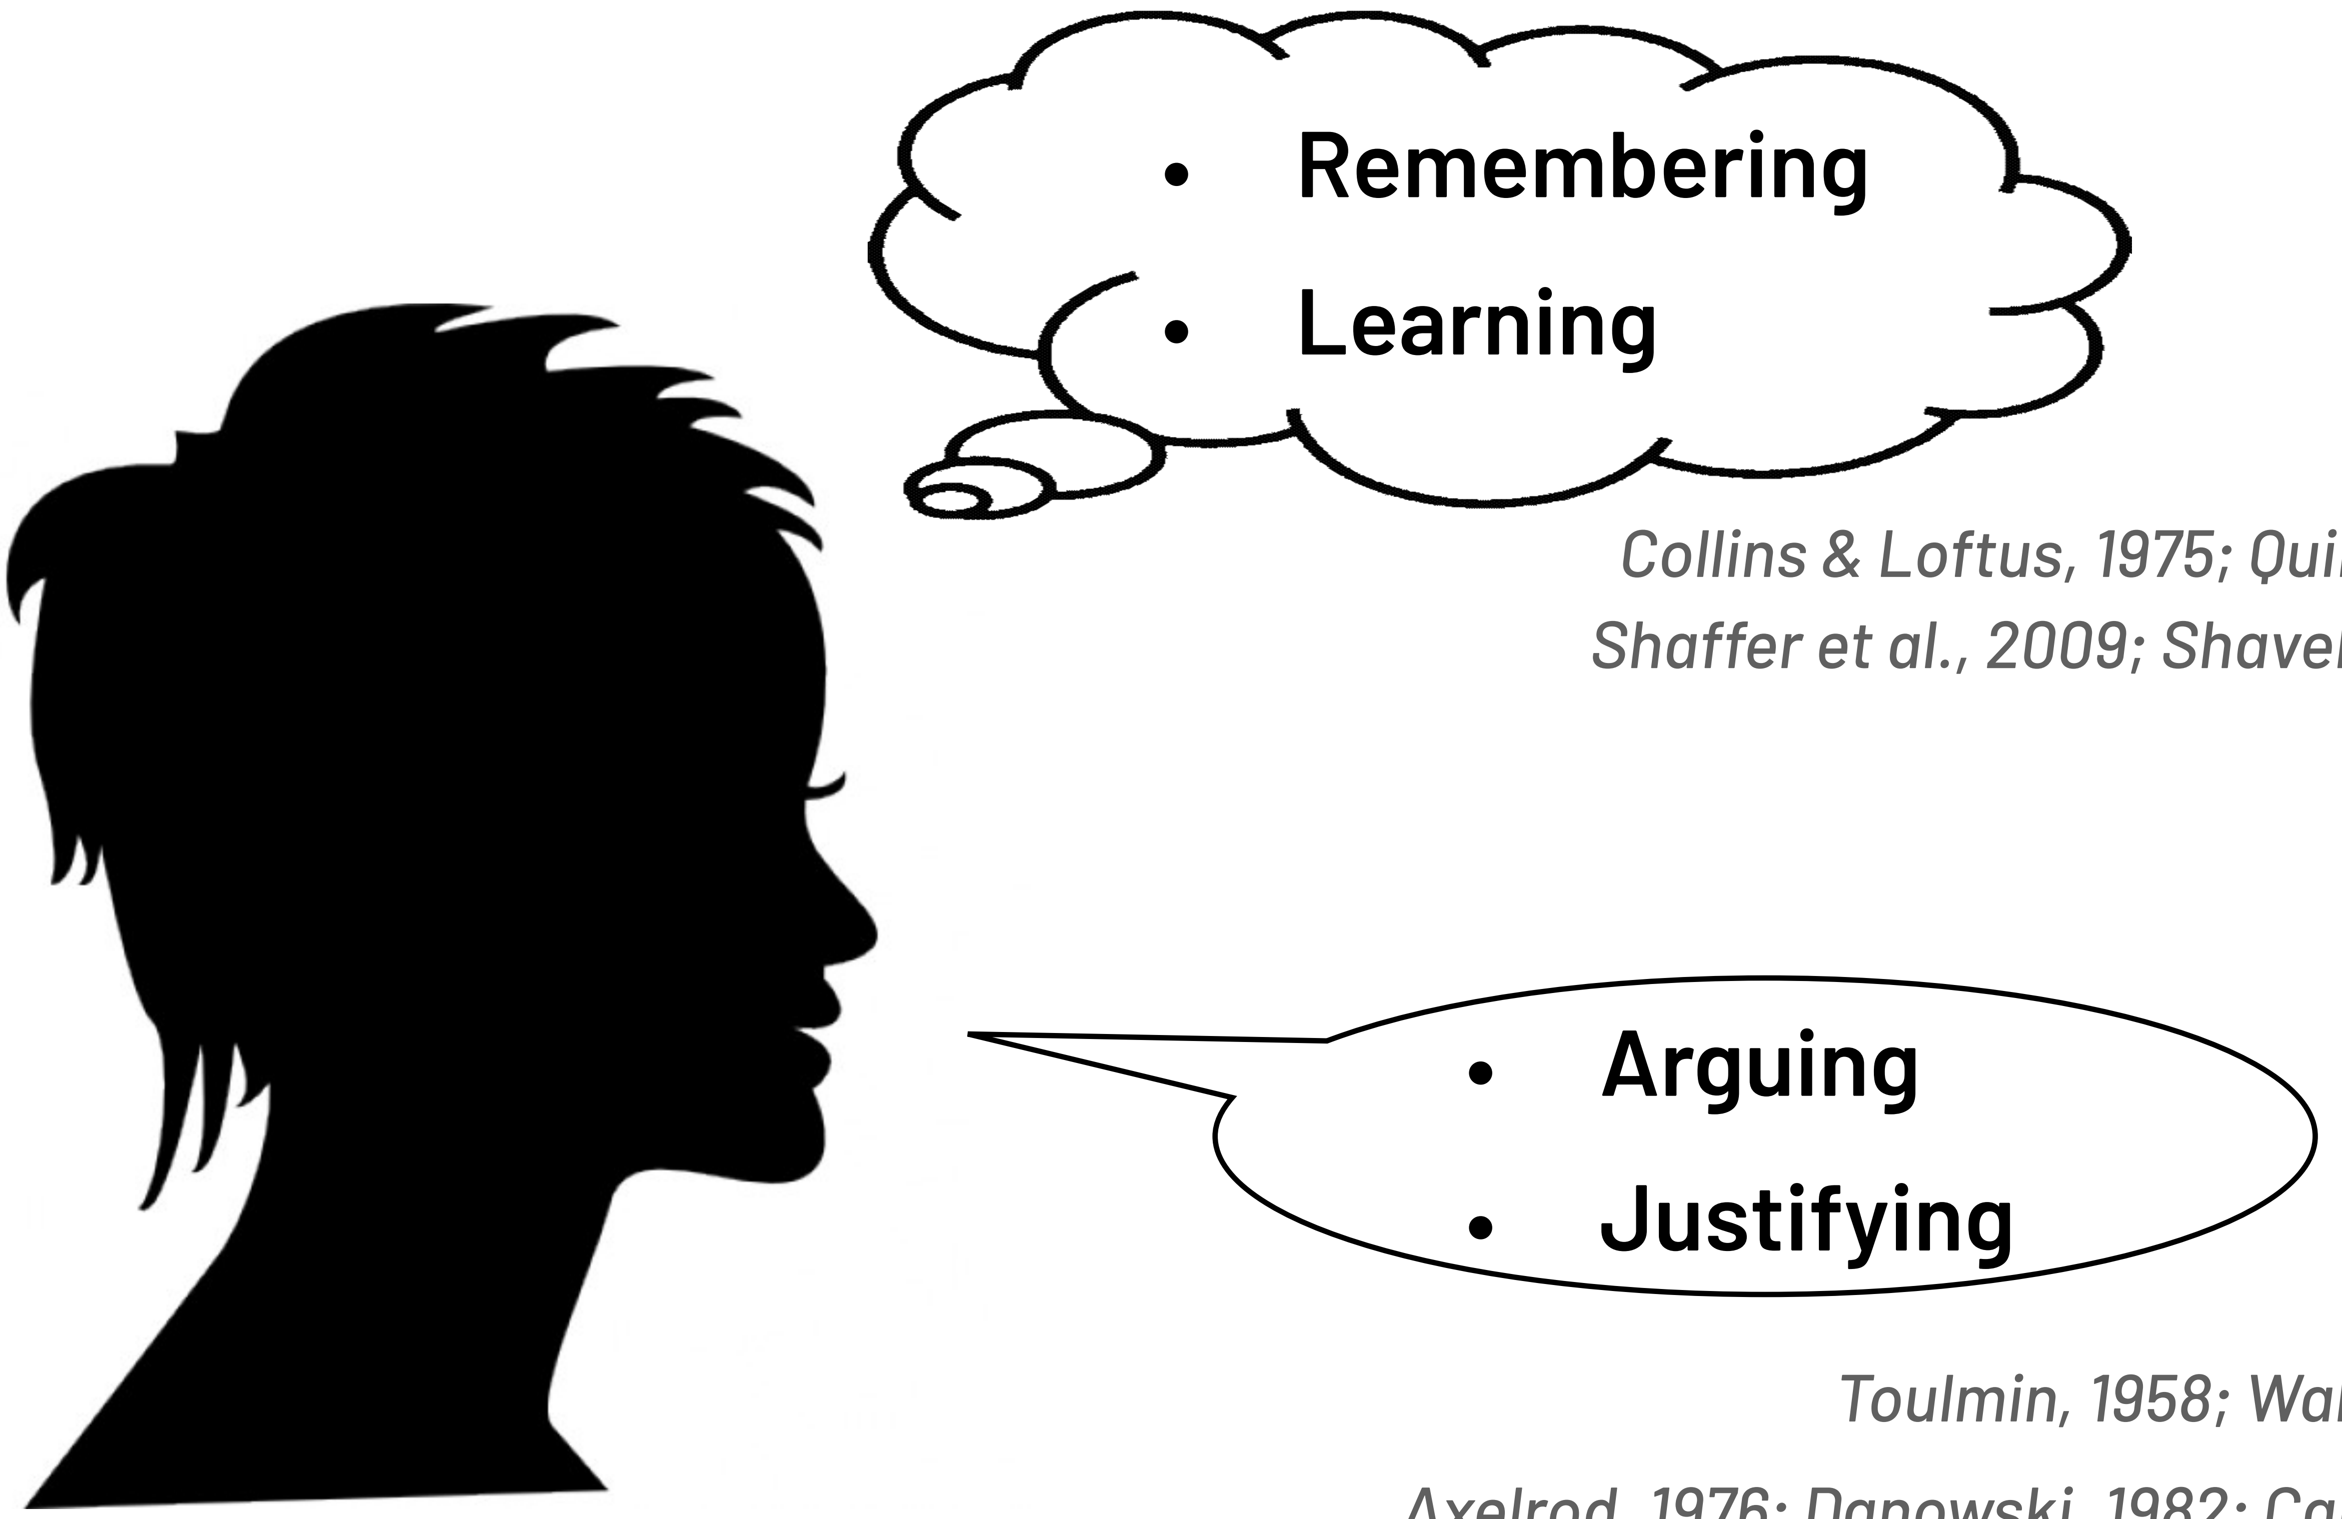

Collins & Loftus, 1975; Quillian, 1967

Shaffer et al., 2009; Shavelson, 1974

Toulmin, 1958; Walton, 1996

Axelrod, 1976; Danowski, 1982; Carley, 1993

1. Political Reasoning is Structured

- Remembering
- Learning

Both have

network structure

- Arguing
- Justifying

1. Political Reasoning is Structured

1. Political Reasoning is Structured

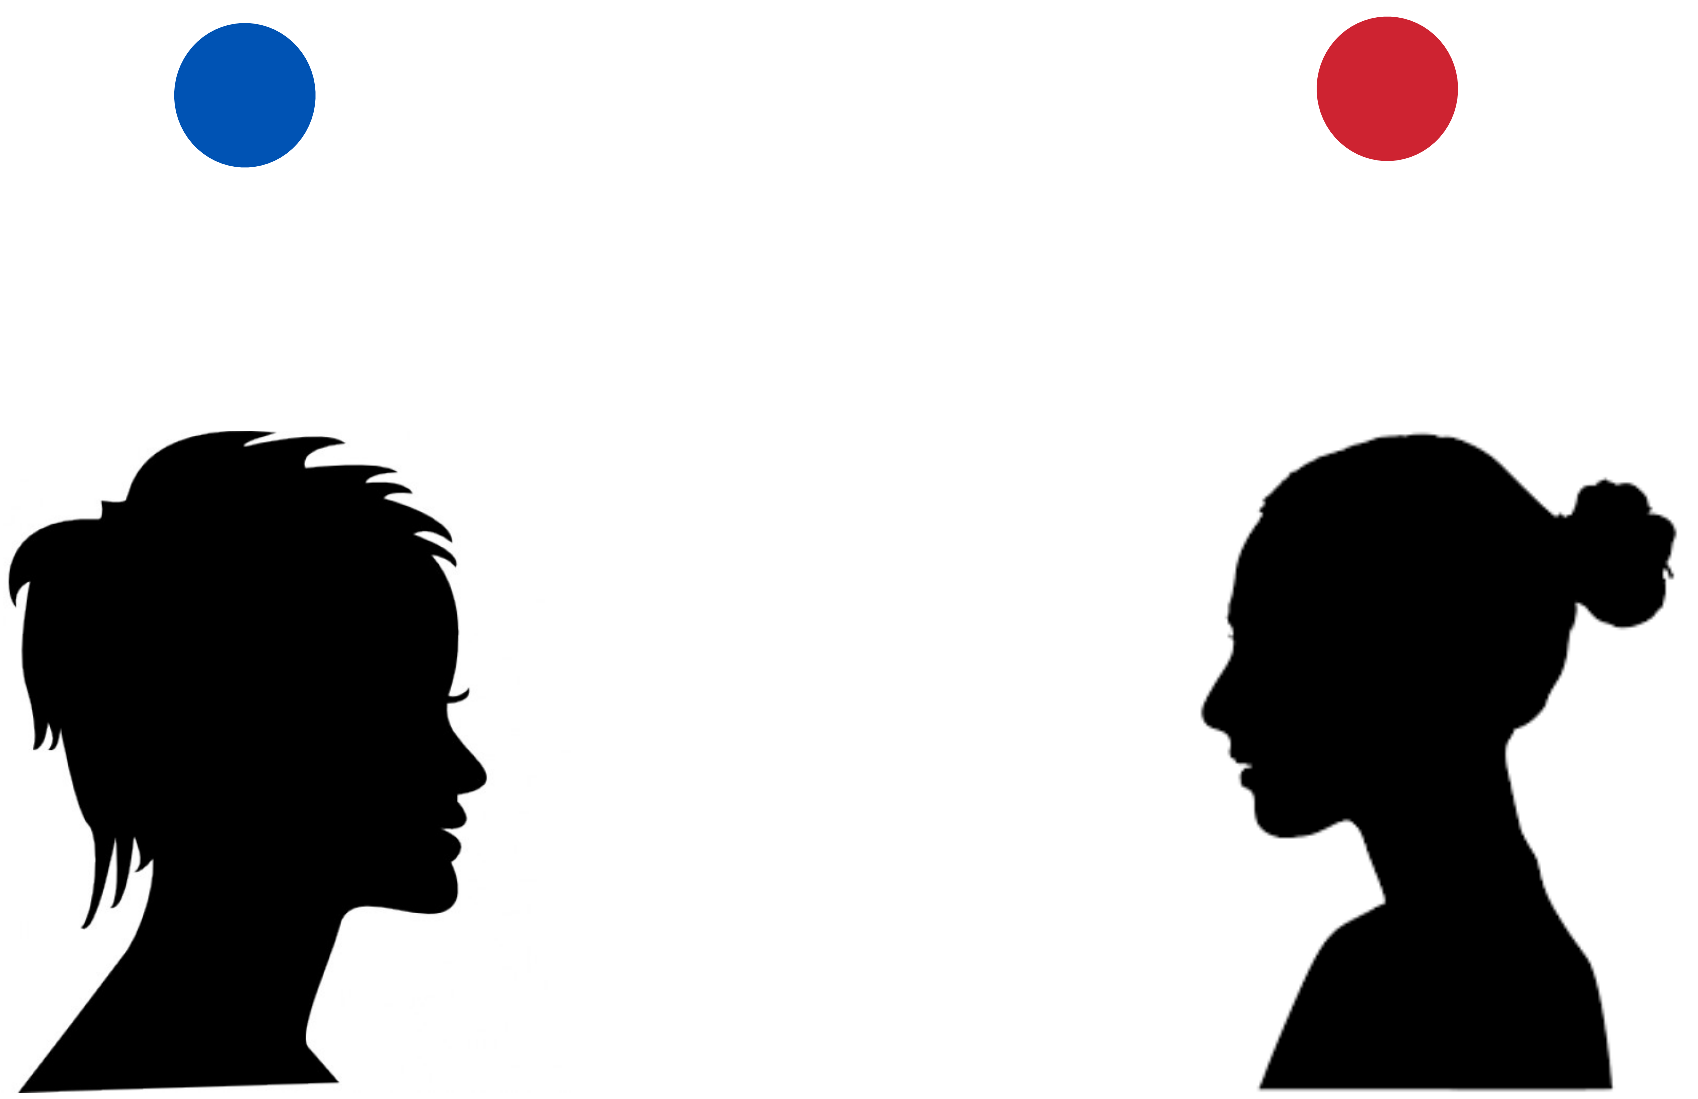

1. Political Reasoning is Structured

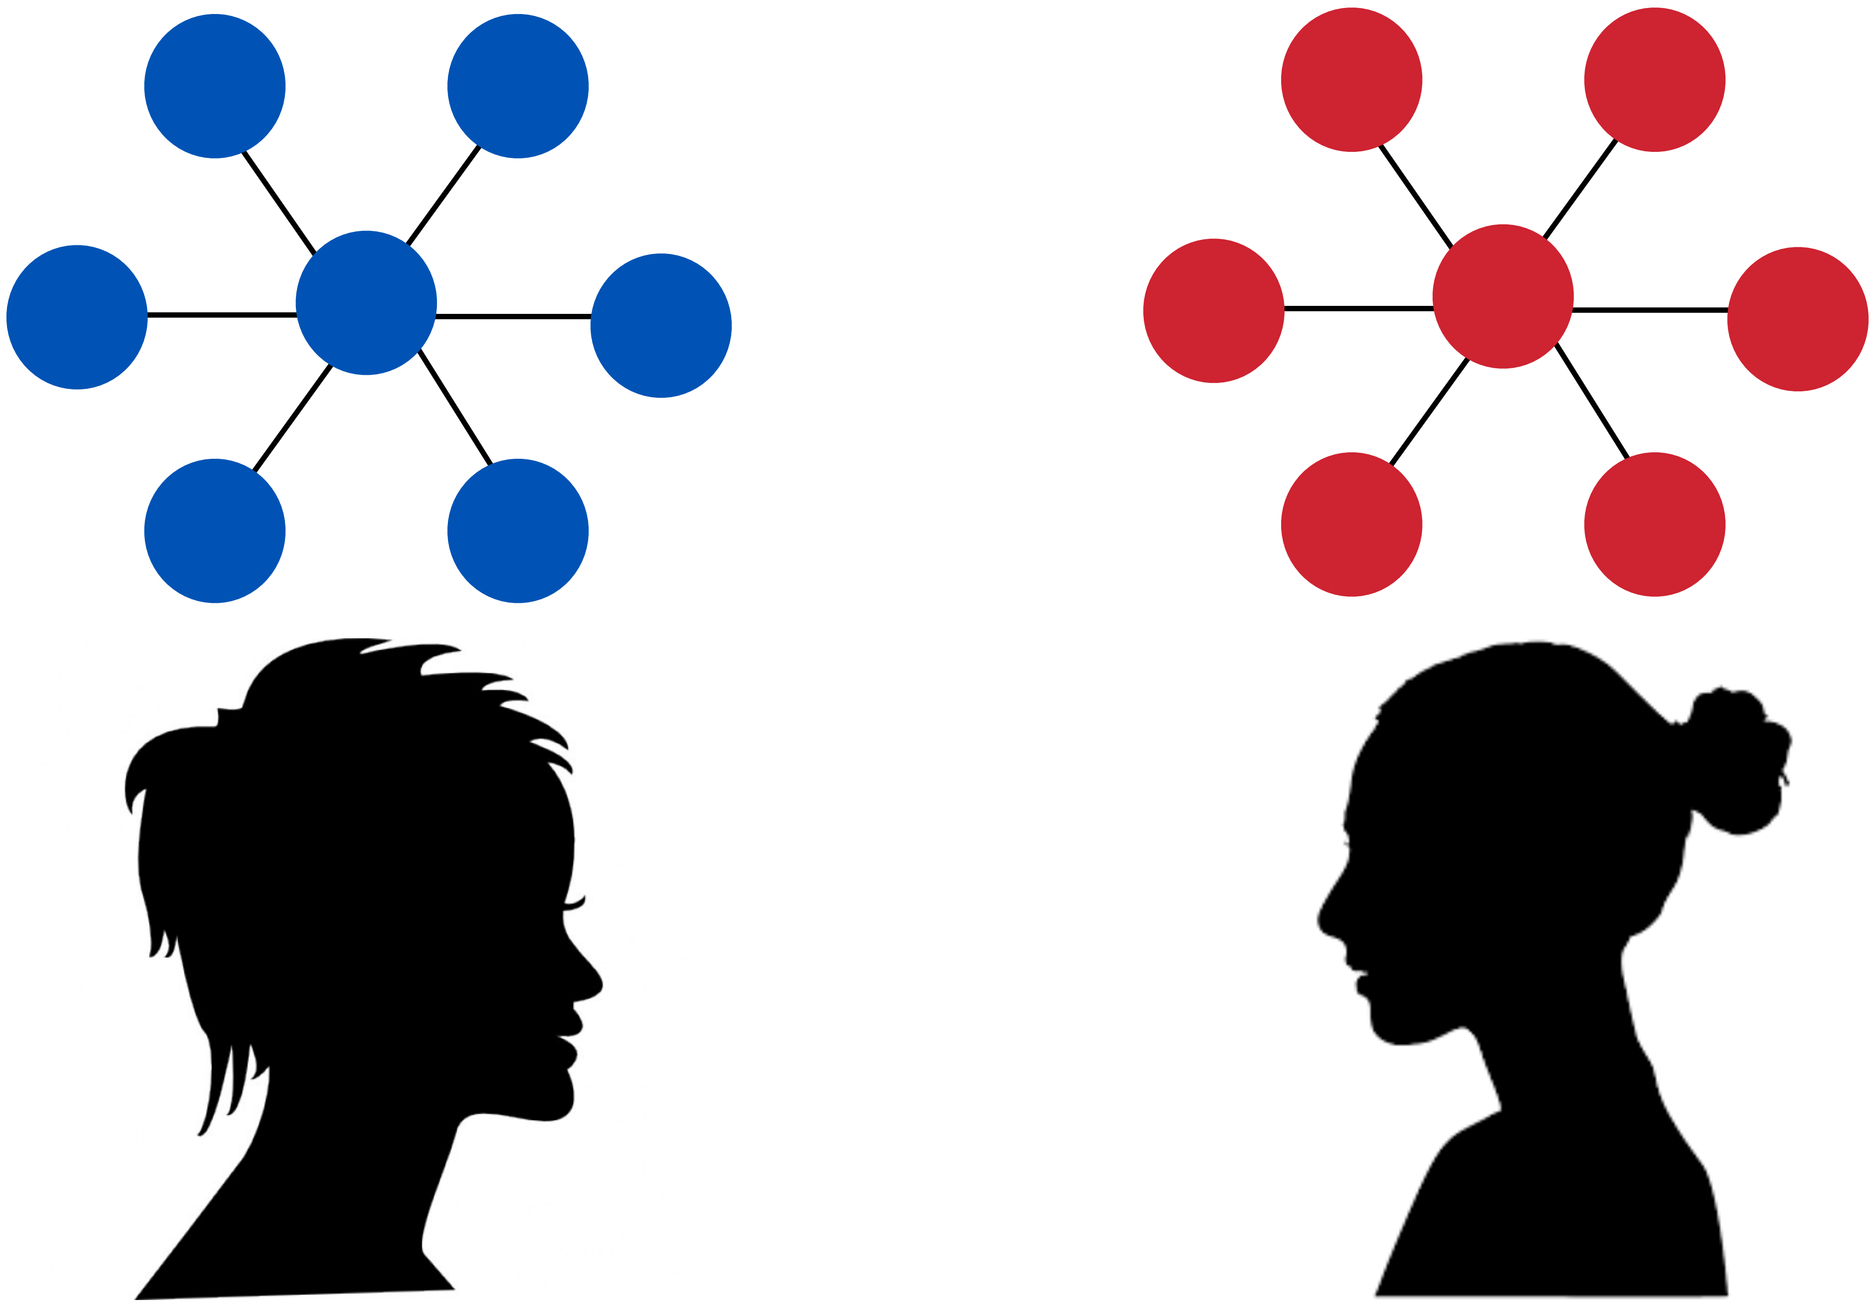

1. Political Reasoning is Structured

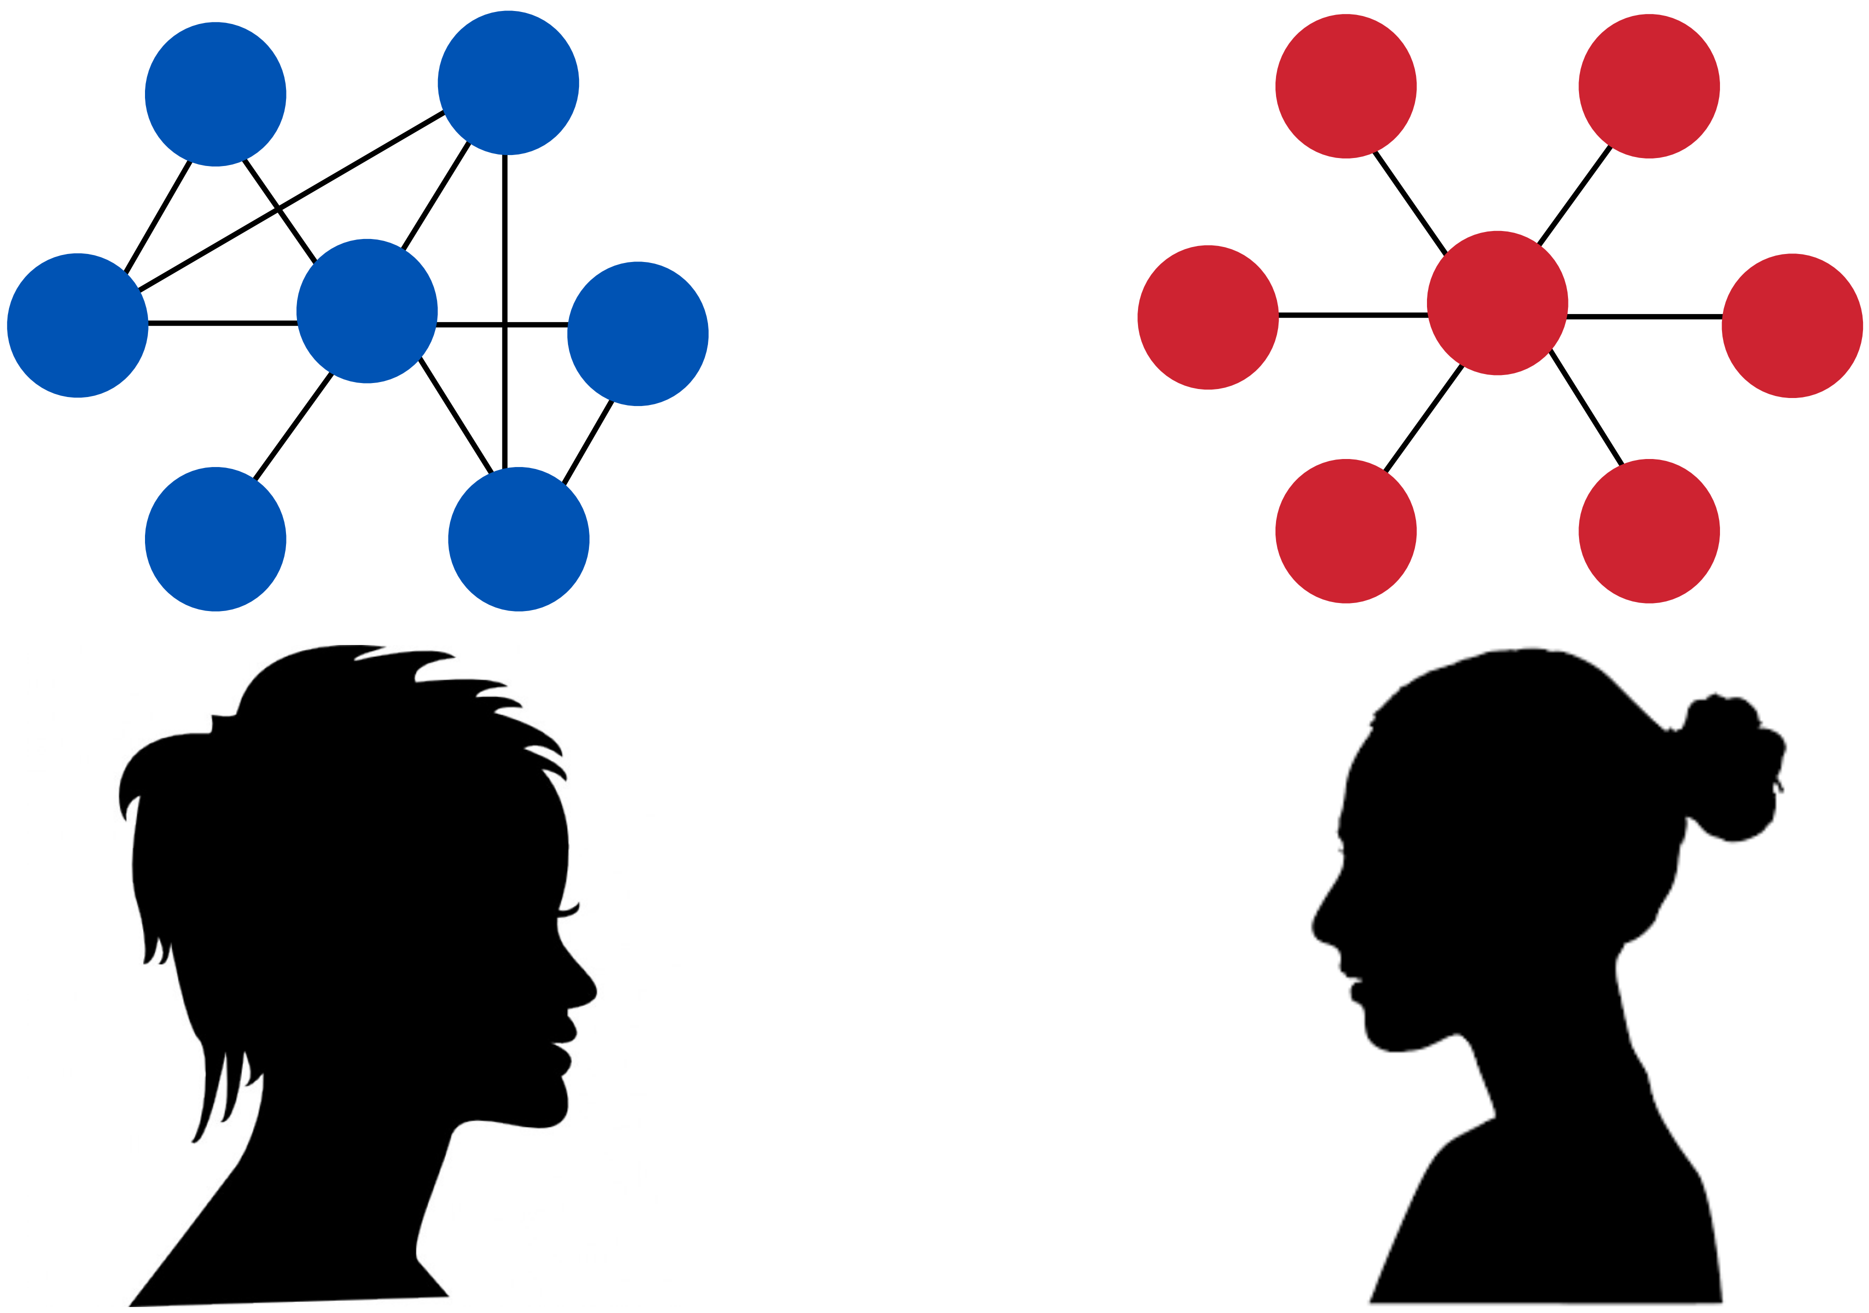

1. Political Reasoning is Structured

Structure and *content* both influence the quality of political talk

1. Political Reasoning is Structured

Structure and *content* both influence the quality of political talk

Structure:

- Sends a signal to interlocutor
- Influences receptivity to new messages
- Represents different philosophical approaches

1. Political Reasoning is Structured

Multiple moral philosophies claim:

Good* reasoning must be coherent*

Sidgwick, 1907; Dancy, 1993

McNaughton & Rawling, 2000; Rawls, 1993

Thagard, 1998; Dorsey, 2006; Berker, 2015

1. Political Reasoning is Structured

Multiple moral philosophies claim:

Good* reasoning
must be coherent*

** For some definitions of*
“good” and “coherent”

1. Political Reasoning is Structured

Multiple moral philosophies claim:

Good* reasoning must be coherent*

** For some definitions of "good" and "coherent"*

** Sidgwick, 1907; Dancy, 1993*

1. Political Reasoning is Structured

Connectivity
Baseline

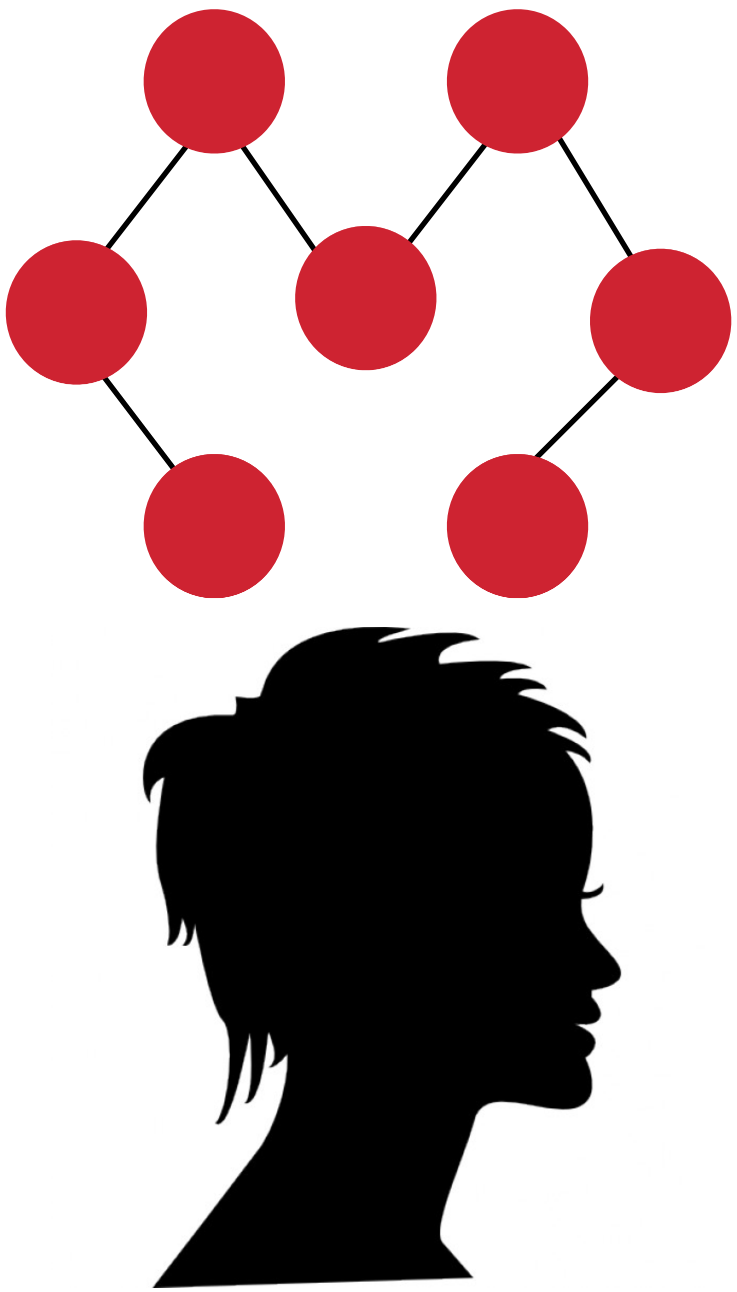

Complexity
Dancy, 1993

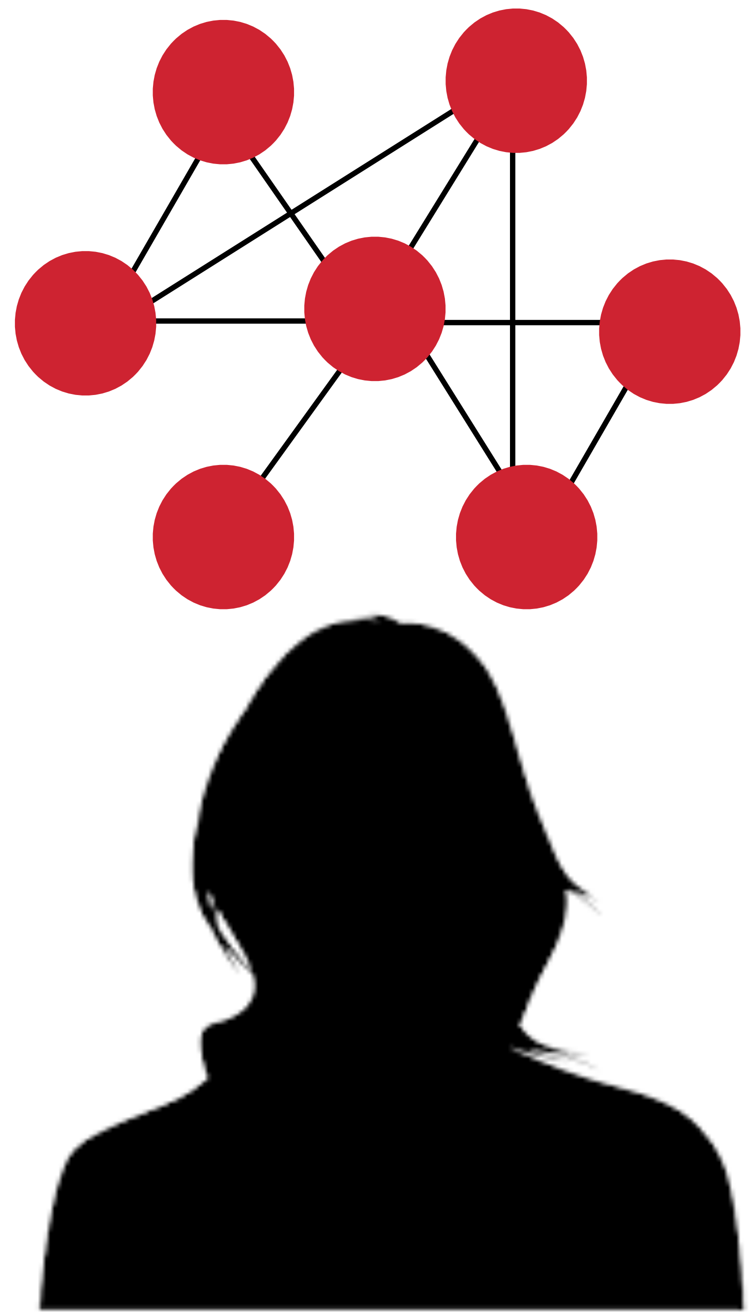

Hierarchy
Sidgwick, 1907

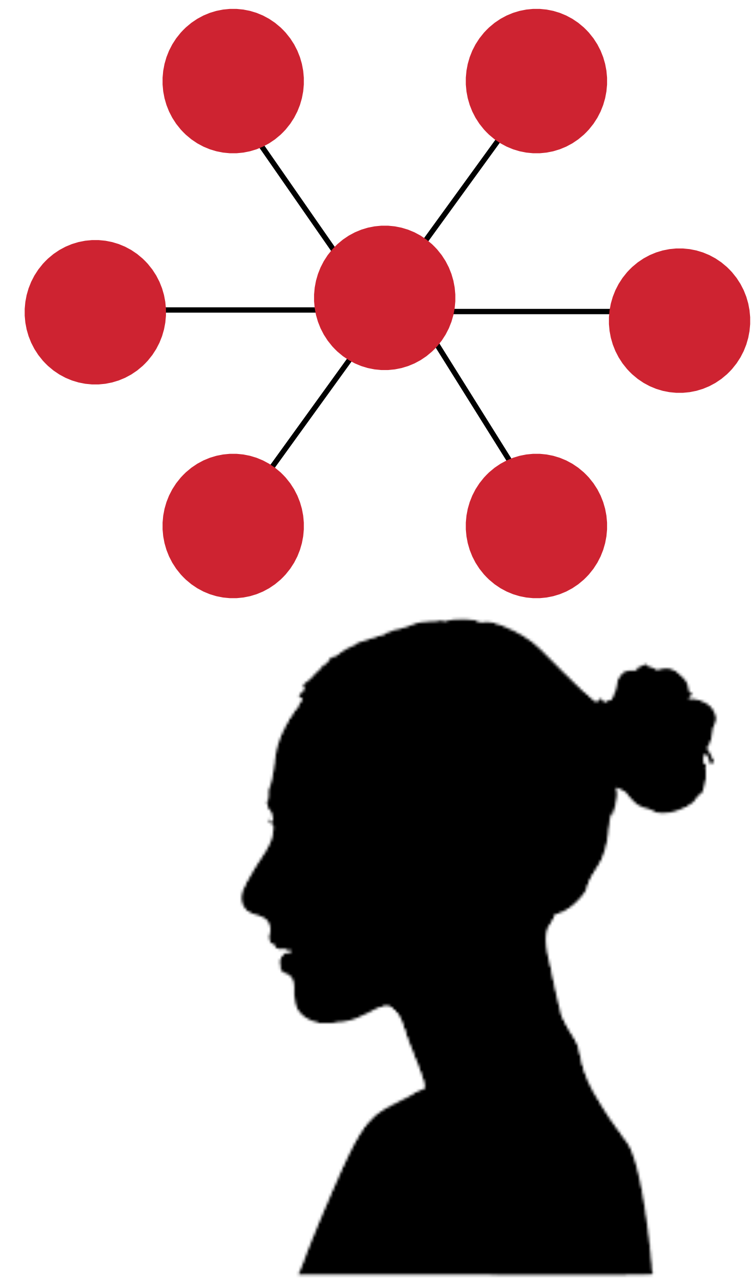

The Structure of Reasoning

Roadmap:

1. Elaborate on “structure”
2. **Define approach for inferring and measuring structure**
3. Demonstrate potential for behavioral insights
   – using two distinct datasets

2. Inferring Network Structure

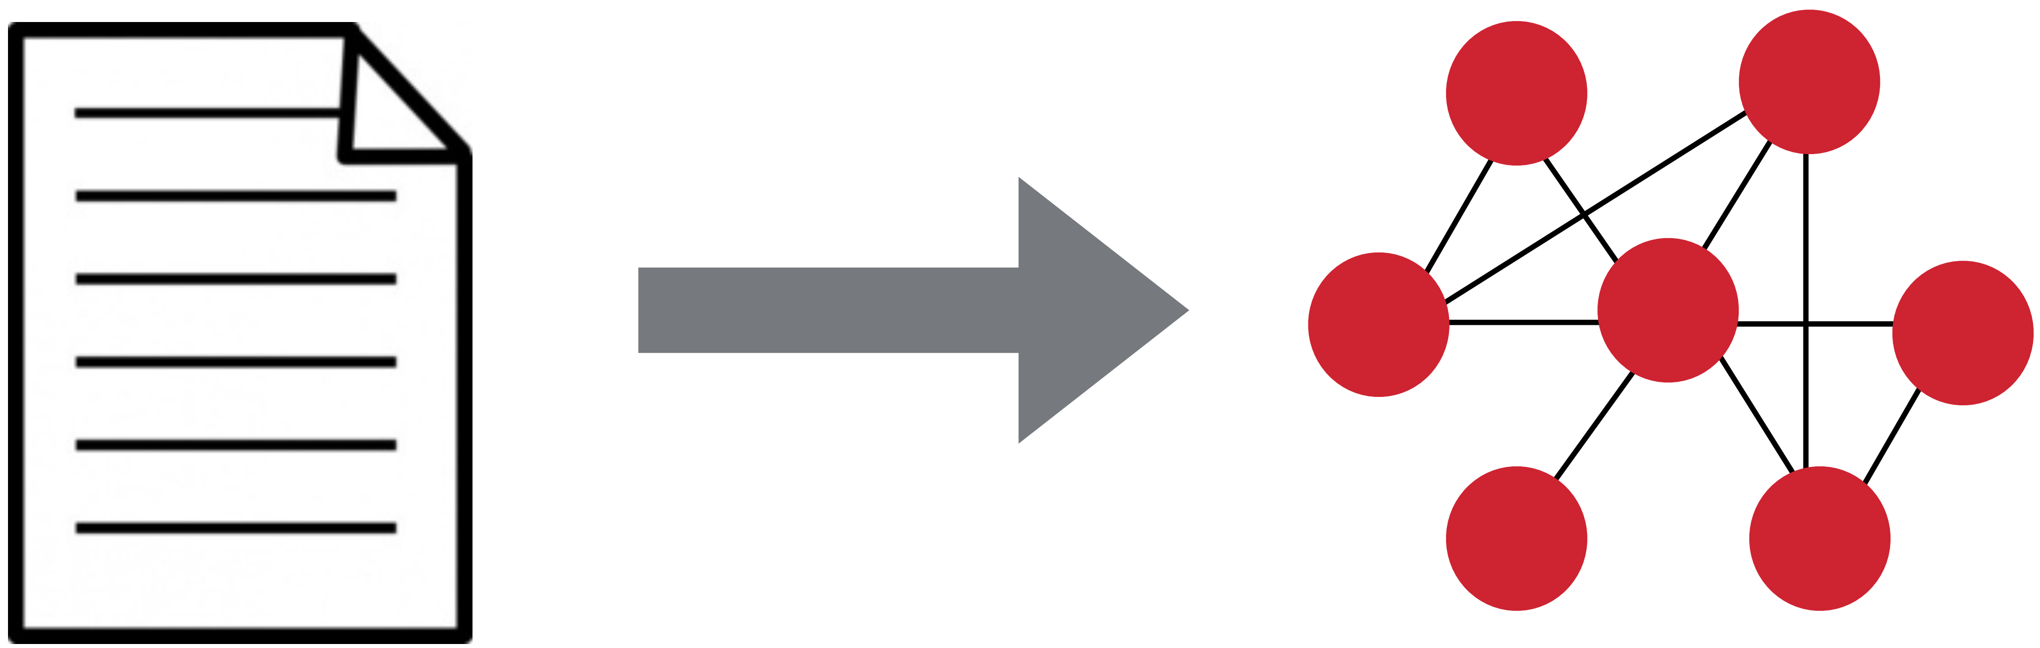

2. Inferring Network Structure

What are the nodes?

What are the edges?

2. Inferring Network Structure

What are the nodes?

- ➡ Concepts

What are the edges?

- ➡ Connections between concepts (??)

2. Inferring Network Structure

What is a "concept" ?

- Compressed representation of information
- Collection of related "things"
- Represented by words

➡ Operationally, a collection of similar words

Collins & Loftus, 1975; Quillian, 1967

2. Inferring Network Structure

Identifying similar words through embeddings:

2. Inferring Network Structure

Identifying similar words through embeddings:

- Words are high dimensional objects and can be embedded in high dimensional space

- Do this in such a way that words which appear in similar contexts are geometrically close

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{-c\leq j\leq c, j\neq 0}\log p\left(w_{t+j|w_t}\right)$$

Mikolov et al, 2013
Spirling and Rodriguez, 2019

2. Inferring Network Structure

Identifying similar words through embeddings:

I took my **dog** to the vet.

I took my **cat** to the vet.

2. Inferring Network Structure

Identifying similar words through embeddings:

I took my **dog** to the vet.

I took my **cat** to the vet.

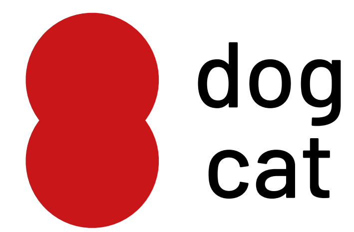

2. Inferring Network Structure

Identifying similar words through embeddings:

I took my dog to the vet.

I took my cat to the vet.

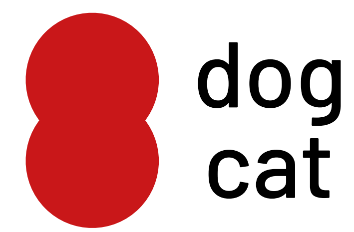

My dog plays fetch.

My cat likes to sleep.

2. Inferring Network Structure

Identifying similar words through embeddings:

I took my dog to the vet.
I took my cat to the vet.

My dog plays fetch.
My cat likes to sleep.

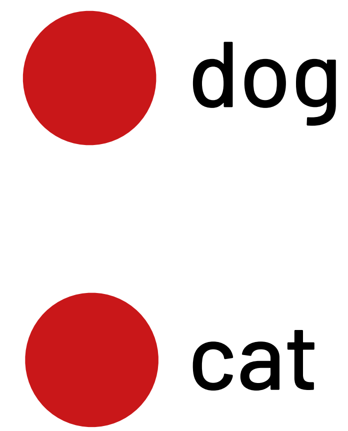

2. Inferring Network Structure

Identifying similar words through embeddings:

I took my dog to the vet.
I took my cat to the vet.

My dog plays fetch.
My cat likes to sleep.

I caught a shuttle from the airport.

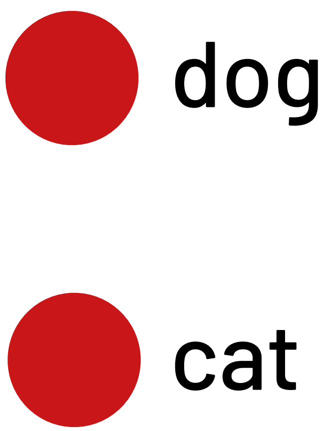

2. Inferring Network Structure

Identifying similar words through embeddings:

I took my **dog** to the vet.

I took my **cat** to the vet.

My **dog** plays fetch.

My **cat** likes to sleep.

I caught a **shuttle** from the airport.

Sidenote: Continuous Bag of Words (CBOW)

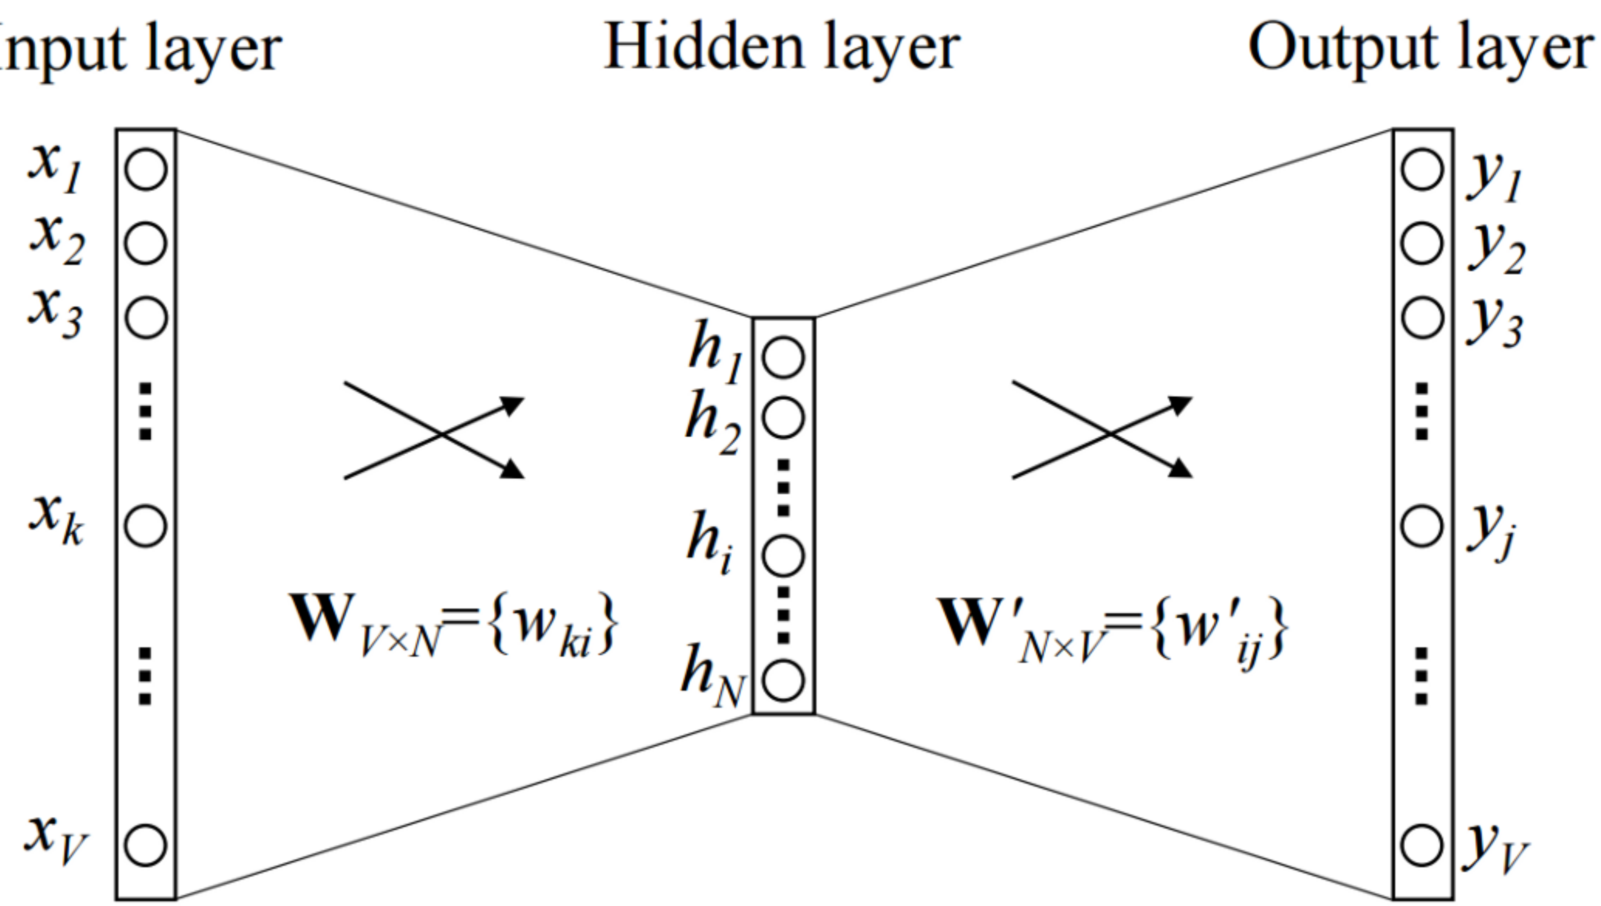

Figure 1: A simple CBOW model with only one word in the context

Mikolov et al, 2013
Rong, 2016

Sidenote: Continuous Bag of Words (CBOW)

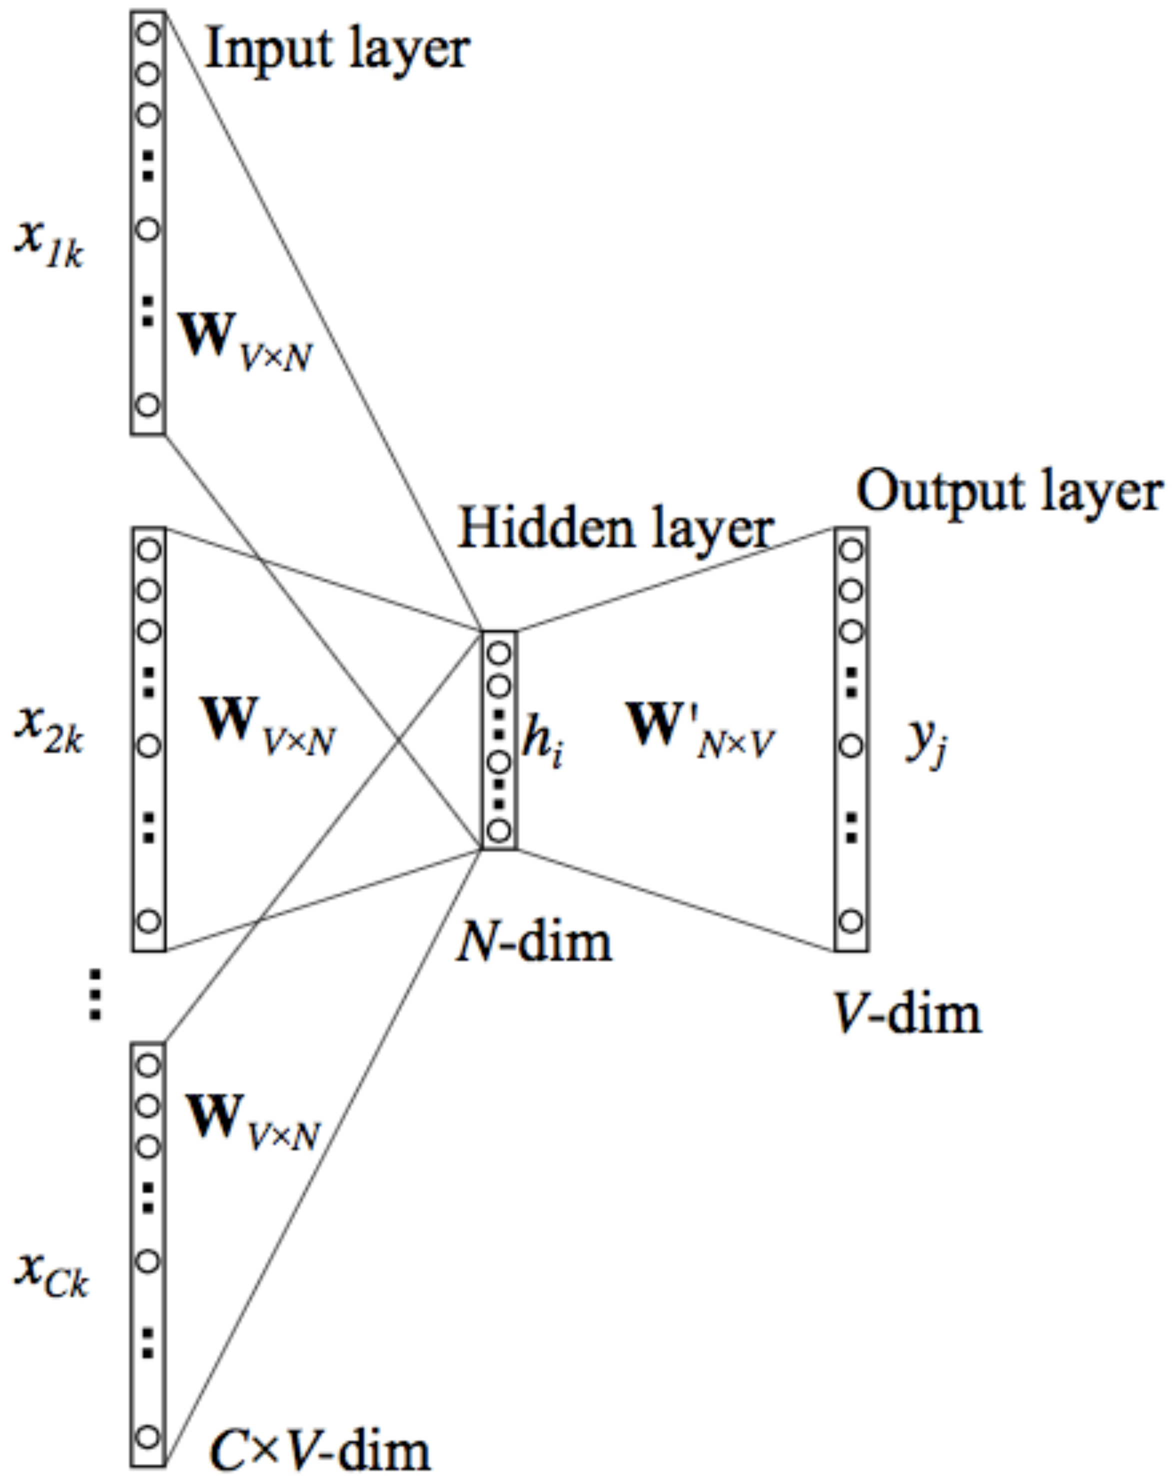

Mikolov et al, 2013
Rong, 2016

2. Inferring Network Structure

Identifying similar words through embeddings:

- Words are high dimensional objects and can be embedded in high dimensional space
- Do this in such a way that words which appear in similar contexts are geometrically close

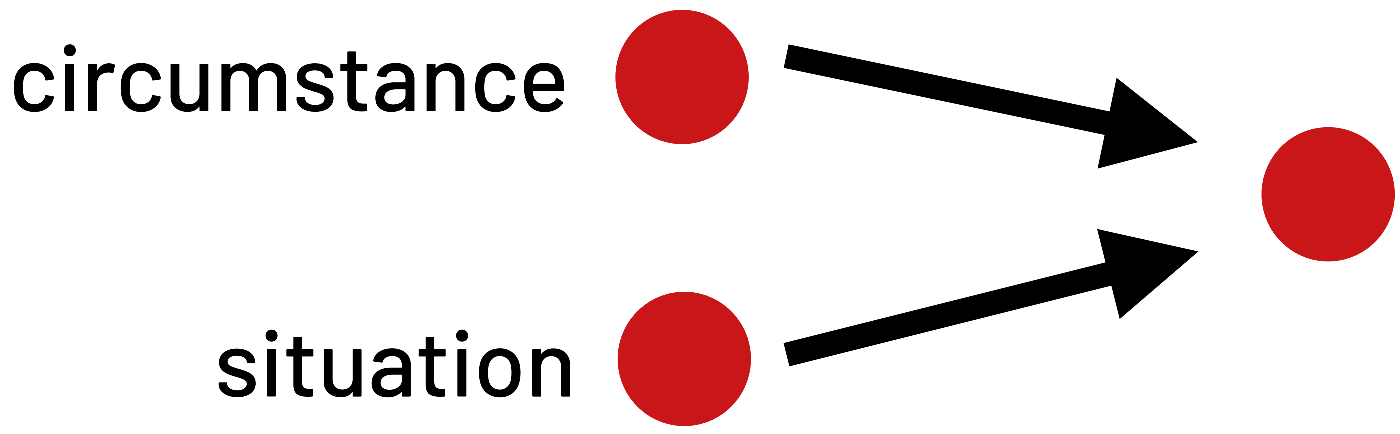

2. Inferring Network Structure

What are the nodes?

➡ Concepts: **“similar words”**

What are the edges?

➡ Connections between concepts (??)

2. Inferring Network Structure

What are the nodes?

➡ Concepts: **“similar words”**

What are the edges?

➡ Connections between **words** (??)

2. Inferring Network Structure

Example:

Bodily autonomy is a basic human right.

2. Inferring Network Structure

Example:

Word co-occurance

Bodily autonomy is a basic human right.

2. Inferring Network Structure

Example:

Word co-occurance

Bodily autonomy is a basic human right.

2. Inferring Network Structure

Example:

Word co-occurance:

Bodily autonomy ~~is a~~ $\underbrace{\text{basic human right.}}$

2. Inferring Network Structure

Example:

Word co-occurance:
Assumes connected concepts are syntactic close

$$\text{Bodily autonomy } \cancel{\text{is a}} \text{ basic } \underbrace{\text{human right.}}$$

2. Inferring Network Structure

Example:

Word co-occurance:
Assumes connected concepts are syntactic close

Bodily autonomy ~~is a~~ basic human right.

2. Inferring Network Structure

Example:

Grammatical structure:

Bodily autonomy is a basic human right.

2. Inferring Network Structure

Example:

Grammatical structure:
Designed to encode implicit connections

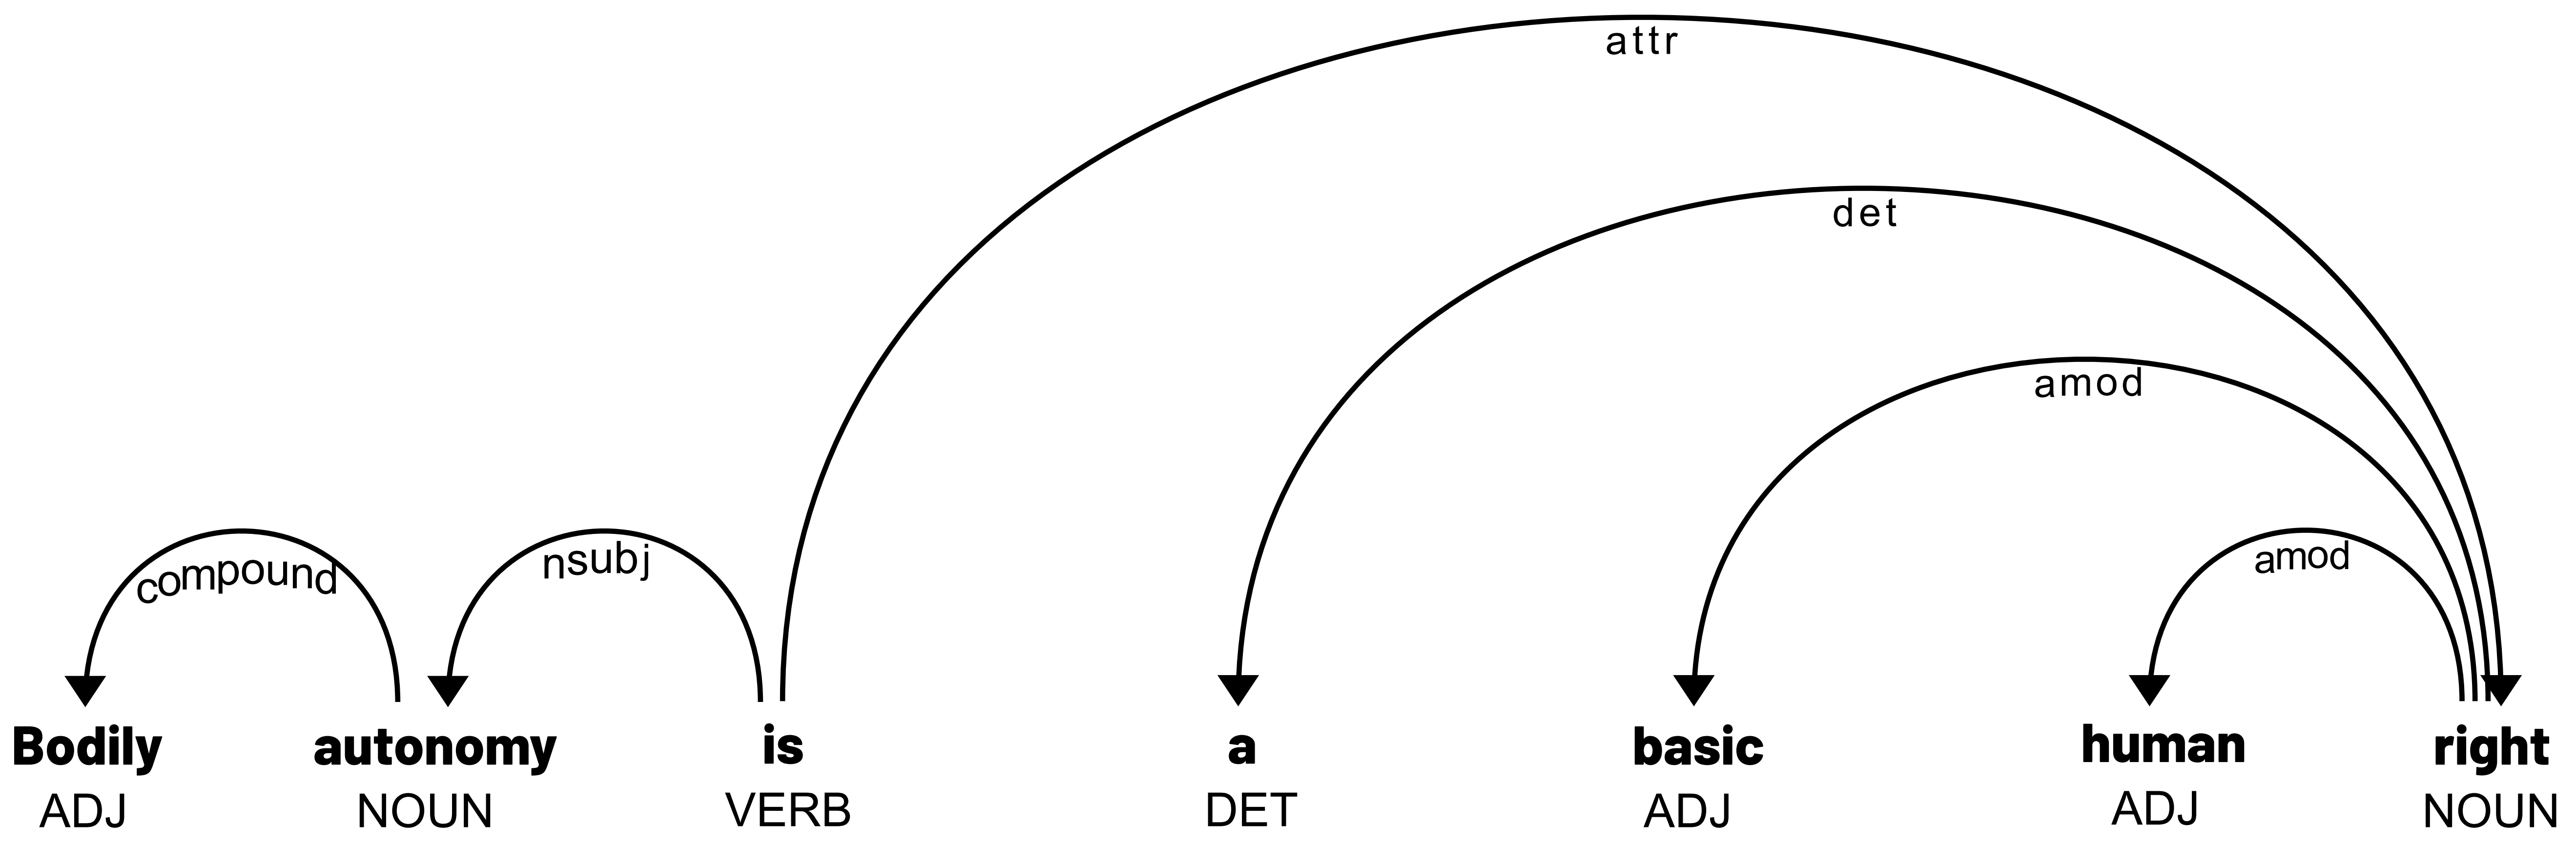

2. Inferring Network Structure

Example:

Grammatical structure:
Designed to encode implicit connections

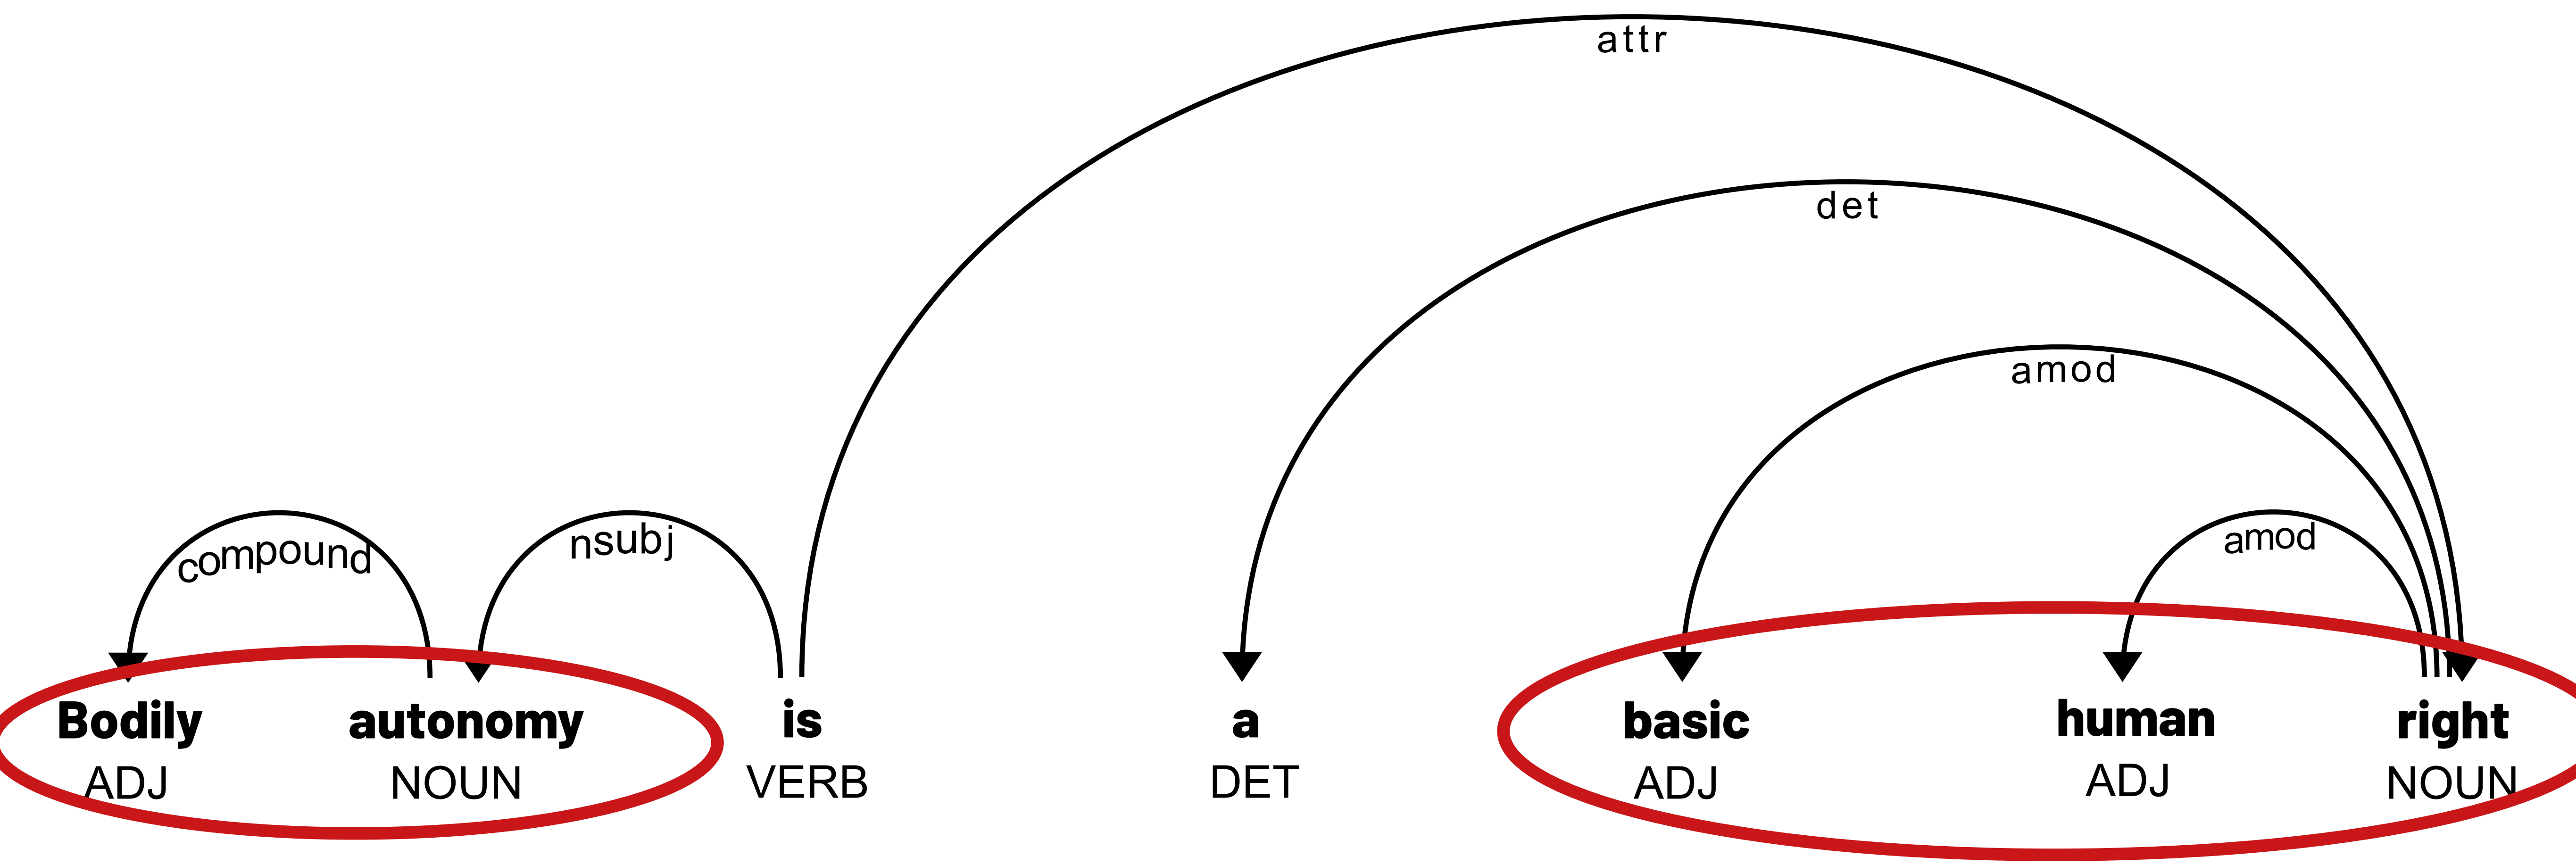

2. Inferring Network Structure

Example:

Grammatical structure:
Designed to encode implicit connections

Bodily autonomy is a basic human right.

2. Inferring Network Structure

Model steps

1. Infer Part of Speech tags and grammatical structure

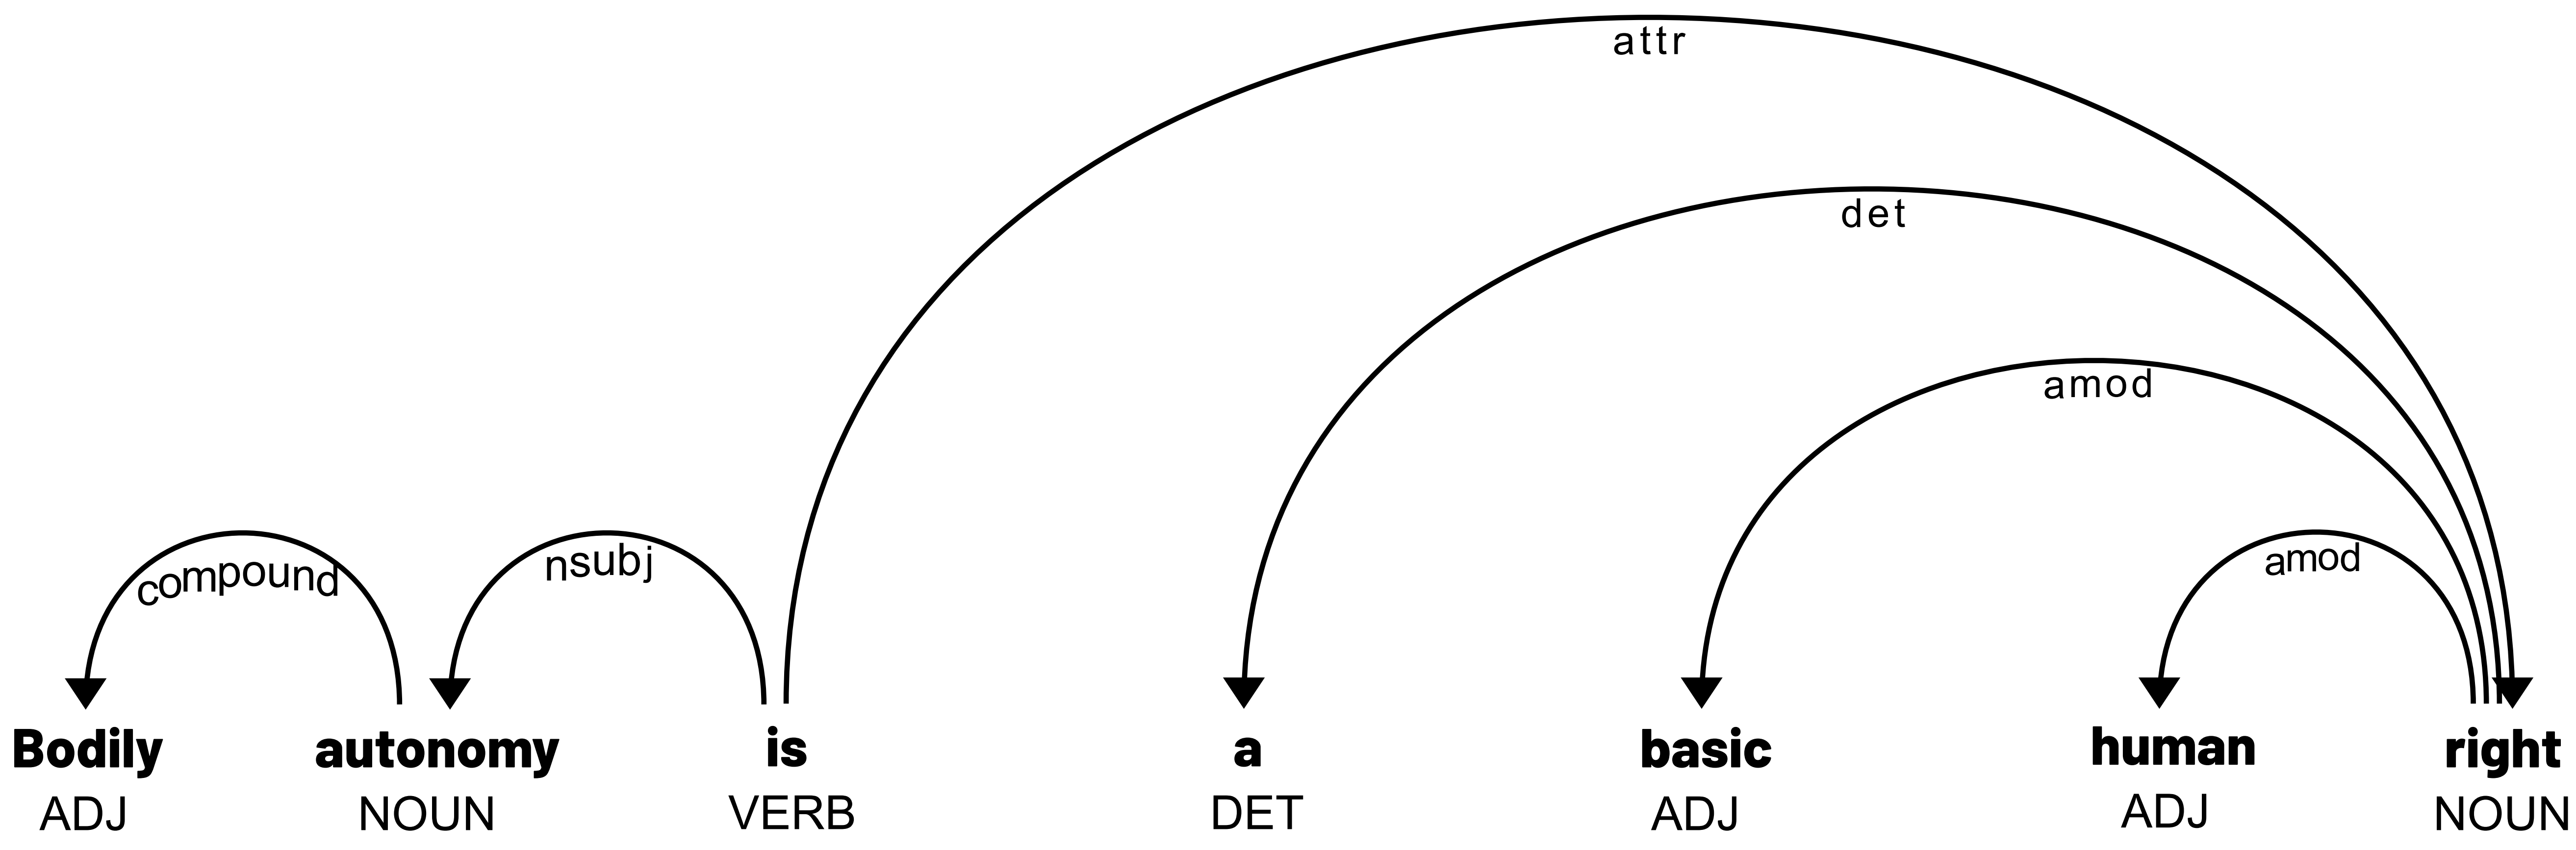

2. Inferring Network Structure

Model steps

1. Infer Part of Speech tags and grammatical structure
2. Turn negative words into negative ties

X is Y

2. Inferring Network Structure

Model steps

1. Infer Part of Speech tags and grammatical structure
2. Turn negative words into negative ties

X is **not** Y

2. Inferring Network Structure

Model steps

1. Infer Part of Speech tags and grammatical structure
2. Turn negative words into negative ties
3. Remove stopwords, maintaining network structure

The X is a Y

2. Inferring Network Structure

Model steps

1. Infer Part of Speech tags and grammatical structure
2. Turn negative words into negative ties
3. Remove stopwords, maintaining network structure

The X is a Y

2. Inferring Network Structure

Model steps

1. Infer Part of Speech tags and grammatical structure
2. Turn negative words into negative ties
3. Remove stopwords, maintaining network structure

~~The~~ X is ~~a~~ Y

2. Inferring Network Structure

Model steps

1. Infer Part of Speech tags and grammatical structure
2. Turn negative words into negative ties
3. Remove stopwords, maintaining network structure
4. Merge similar words using embeddings

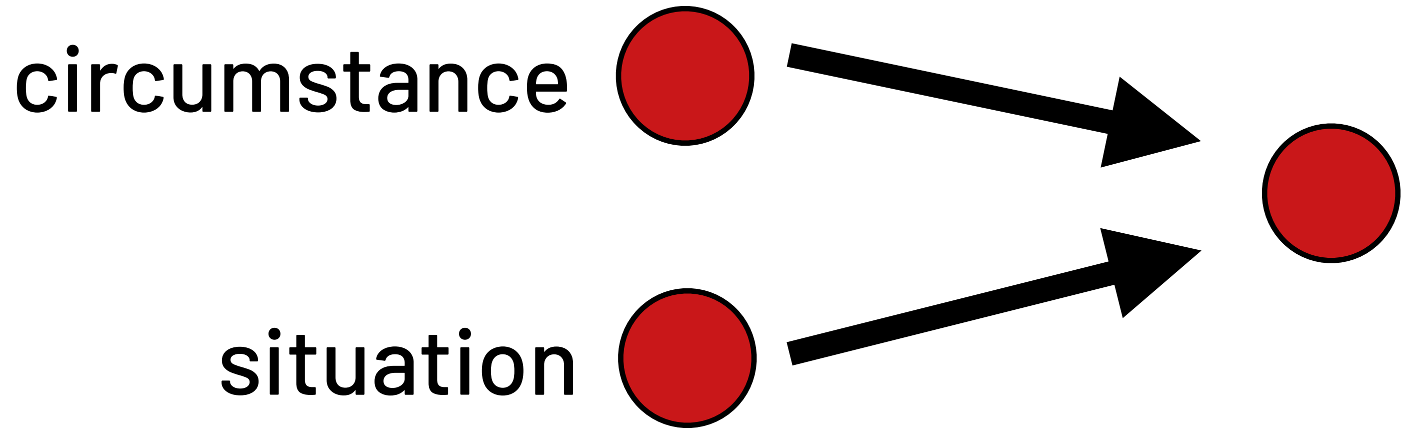

Sample Inferred Networks

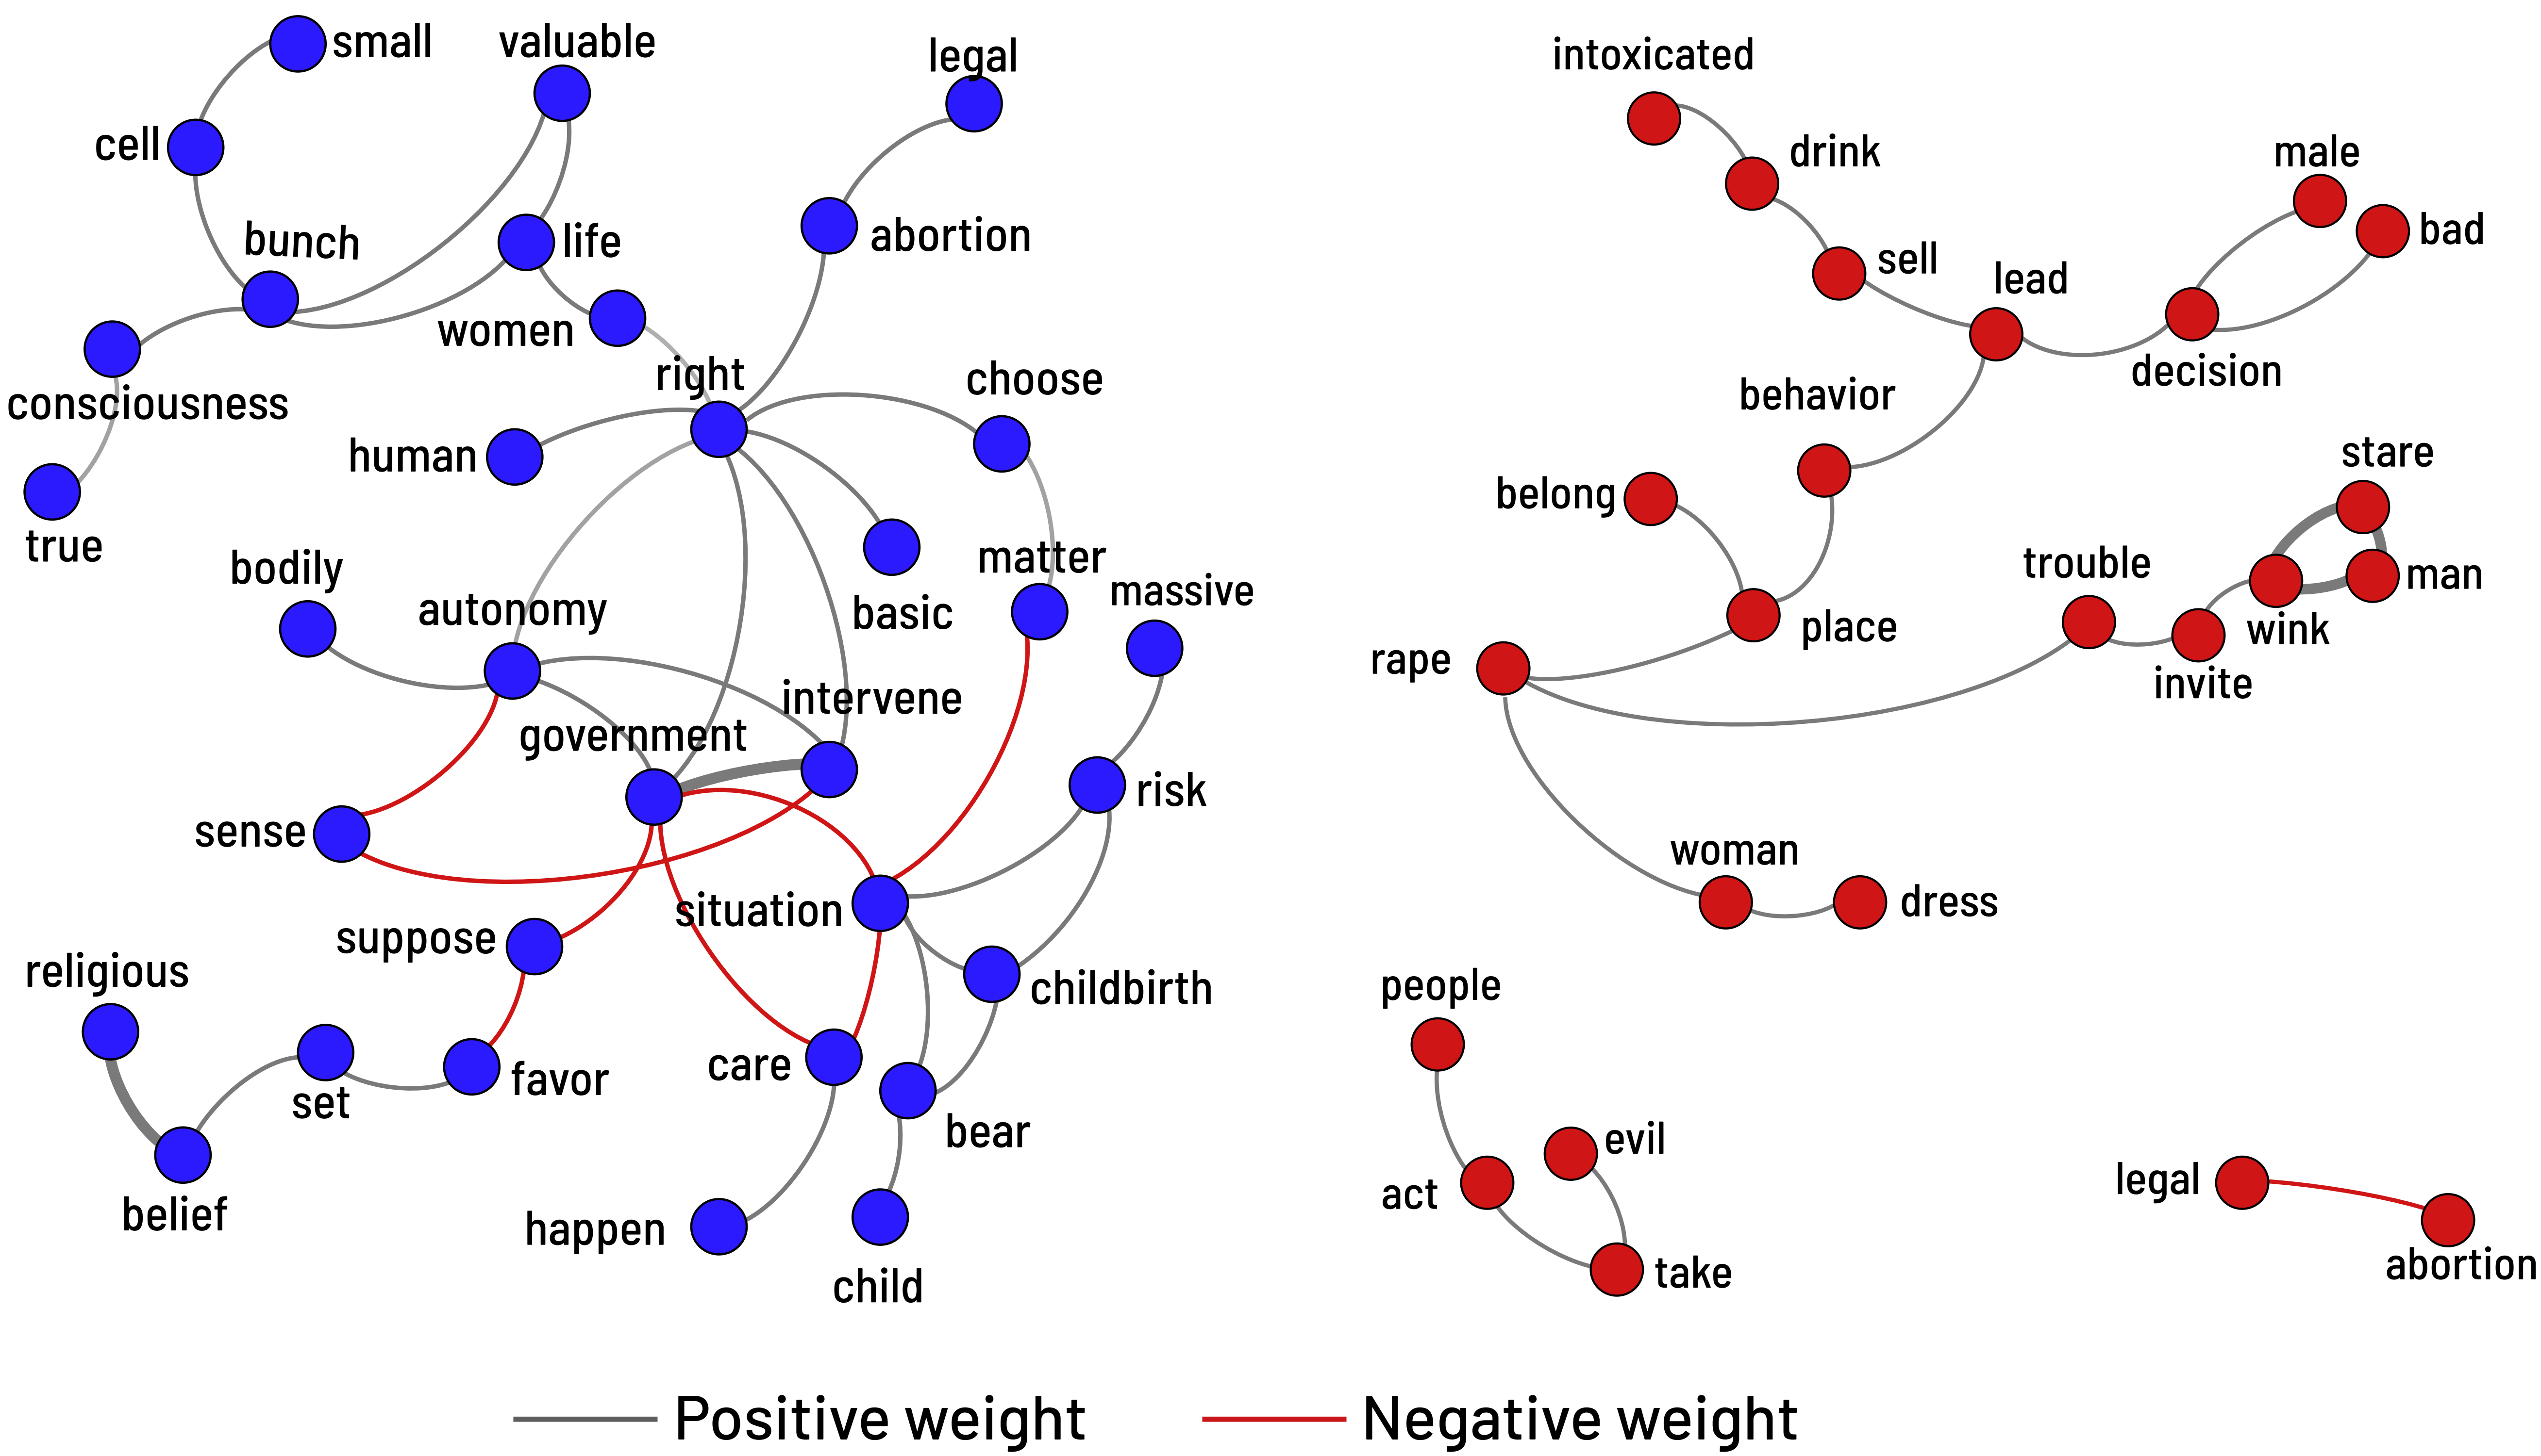

Measuring Network Similarity

Measuring Network Similarity

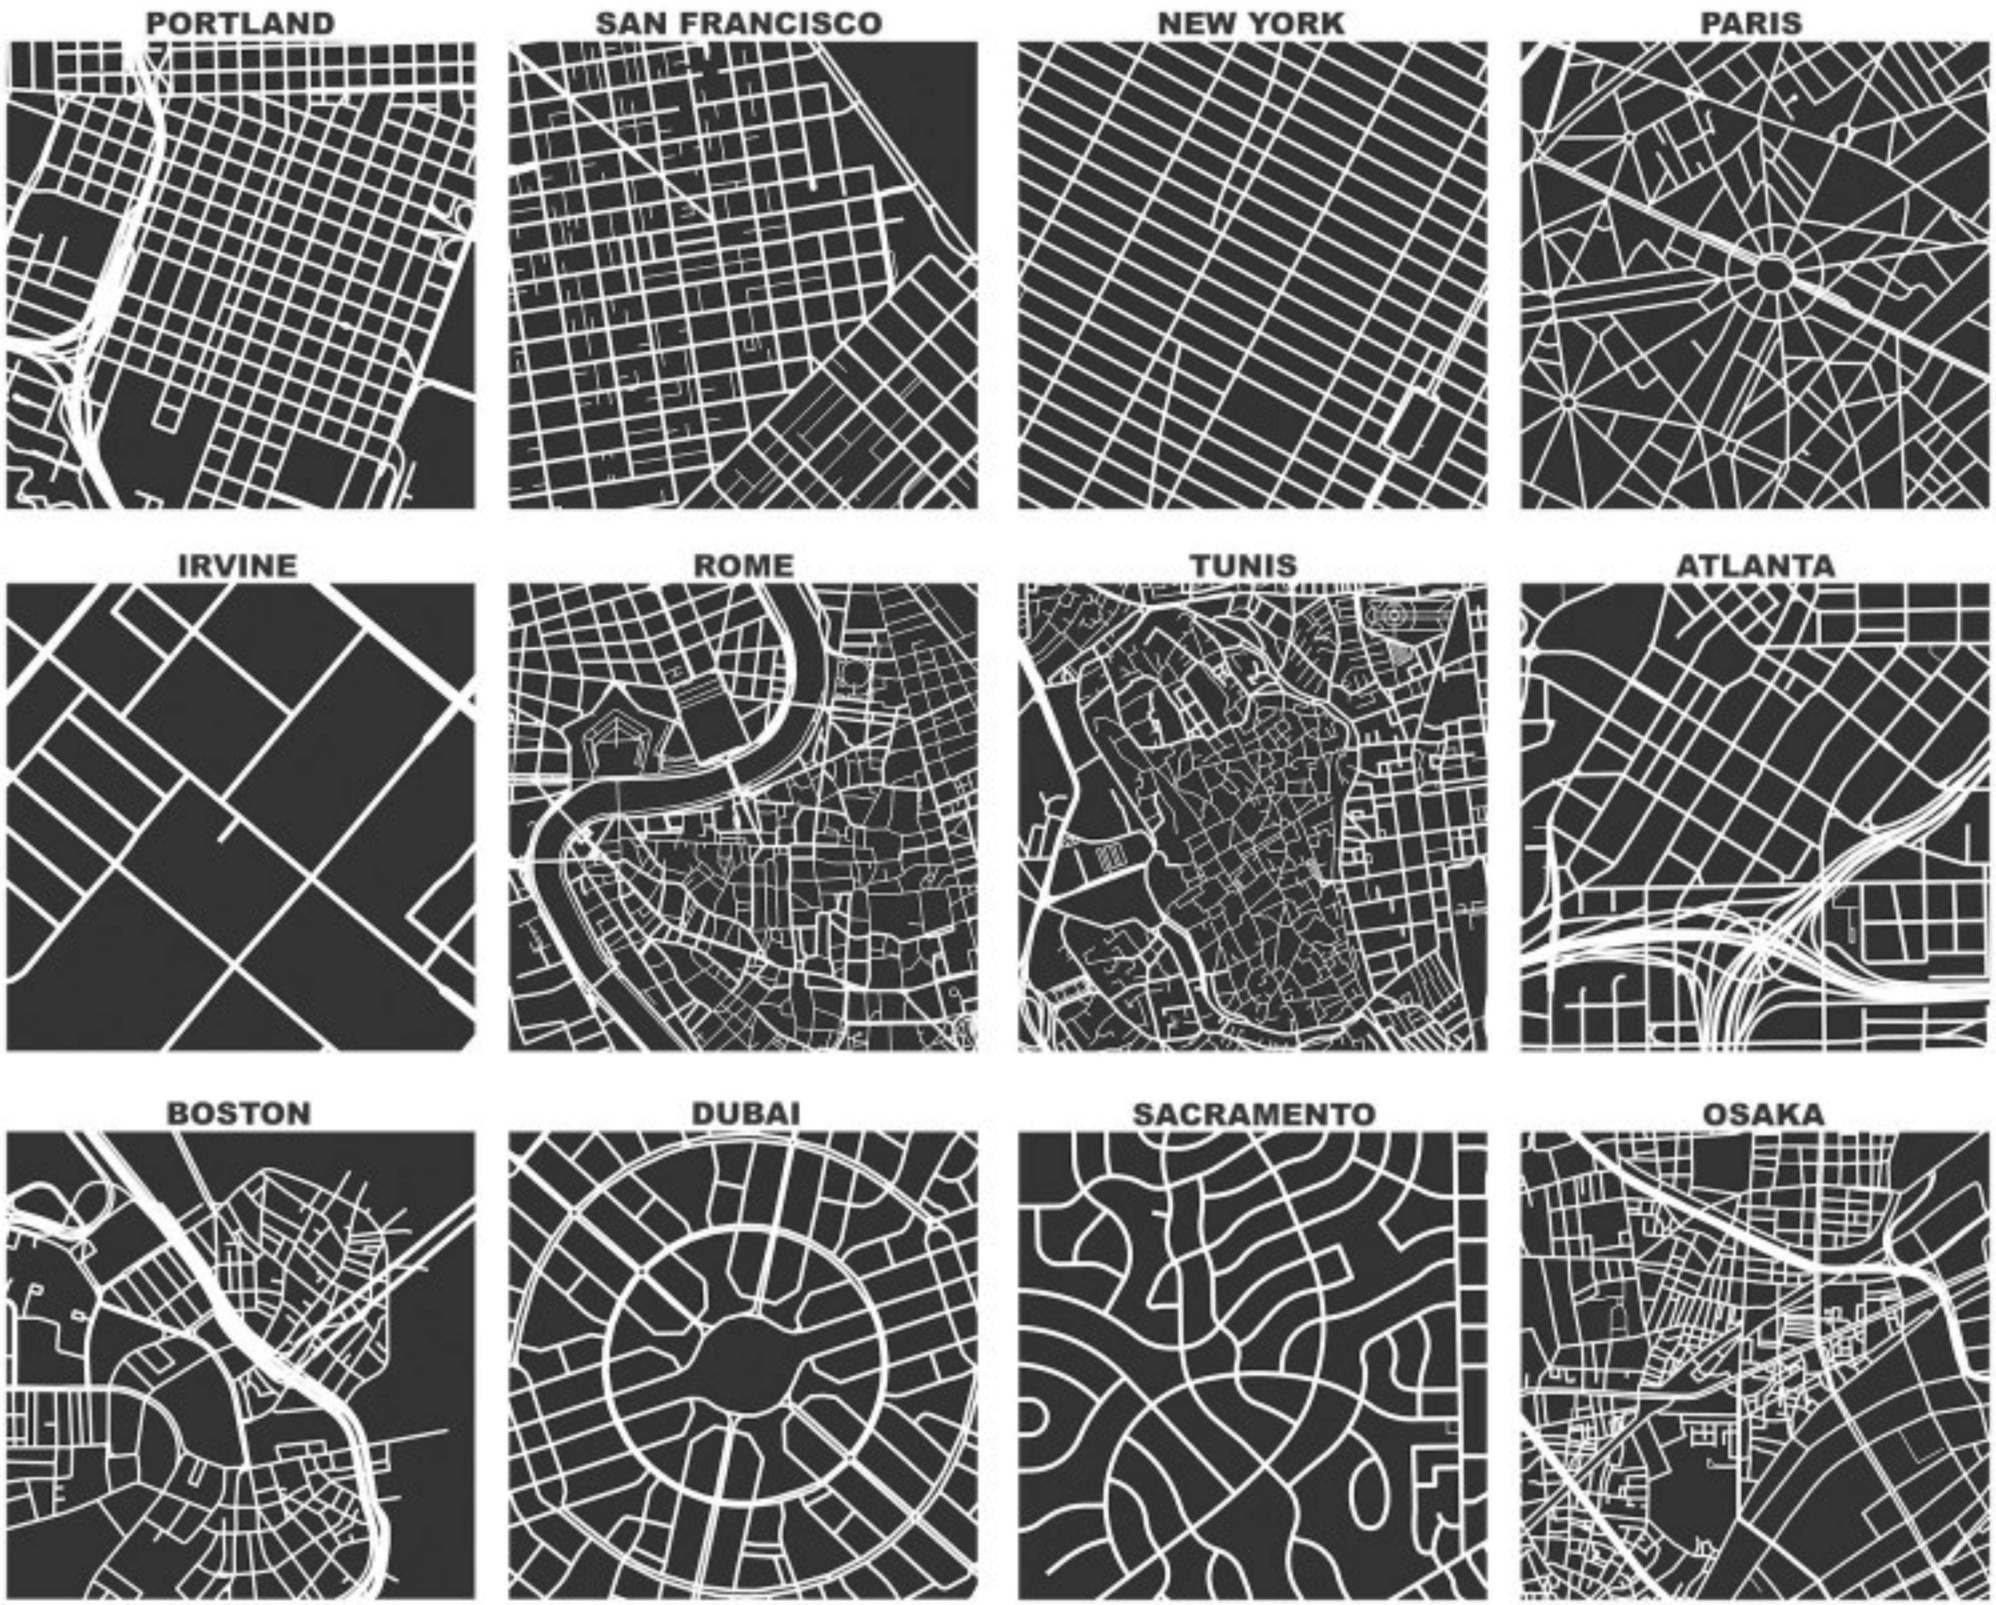

Boeing, 2017

Measuring Network Similarity

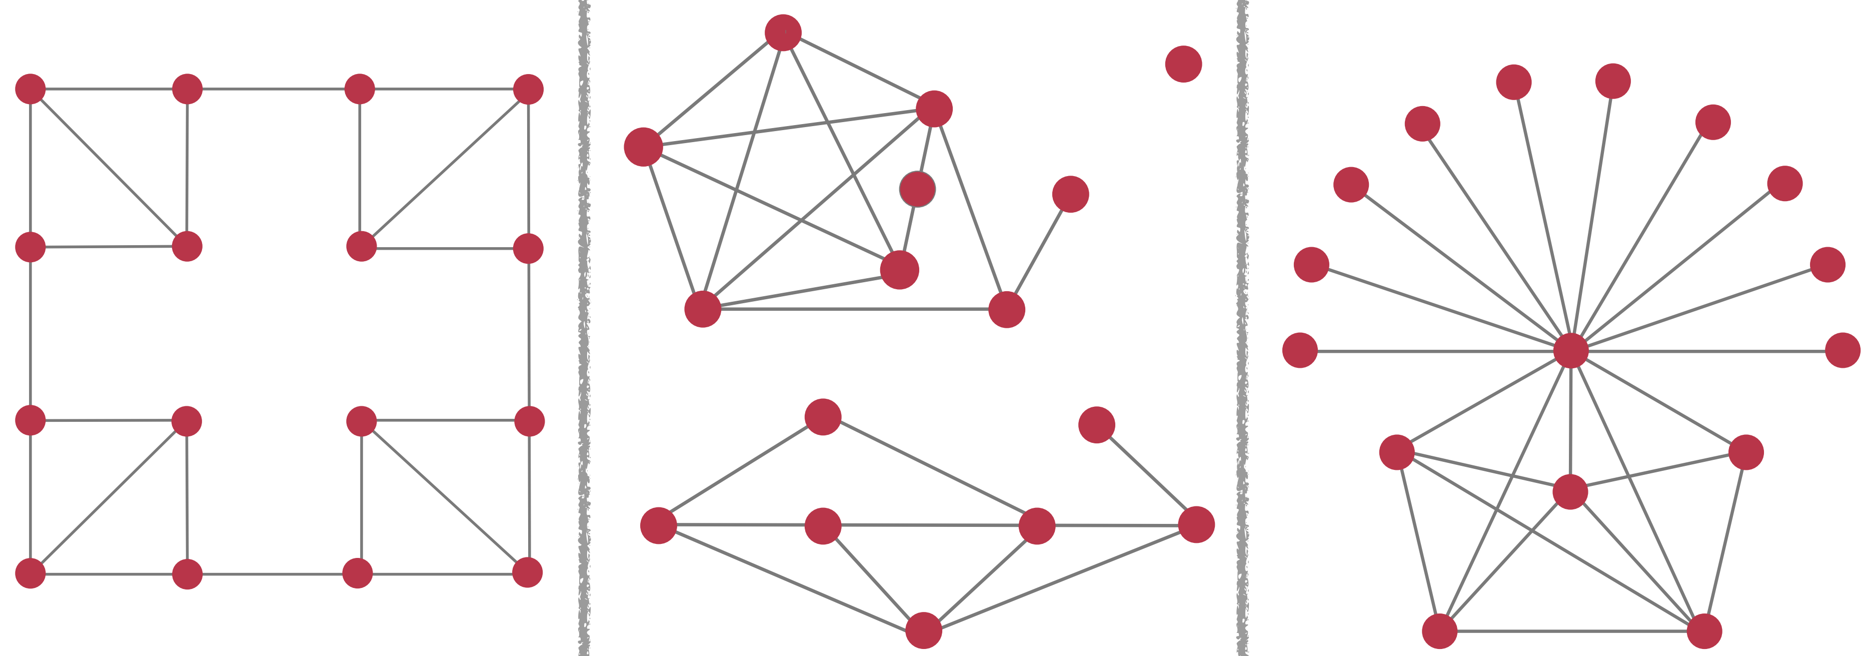

Homogeneous **Heterogenous**

density	0.2	0.2	0.2
k avg	3.0	3.0	3.0
clustering	0.2	0.4	0.3
giant component	1.0	0.4	1.0
entropy	0.0	1.5	1.0
disassortativity	-1.0	-0.1	0.7
k std	0.0	3.0	3.5

The Structure of Reasoning

Roadmap:

1. Elaborate on “structure”
2. Define approach for inferring and measuring structure
3. **Demonstrate potential for behavioral insights — using two distinct datasets**

The Structure of Reasoning

Roadmap:

1. Elaborate on “structure”
2. Define approach for inferring and measuring structure
3. **Demonstrate potential for behavioral insights — using two distinct datasets**

Does the structure of expressed reasons convey useful information?

Data

Data

1. Experiment and survey
   - 100 subjects, recruited through MTurk
   - Three methods of inferring networks, for two of three topics: (1) abortion (2) healthcare (3) childrearing
   - Extensive demographic and personality survey

Shugars, Beauchamp, and Levine; 2019

Data

1. Experiment and survey
   - 100 subjects, recruited through MTurk
   - Three methods of inferring networks, for two of three topics: (1) abortion (2) healthcare (3) childrearing
   - Extensive demographic and personality survey

2. Ideological “Turing test"
   - 1000 subjects, recruited by YouGov
   - Asked to provide “liberal” and “conservative” positions on one of three topics
     (1) abortion (2) minimum wage (3) national defense

Hopkins and Noel, 2016

Data

1. Experiment and survey
    - 100 subjects, recruited through MTurk
    - Three methods of inferring networks, for two of three topics: (1) abortion (2) healthcare (3) childrearing
    - Extensive demographic and personality survey

Research Questions

- Does structure meaningfully correlate to known personality traits?

Shugars, Beauchamp, and Levine; 2019

3. Potential for Behavioral Insights

Research Questions

- Does structure meaningfully correlate to known personality traits?

  - Purity (Moral Foundations)
  - Authority (Moral Foundations)
  - Ingroup (Moral Foundations)
  - Harm (Moral Foundations)
  - Fairness (Moral Foundations)
  - Progressivism (Moral Foundations)
  - Extroversion (Big 5)
  - Agreeableness (Big 5)
  - Neuroticism (Big 5)
  - Conscientiousness (Big 5)
  - Openness (Big 5)

  - Ideology: Conservative
  - Political Knowledge
  - Deliberativeness

Haidt & Joseph, 2008; John & Srivastava, 1999

Gastil et al., 2012; Carpini & Keeter, 1993 ; Pew, 2017

3. Potential for Behavioral Insights

Research Questions

- Does structure meaningfully correlate to known personality traits?

$$s = \beta p + \alpha_t + \epsilon$$

3. Potential for Behavioral Insights

Research Questions

- Does structure meaningfully correlate to known personality traits?

$$\underbrace{s}_{\text{Network statistic}} = \beta \underbrace{p}_{\text{Personality measure}} + \underbrace{\alpha_t}_{\text{Topic fixed effect}} + \epsilon$$

3. Potential for Behavioral Insights

Research Questions

- Does structure meaningfully correlate to known personality traits?

$$\underbrace{s}_{\text{Network statistic}} = \beta \underbrace{p}_{\text{Personality measure}} + \underbrace{\alpha_t}_{\text{Topic fixed effect}} + \epsilon$$

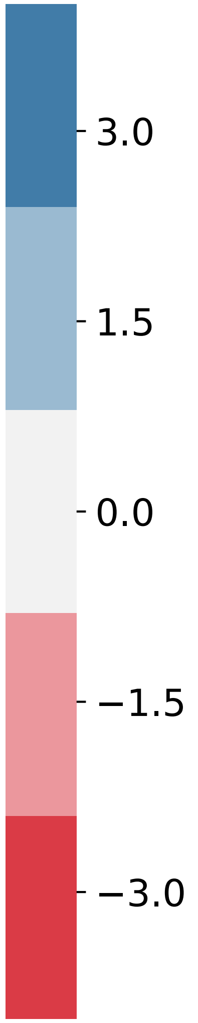

3. Potential for Behavioral Insights

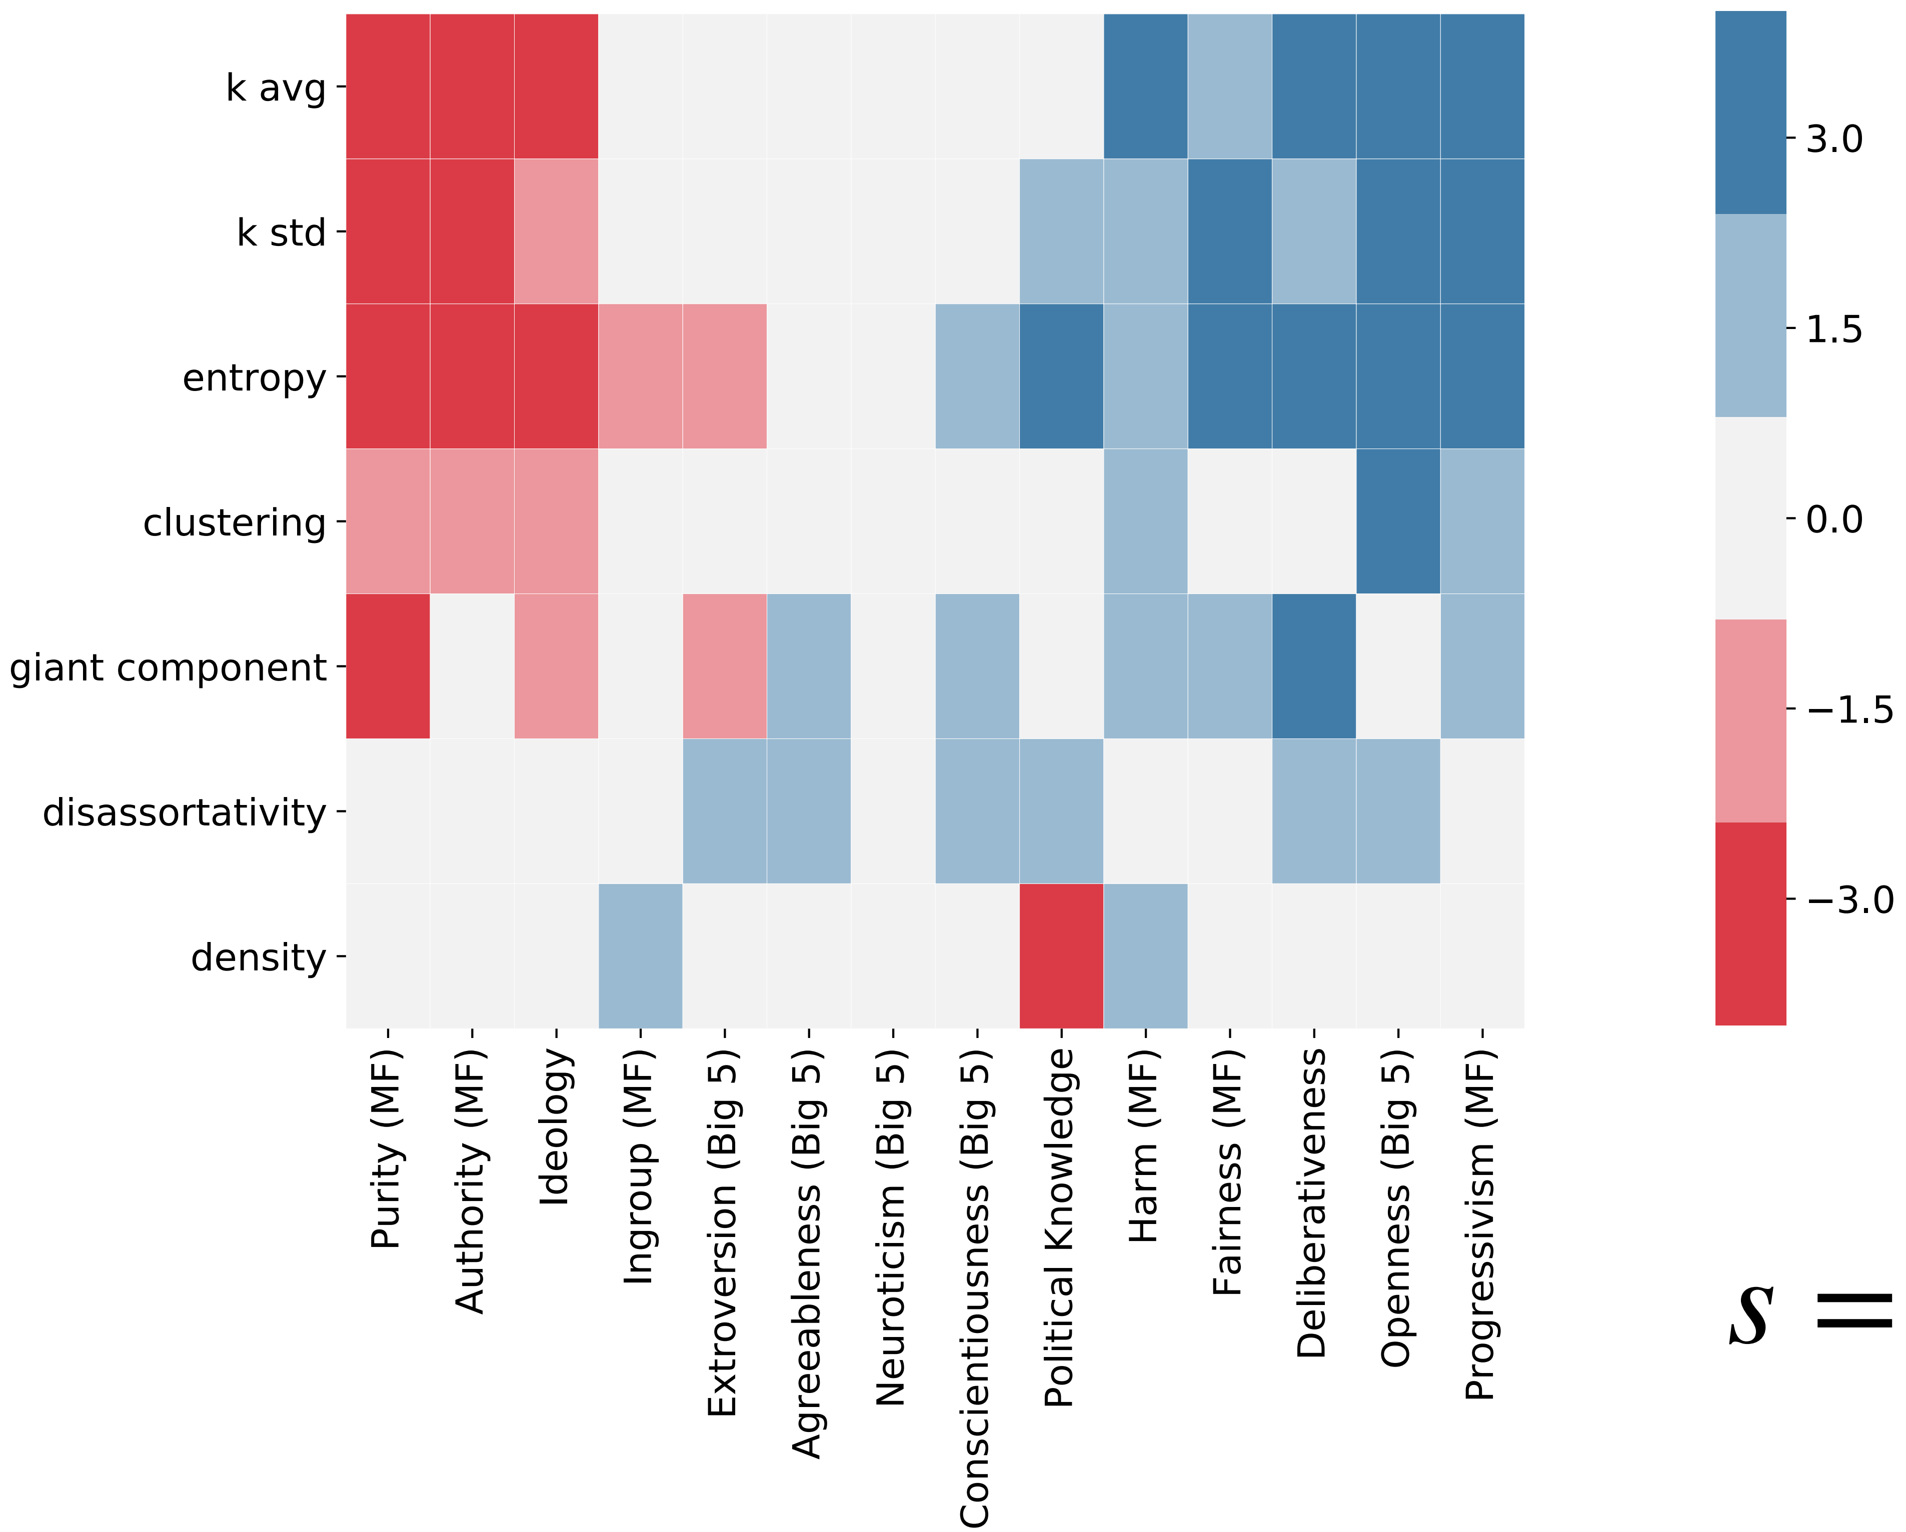

$$s = \beta p + \alpha_t + \epsilon$$

3. Potential for Behavioral Insights

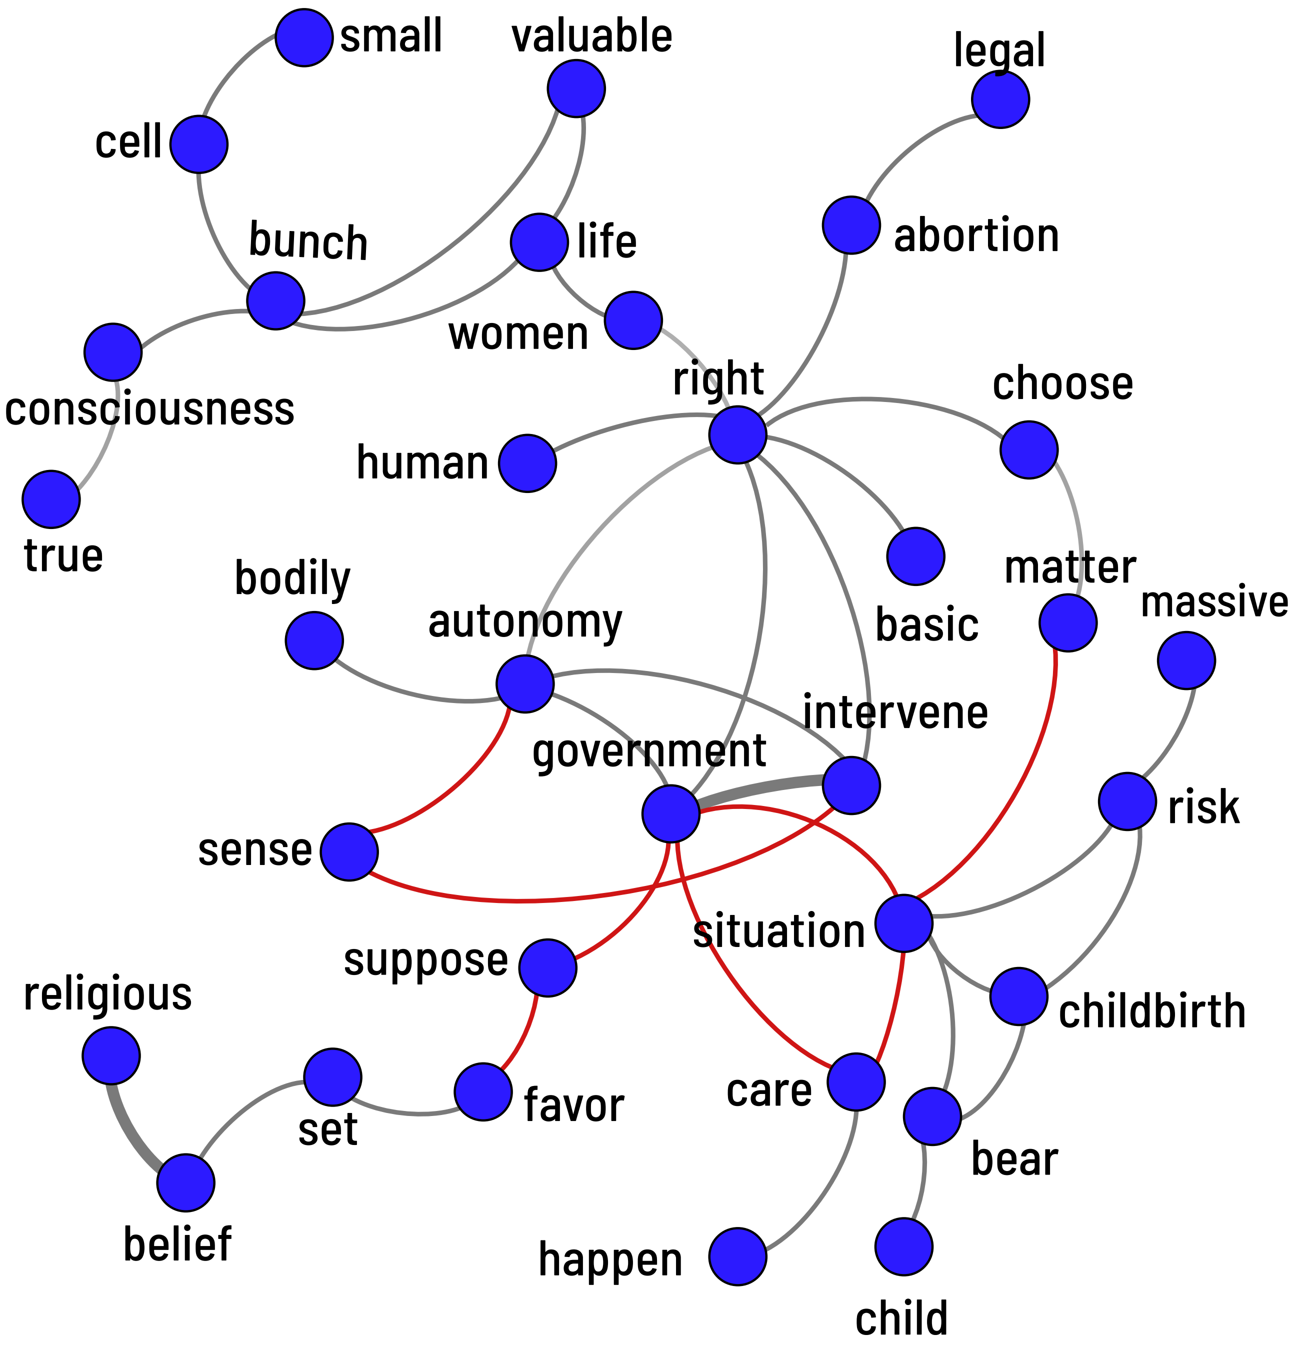

Complexity

Hierarchy

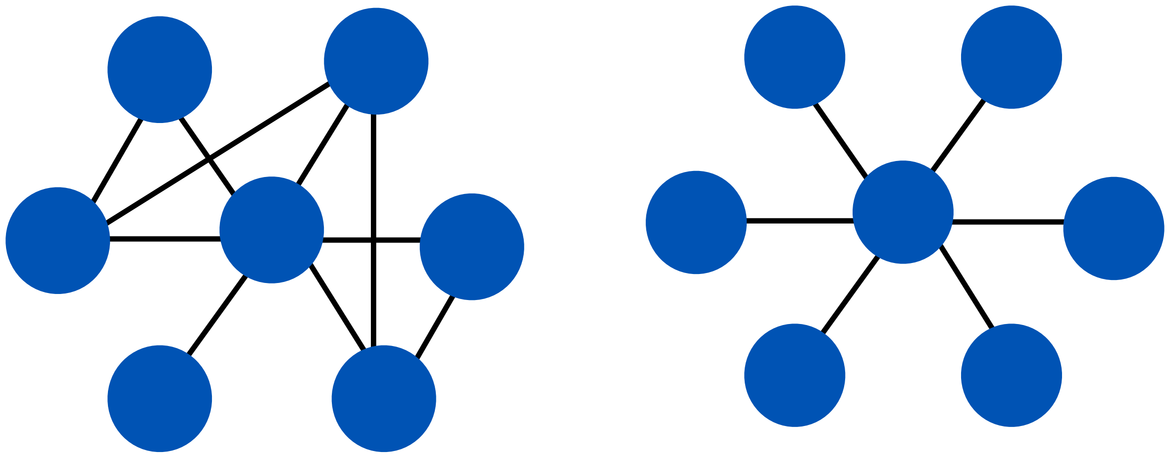

3. Potential for Behavioral Insights

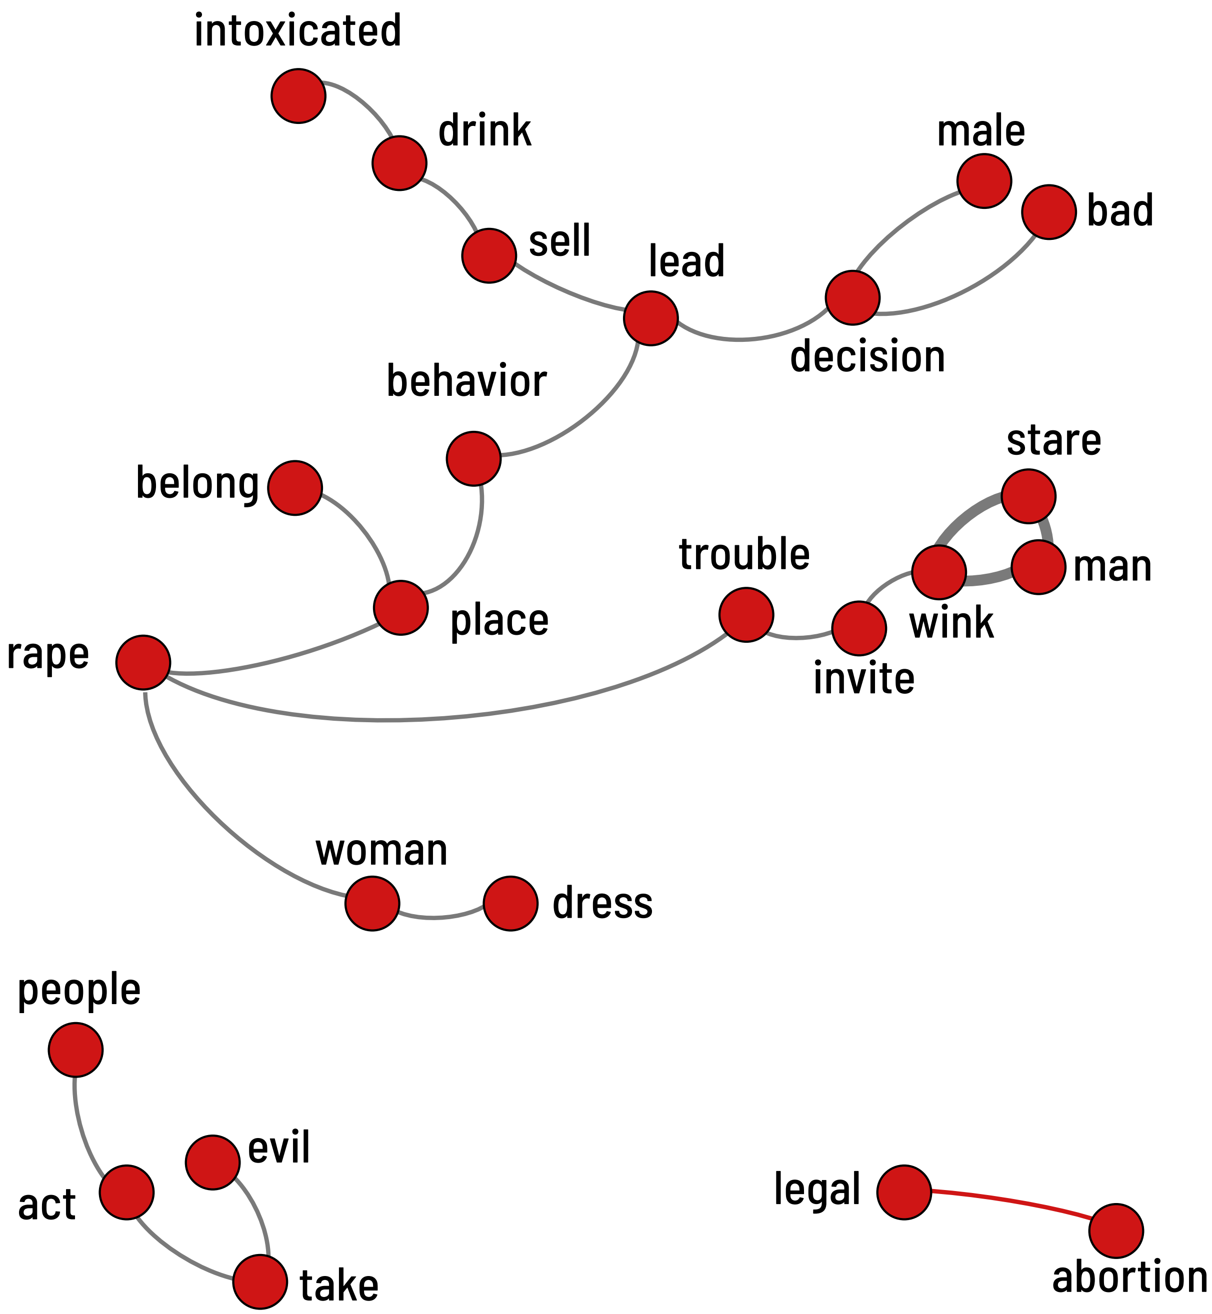

(lack of)
Connectivity

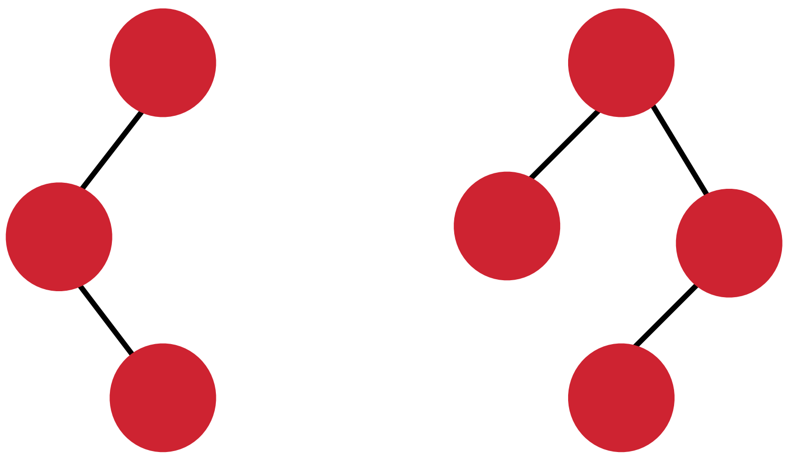

Data

1. Experiment and survey
    - 100 subjects, recruited through MTurk
    - Three methods of inferring networks, for two of three topics: (1) abortion (2) healthcare (3) childrearing
    - Extensive demographic and personality survey

Research Questions

- Does structure meaningfully correlate to known personality traits? **Yes.**

Shugars, Beauchamp, and Levine; 2019

Data

2. Ideological “Turing test"

- 1000 subjects, recruited by YouGov
- Asked to provide “liberal” and “conservative” positions on one of three topics
(1) abortion (2) minimum wage (3) national defense

Hopkins and Noel, 2016

Data

2. Ideological “Turing test"
    - 1000 subjects, recruited by YouGov
    - Asked to provide “liberal” and “conservative” positions on one of three topics
      (1) abortion (2) minimum wage (3) national defense

Research Questions

- Is structure driven by ideology or by individual traits
- Does structure suggest argument quality?

Hopkins and Noel, 2016

3. Potential for Behavioral Insights

Each subject provided 2 networks:

- Liberal position & conservative position

My liberal essay v.
My conservative essay

My liberal essay v.
Your liberal essay

Which are more similar?

3. Potential for Behavioral Insights

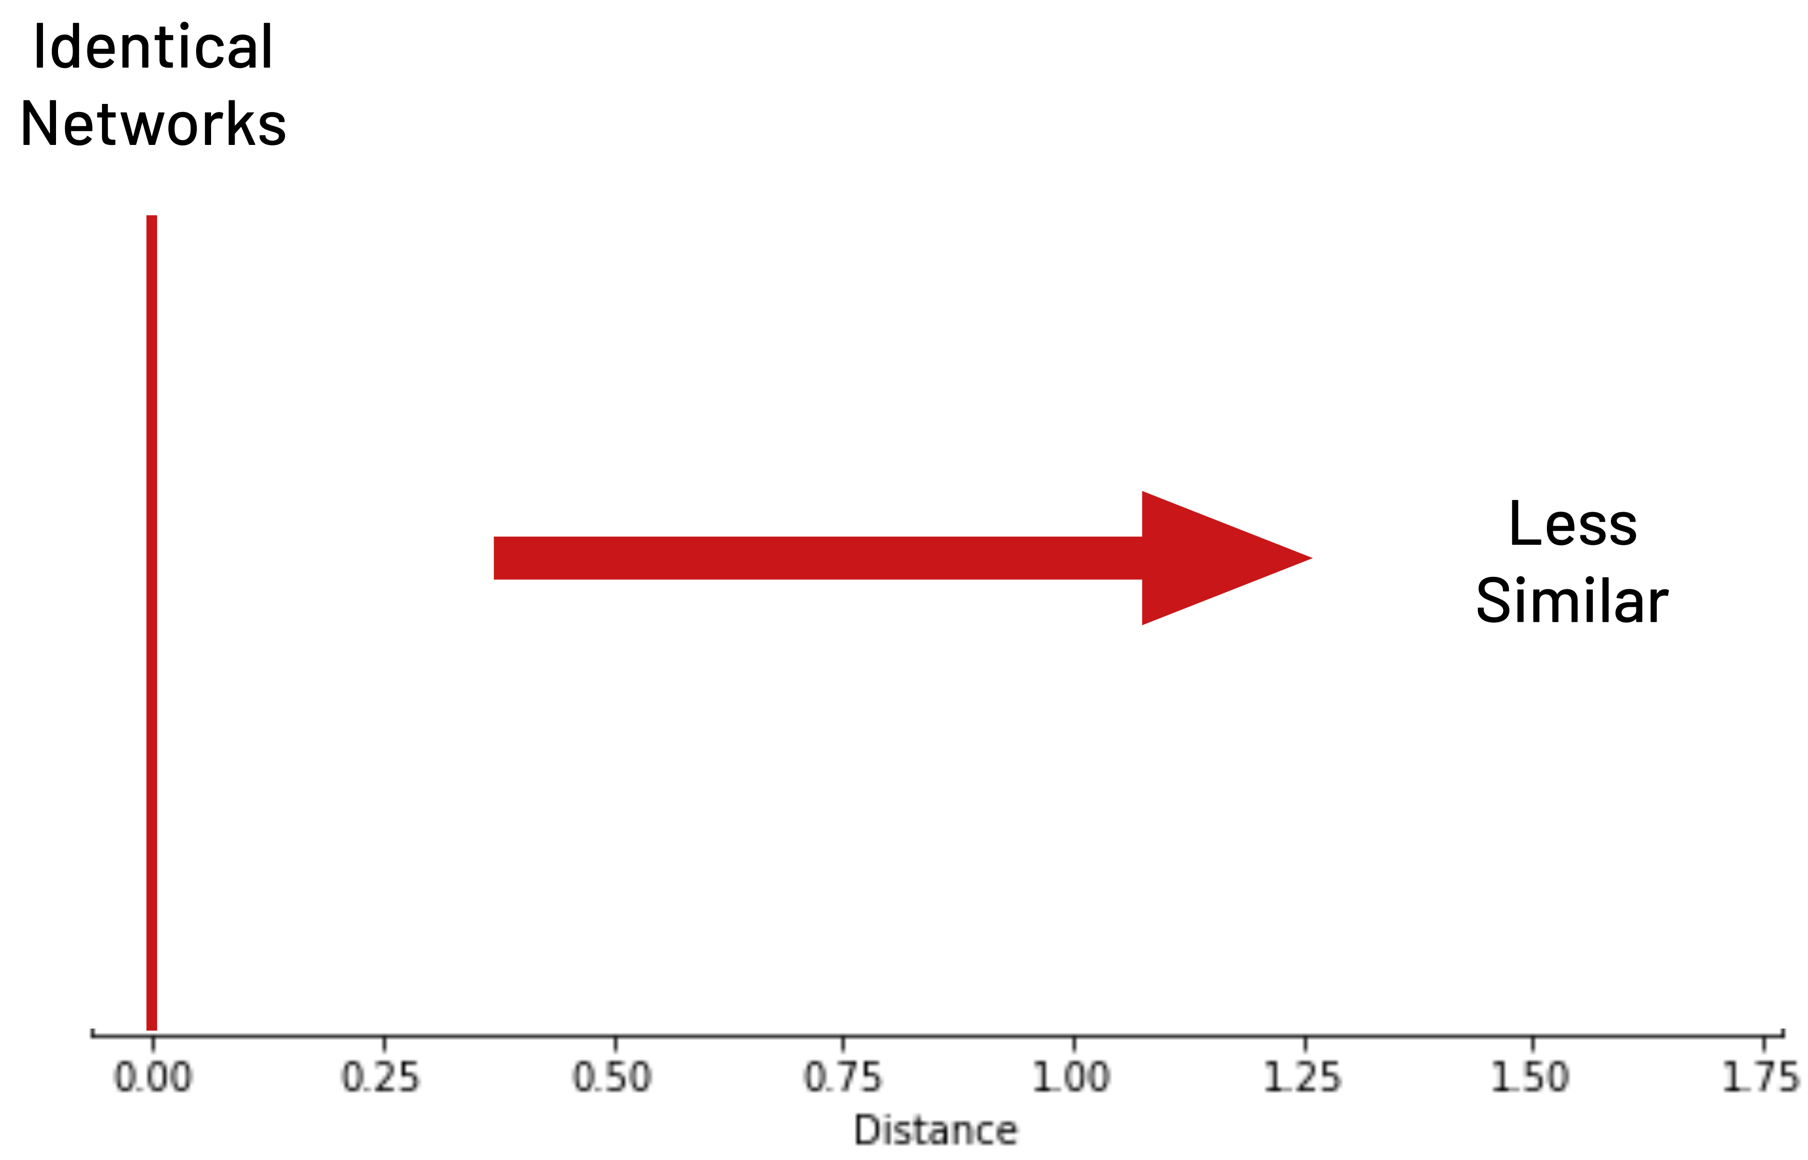

3. Potential for Behavioral Insights

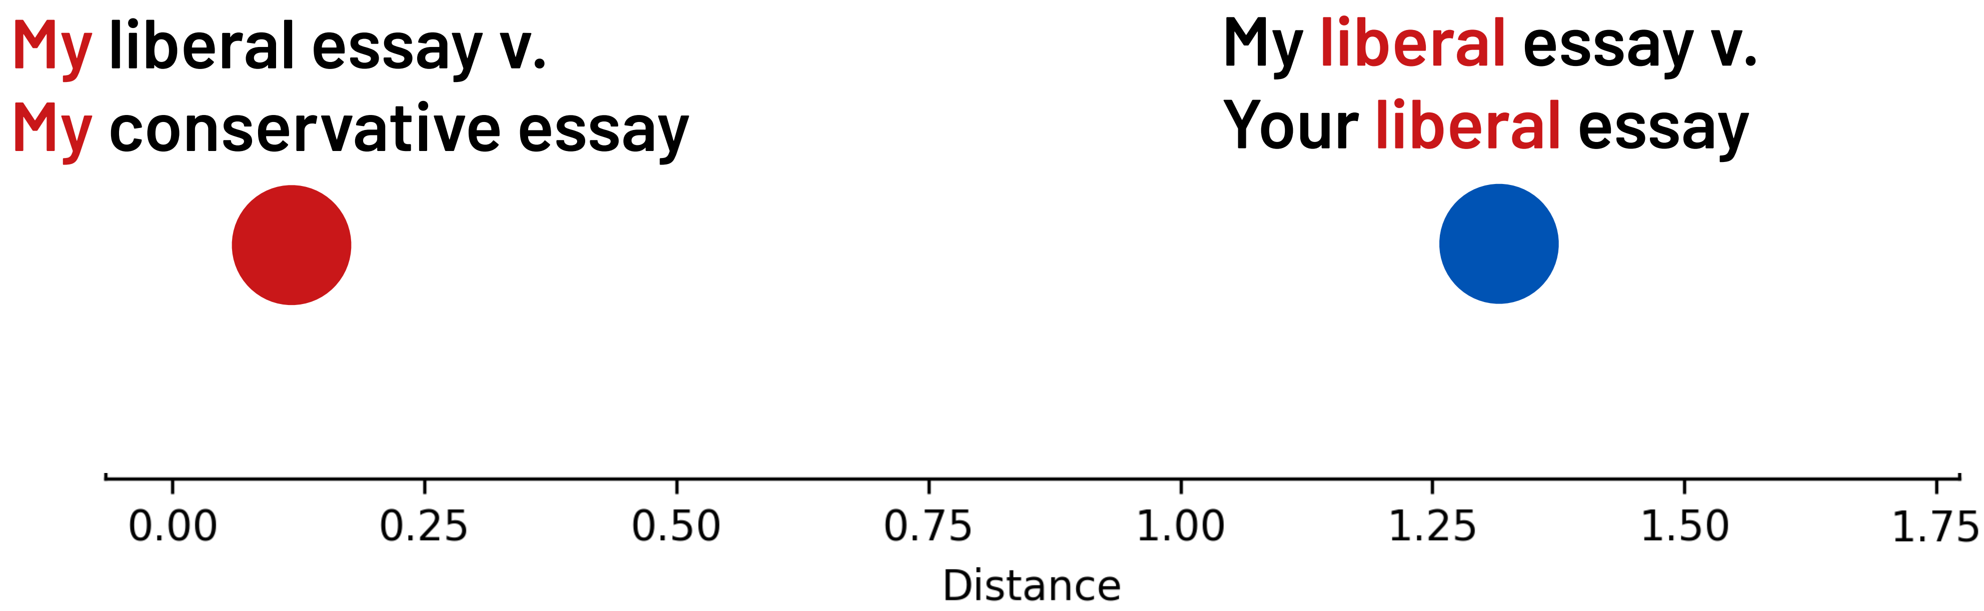

3. Potential for Behavioral Insights

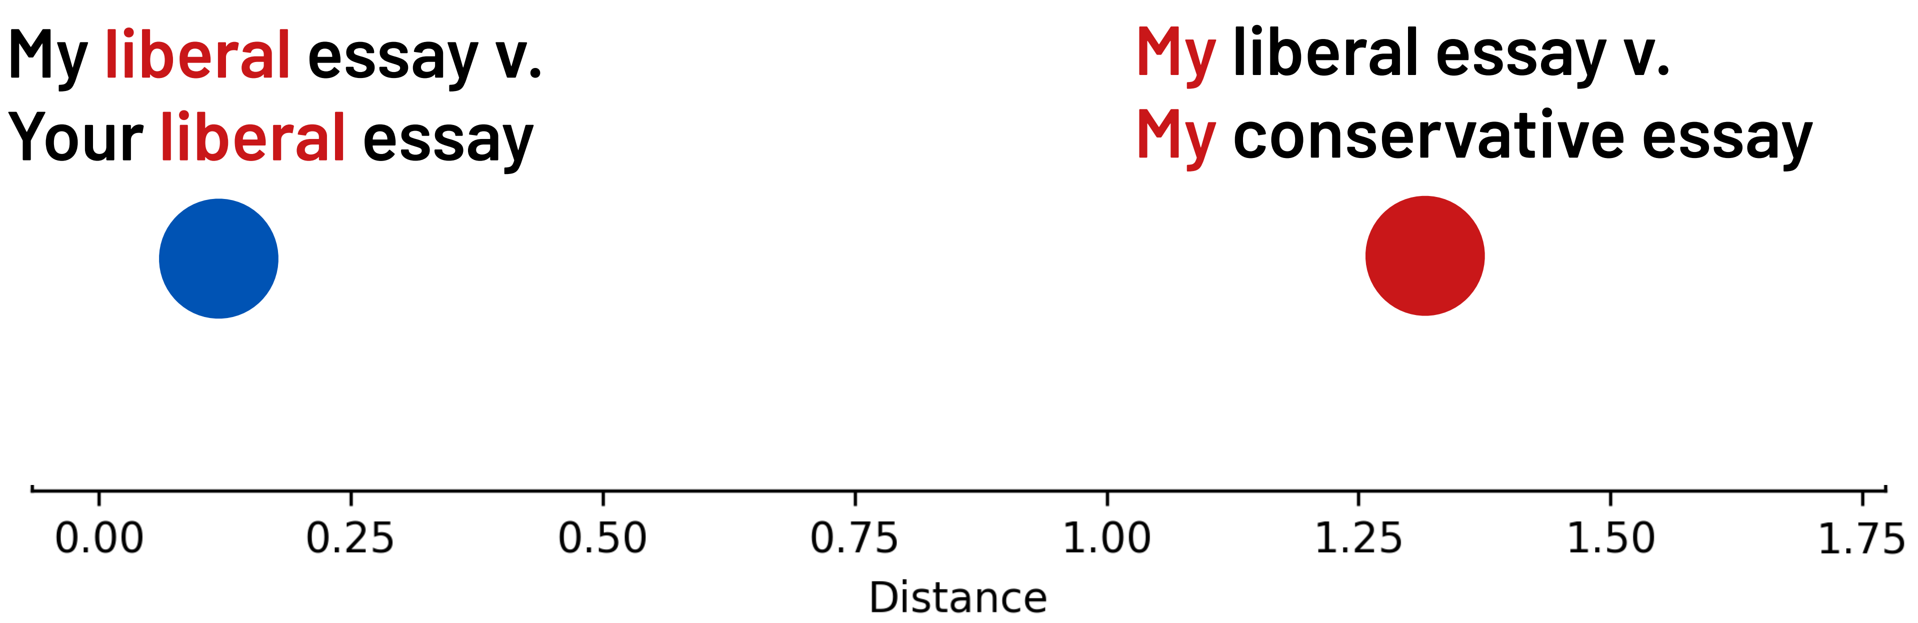

R2: Sources of Similarity

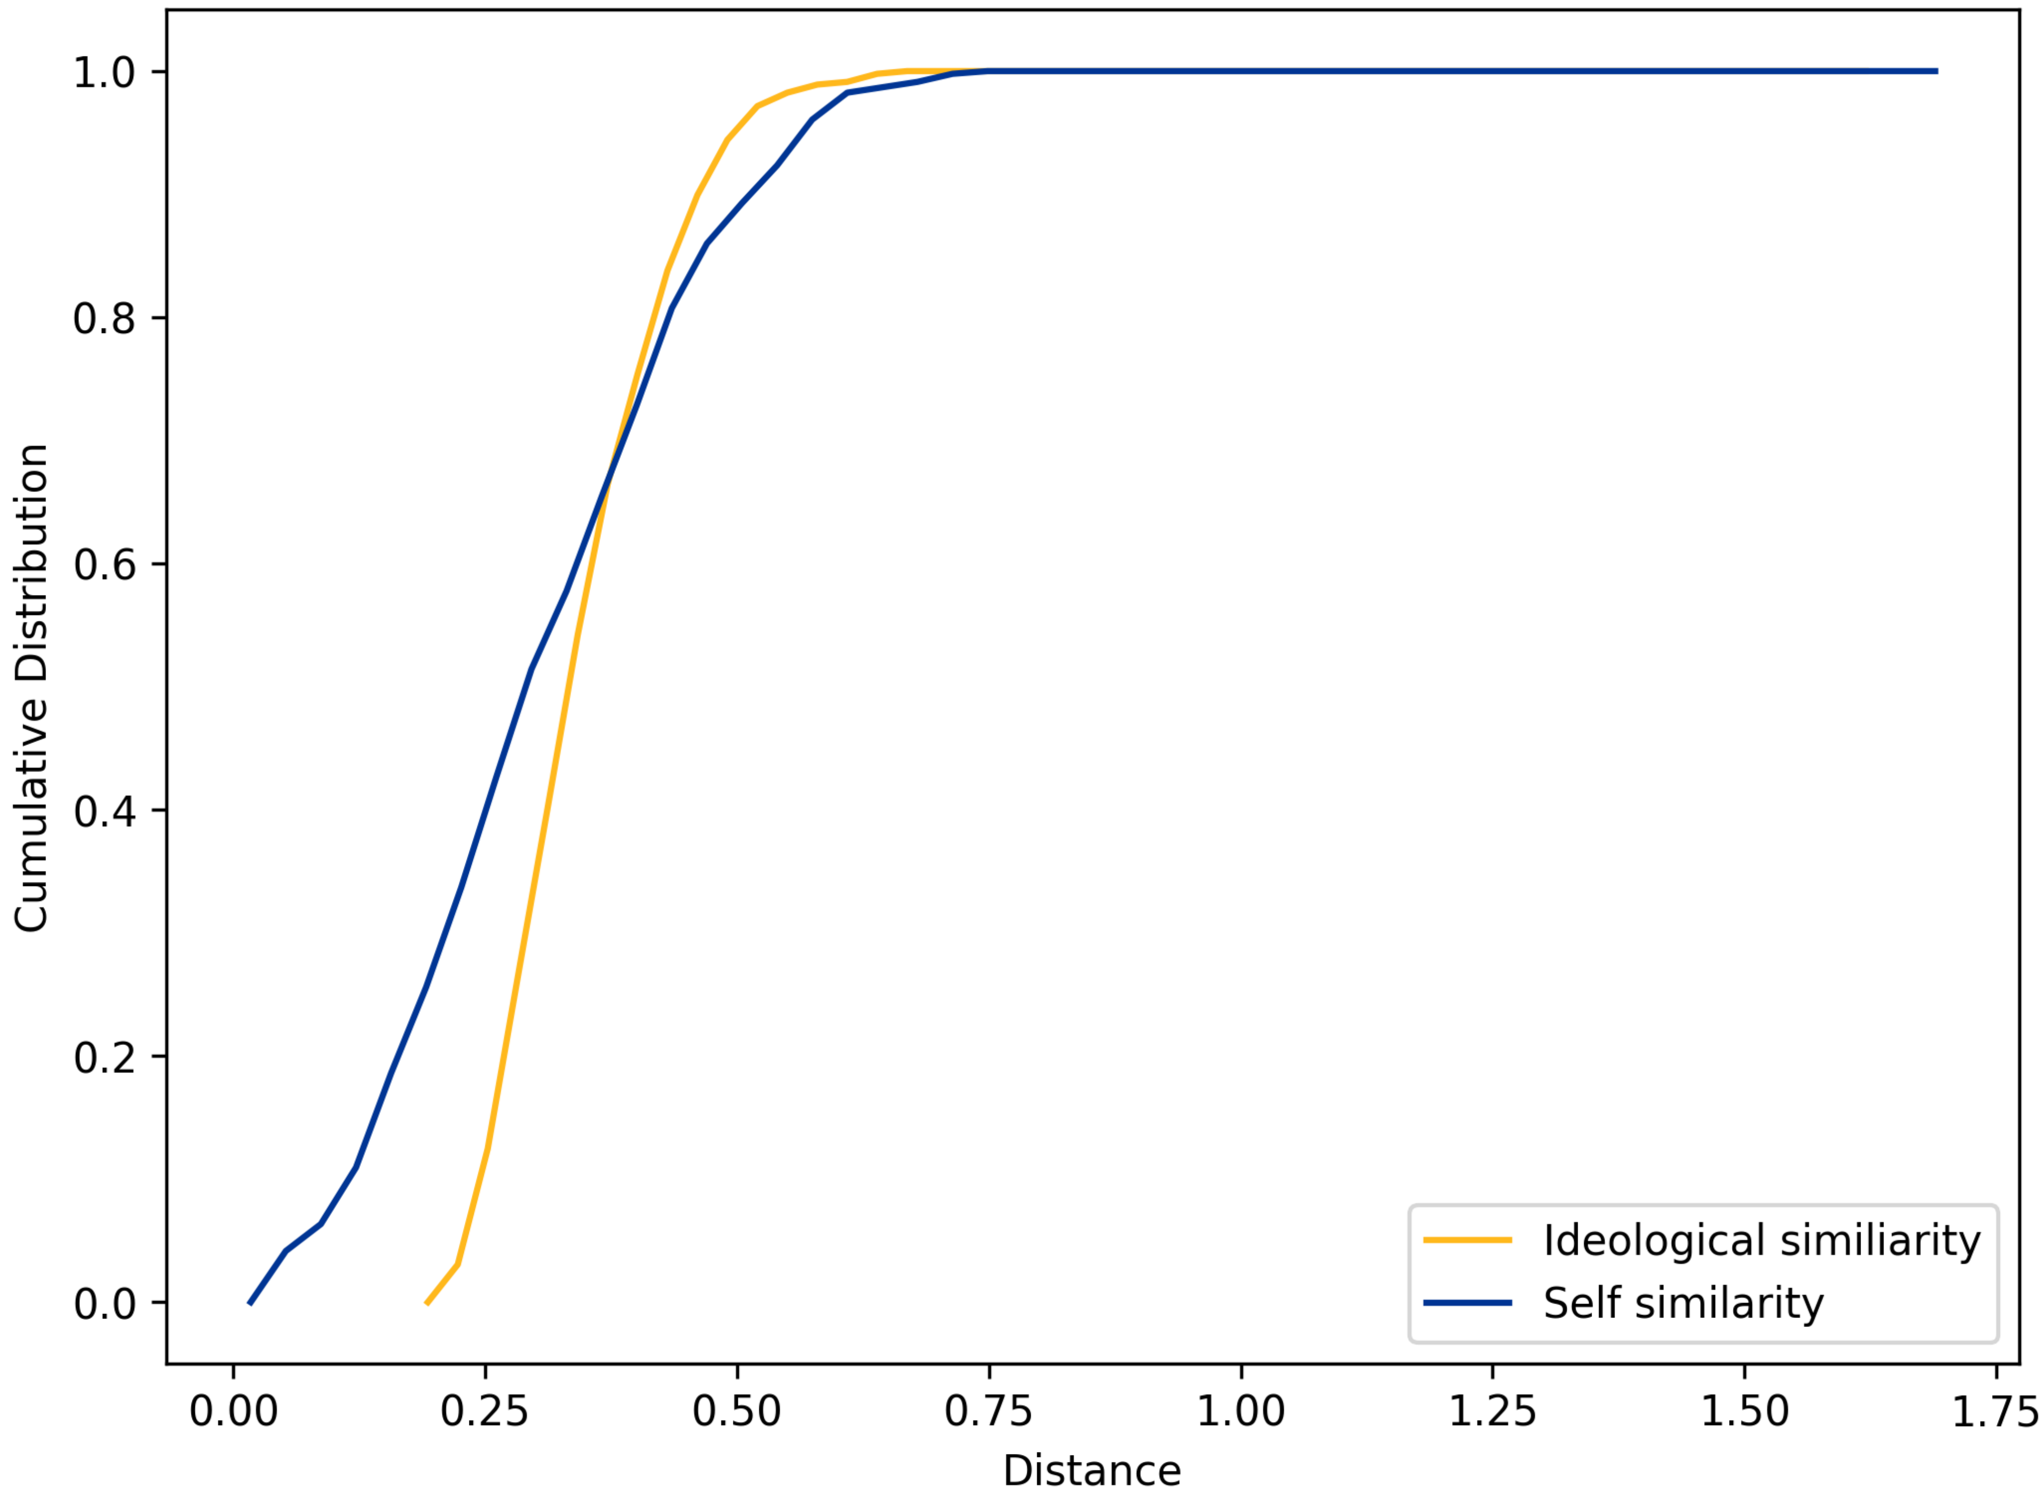

$p << 0.05$

R2: Sources of Similarity

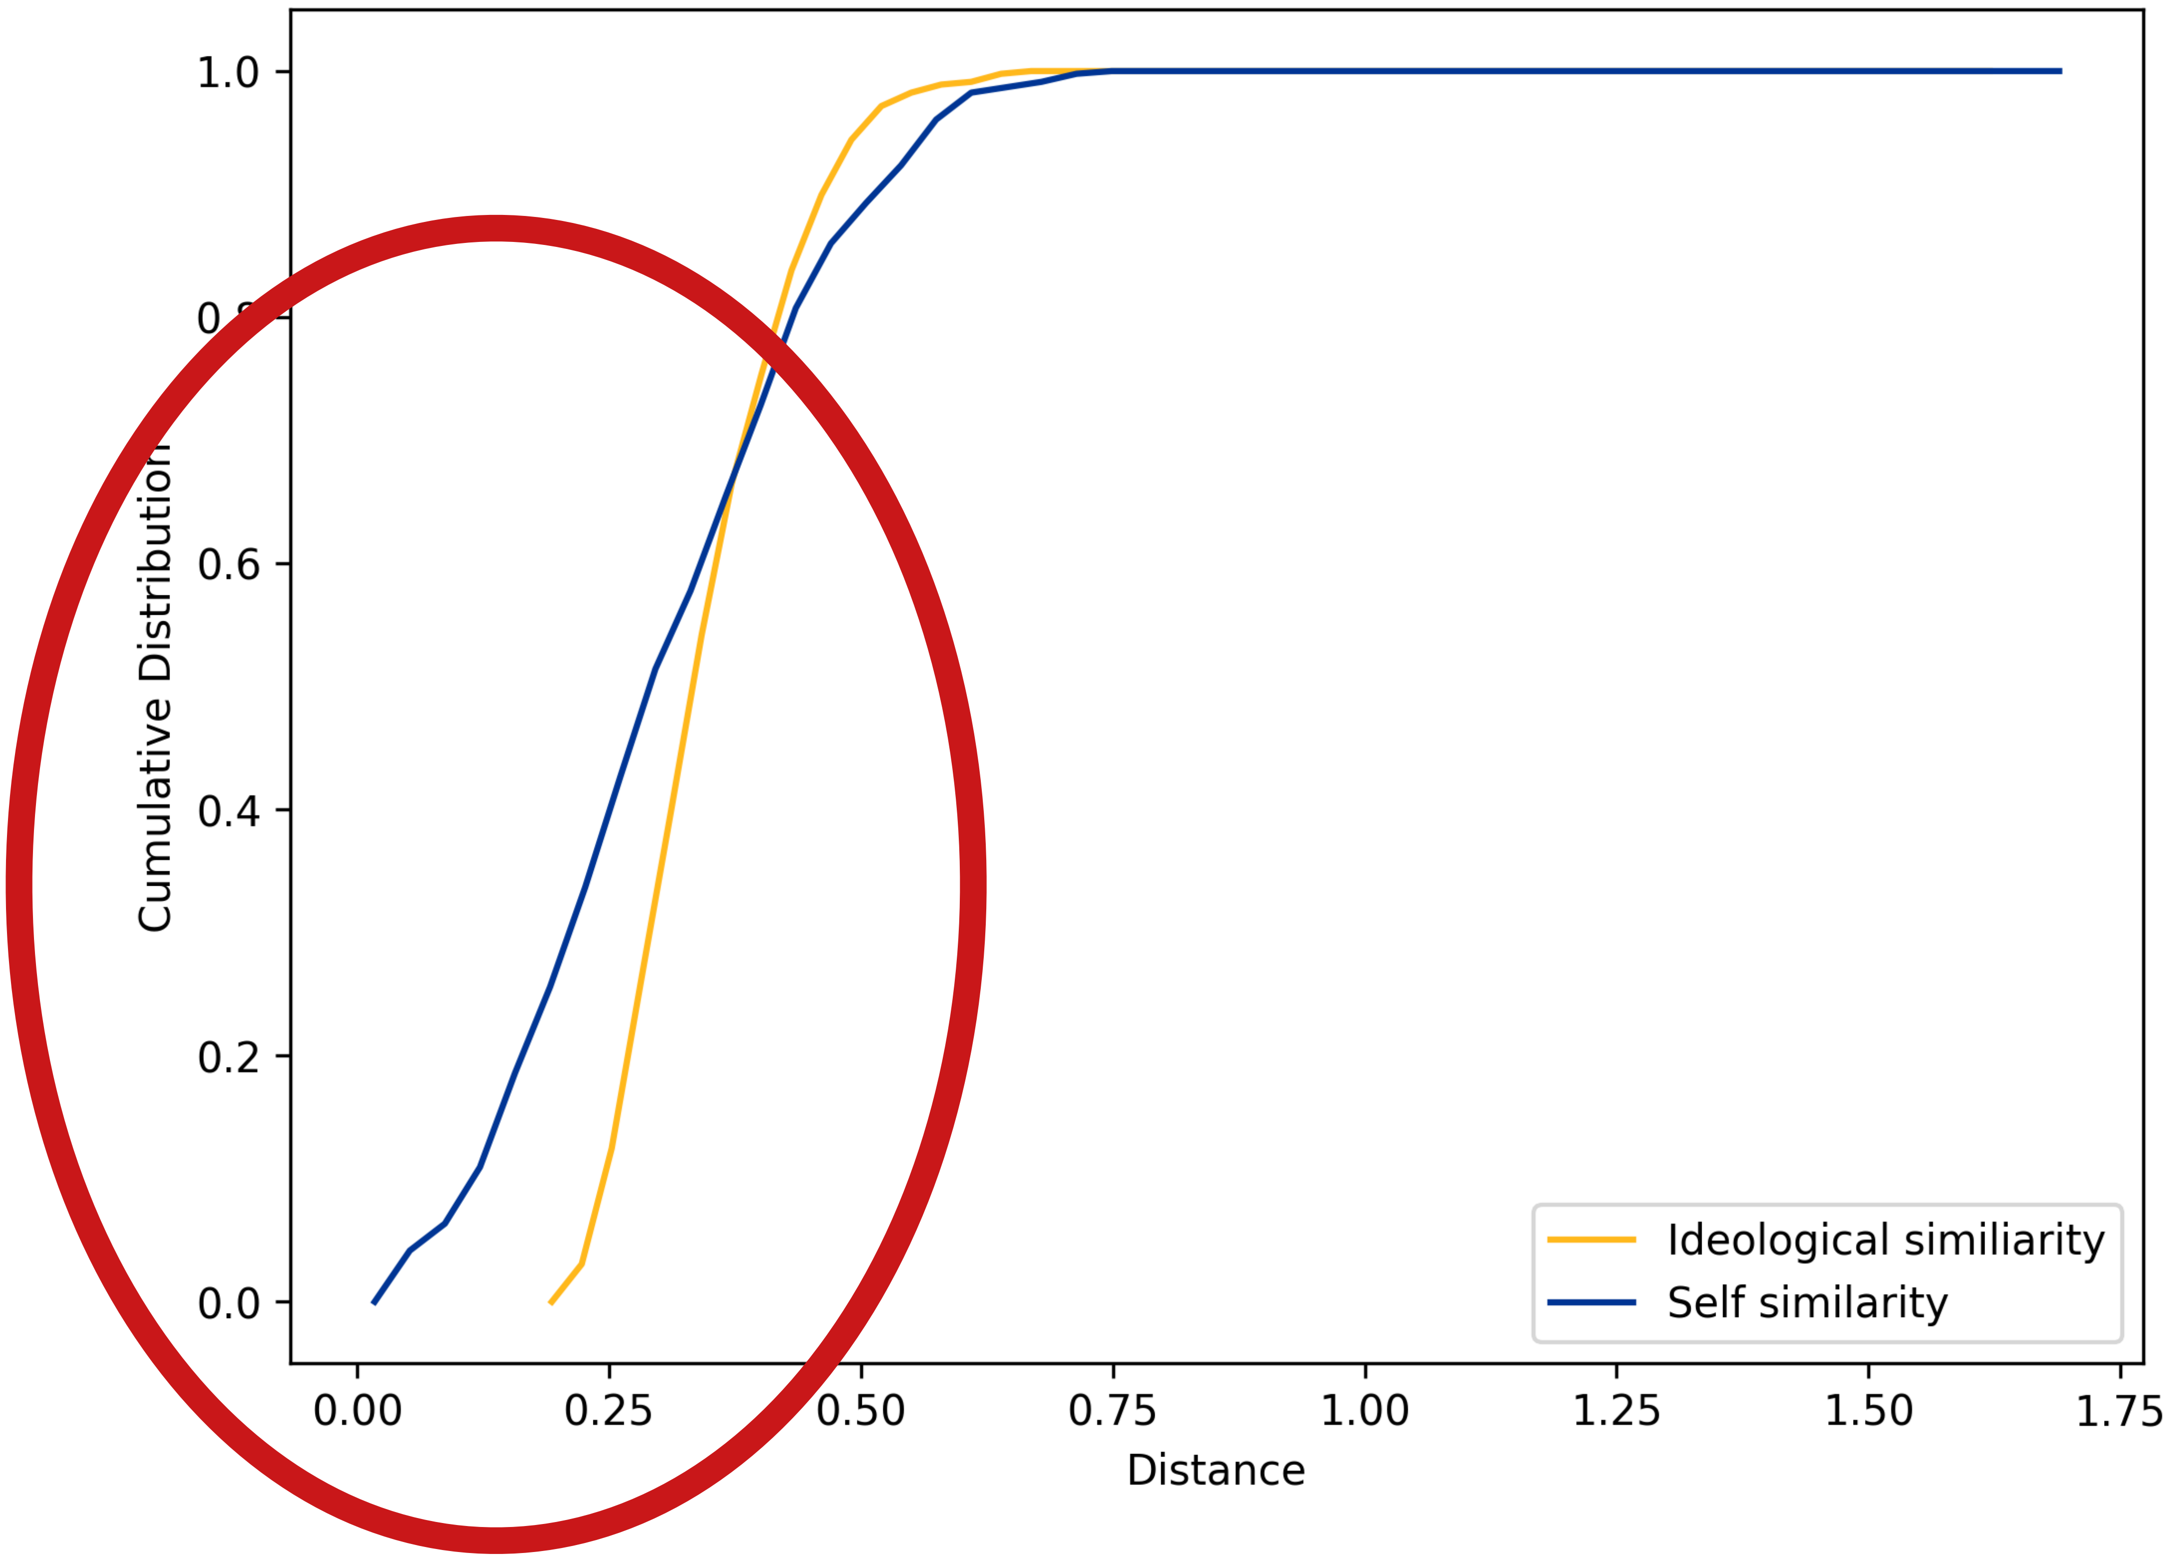

p << 0.05

Data

2. Ideological "Turing test"
   - 1000 subjects, recruited by YouGov
   - Asked to provide "liberal" and "conservative" positions on one of three topics (1) abortion (2) minimum wage (3) national defense

Research Questions

- Is structure driven by ideology or by individual traits? **Individual traits.**
- Does structure suggest argument quality?

Hopkins and Noel, 2016

Data – Guess That Ideology!

The conservative / liberal position on abortion is:

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The liberal position on abortion is:

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The liberal position on abortion is:

A woman has the right to determine what happens to her body

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The liberal position on abortion is:

A woman has the right to determine what happens to her body

Coding = 1
Authentic

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The liberal position on abortion is:

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The liberal position on abortion is:

It is okay to murder

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The liberal position on abortion is:

It is okay to murder

Coding = 0
Ironic

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The conservative position on abortion is:

This text was written by a:

conservative

liberal

Data — Guess That Ideology!

The conservative position on abortion is:

Women need guidance from more superior men!

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The conservative position on abortion is:

Women need guidance from more superior men!

Coding = 0
Ironic

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The conservative position on abortion is:

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The conservative position on abortion is:

All life is sacred.

This text was written by a:

conservative

liberal

Data – Guess That Ideology!

The conservative position on abortion is:

All life is sacred.

Coding = 1
Authentic

This text was written by a:

conservative

liberal

3. Potential for Behavioral Insights

Does structure suggest argument quality?

- ➡ Can we tell “authentic” from “ironic” responses?

3. Potential for Behavioral Insights

Does structure suggest argument quality?

➡ Can we tell “authentic” from “ironic” responses?

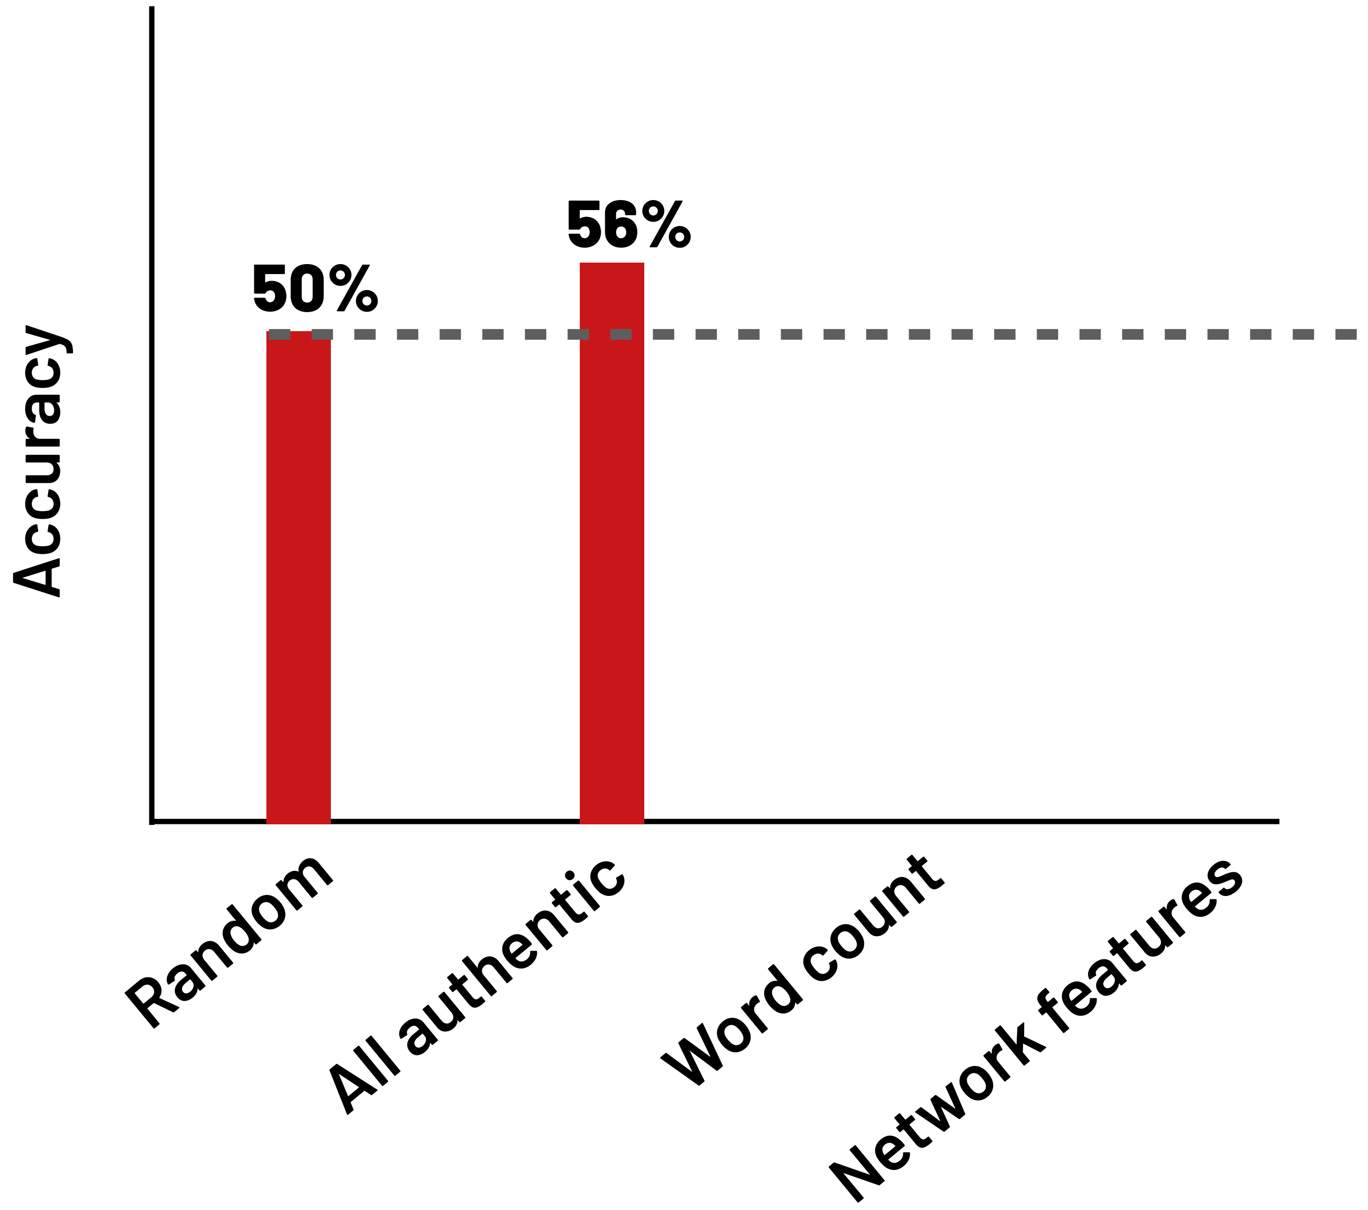

3. Potential for Behavioral Insights

Does structure suggest argument quality?

➡ Can we tell “authentic” from “ironic” responses?

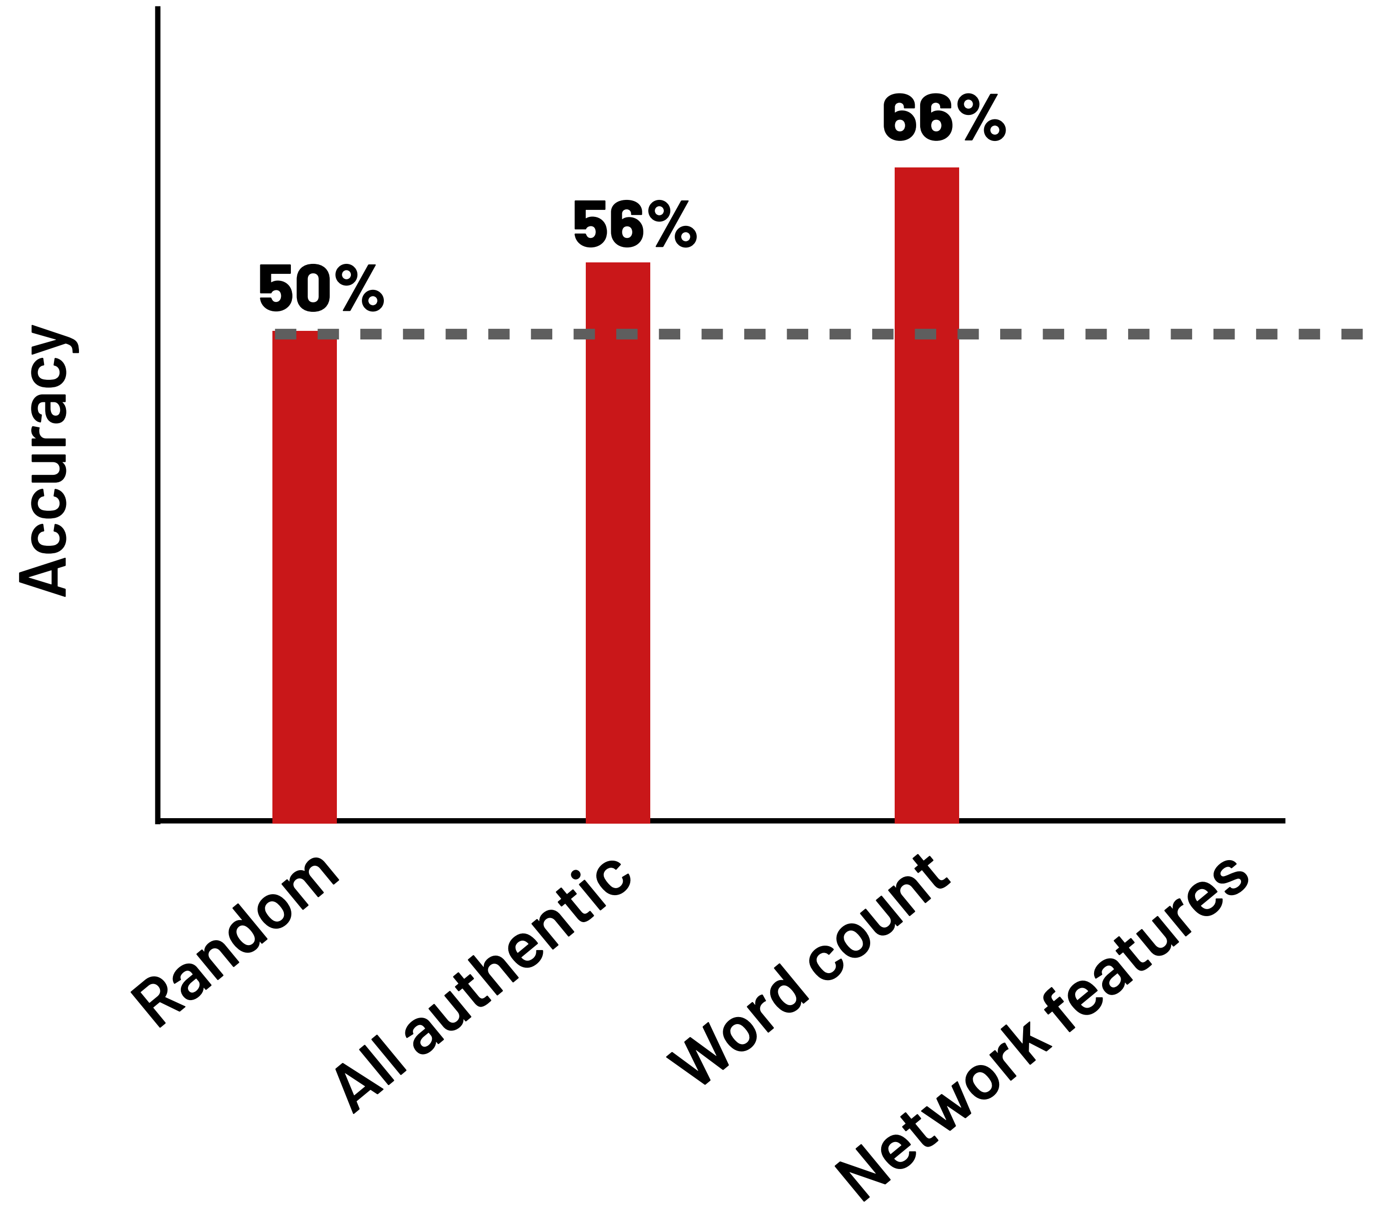

3. Potential for Behavioral Insights

Does structure suggest argument quality?

➡ Can we tell "authentic" from "ironic" responses?

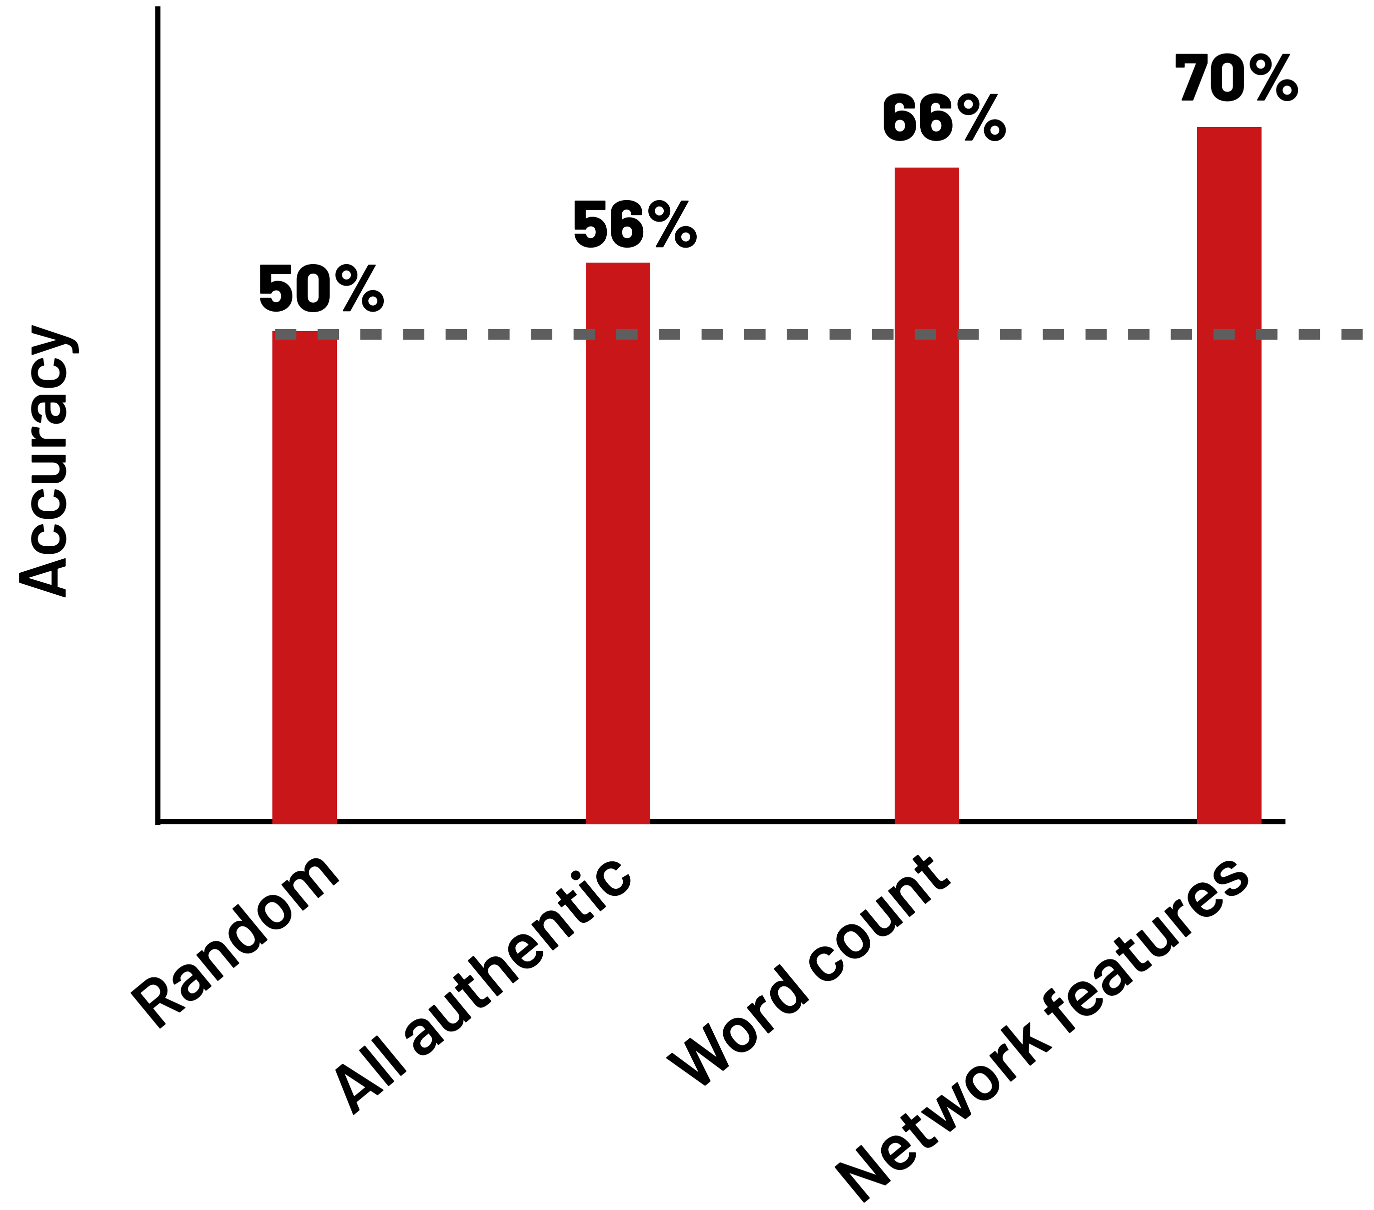

3. Potential for Behavioral Insights

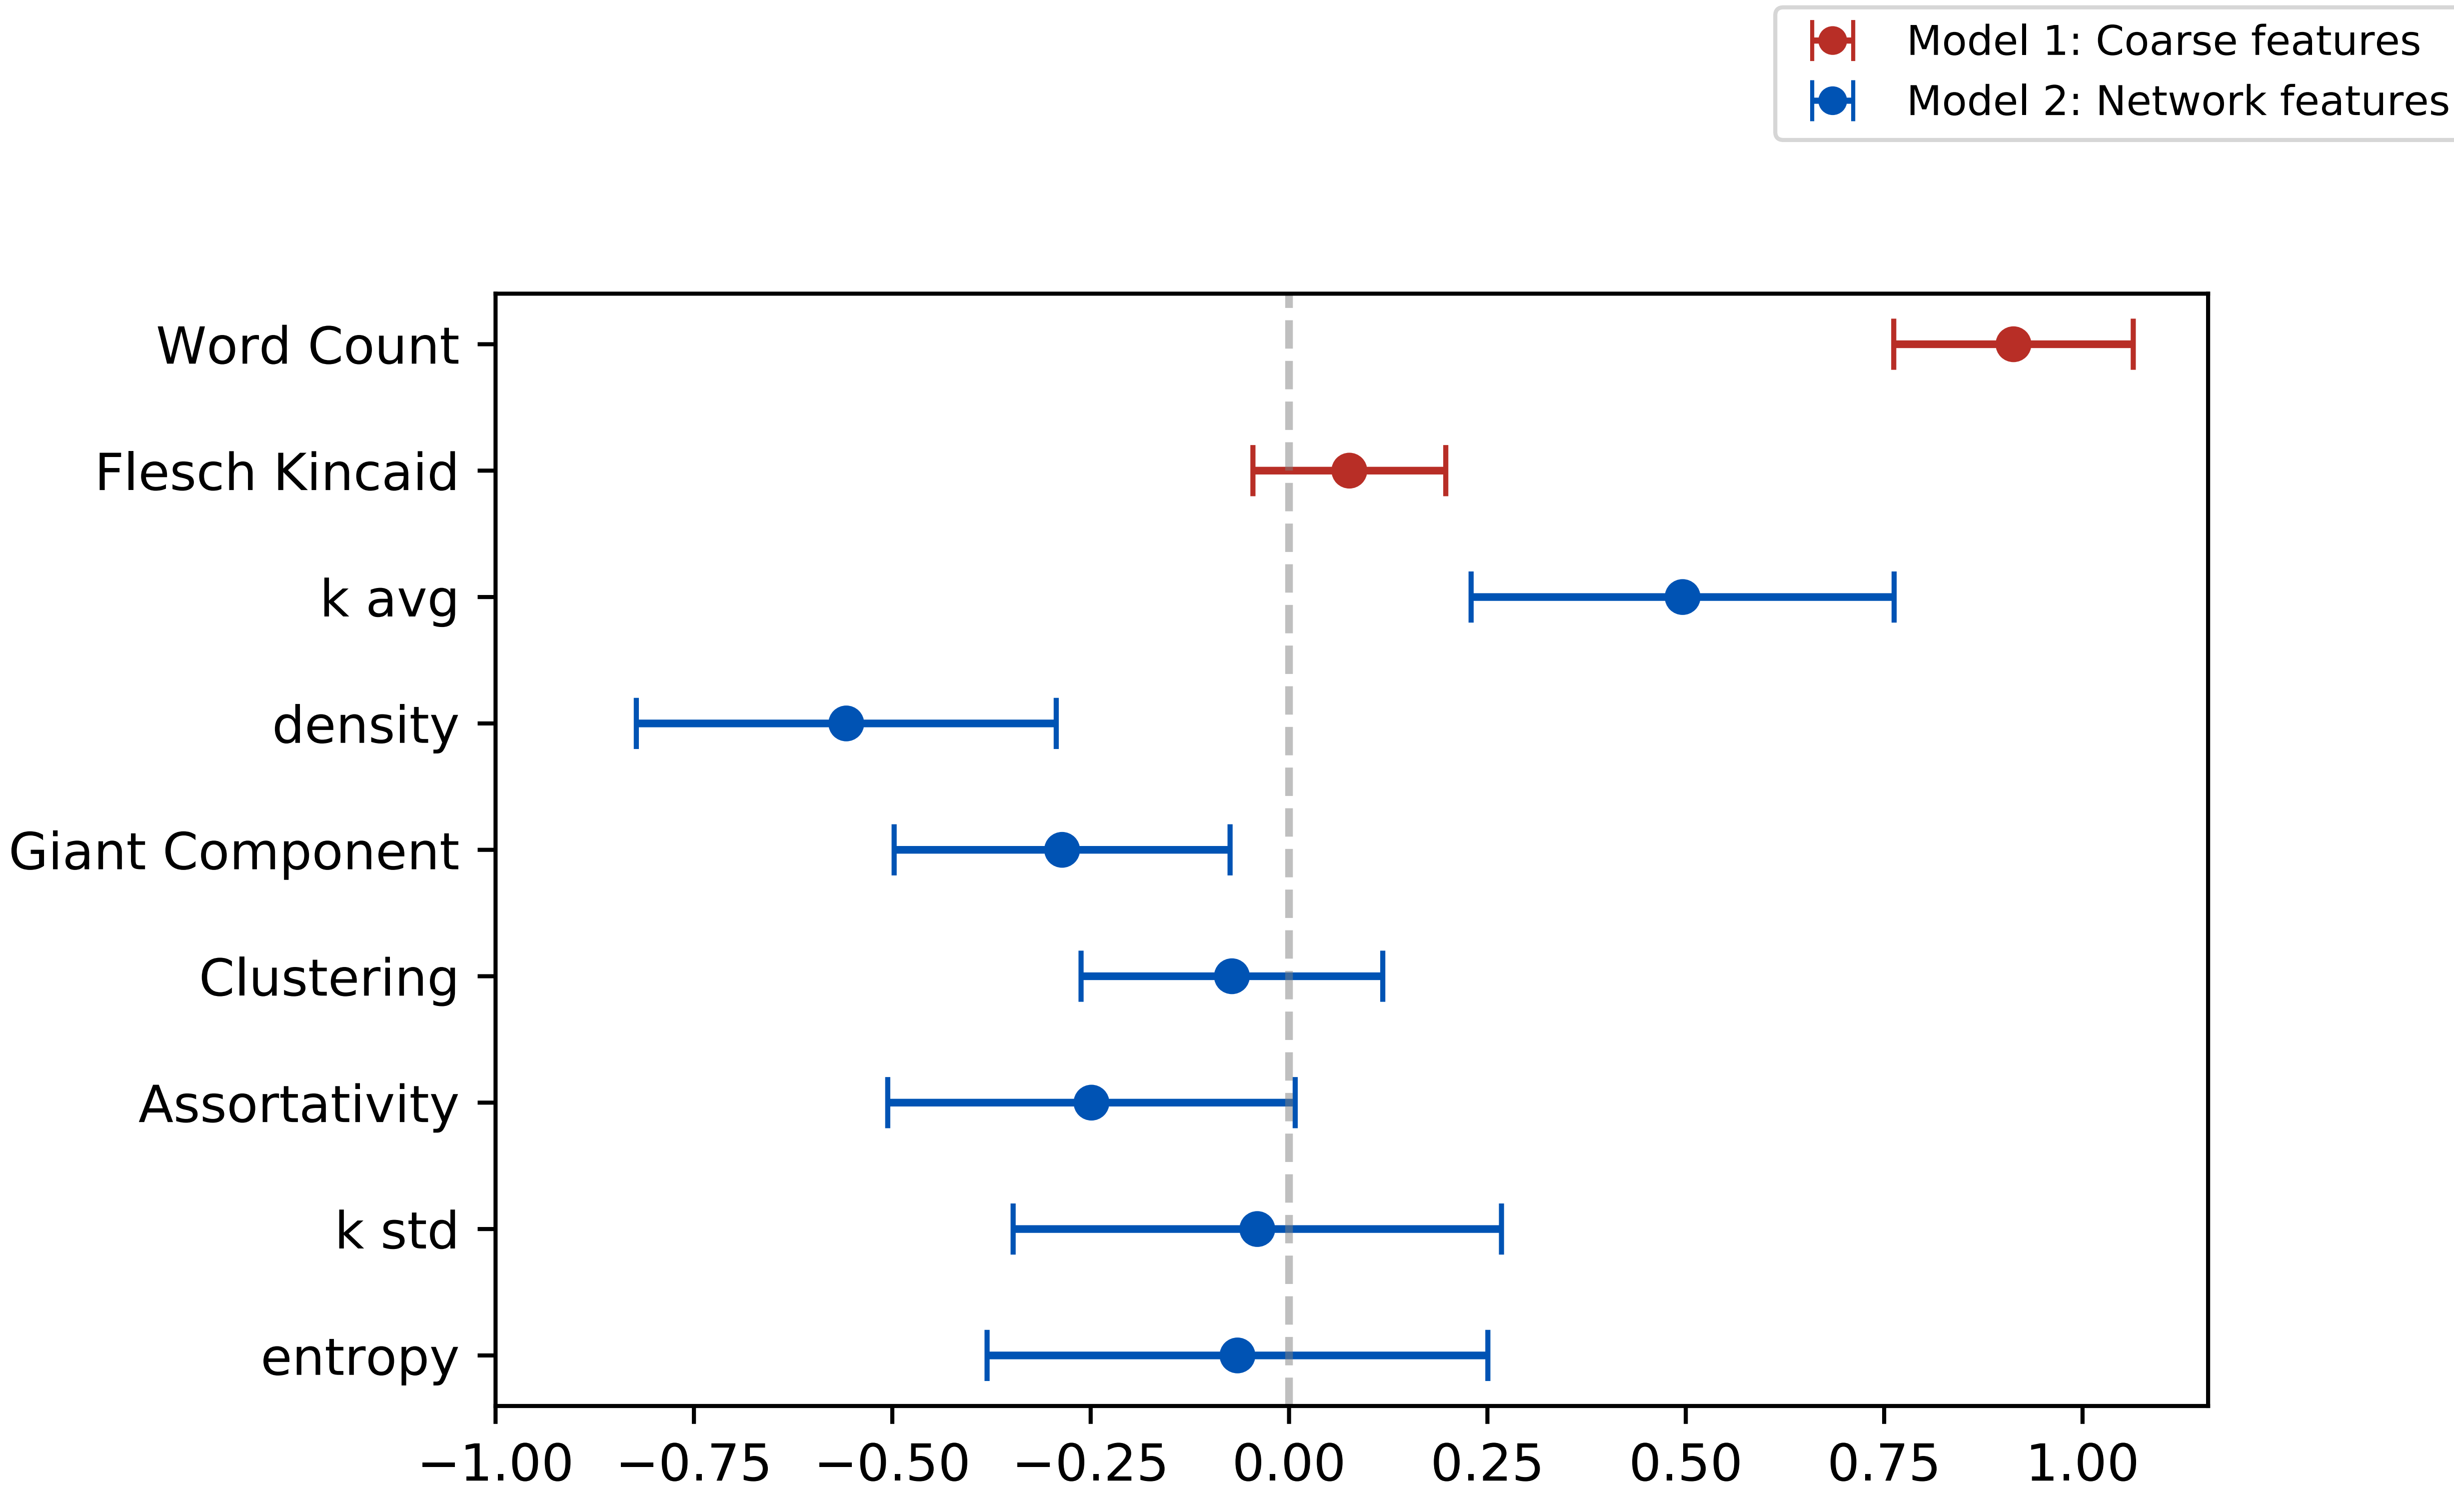

3. Potential for Behavioral Insights

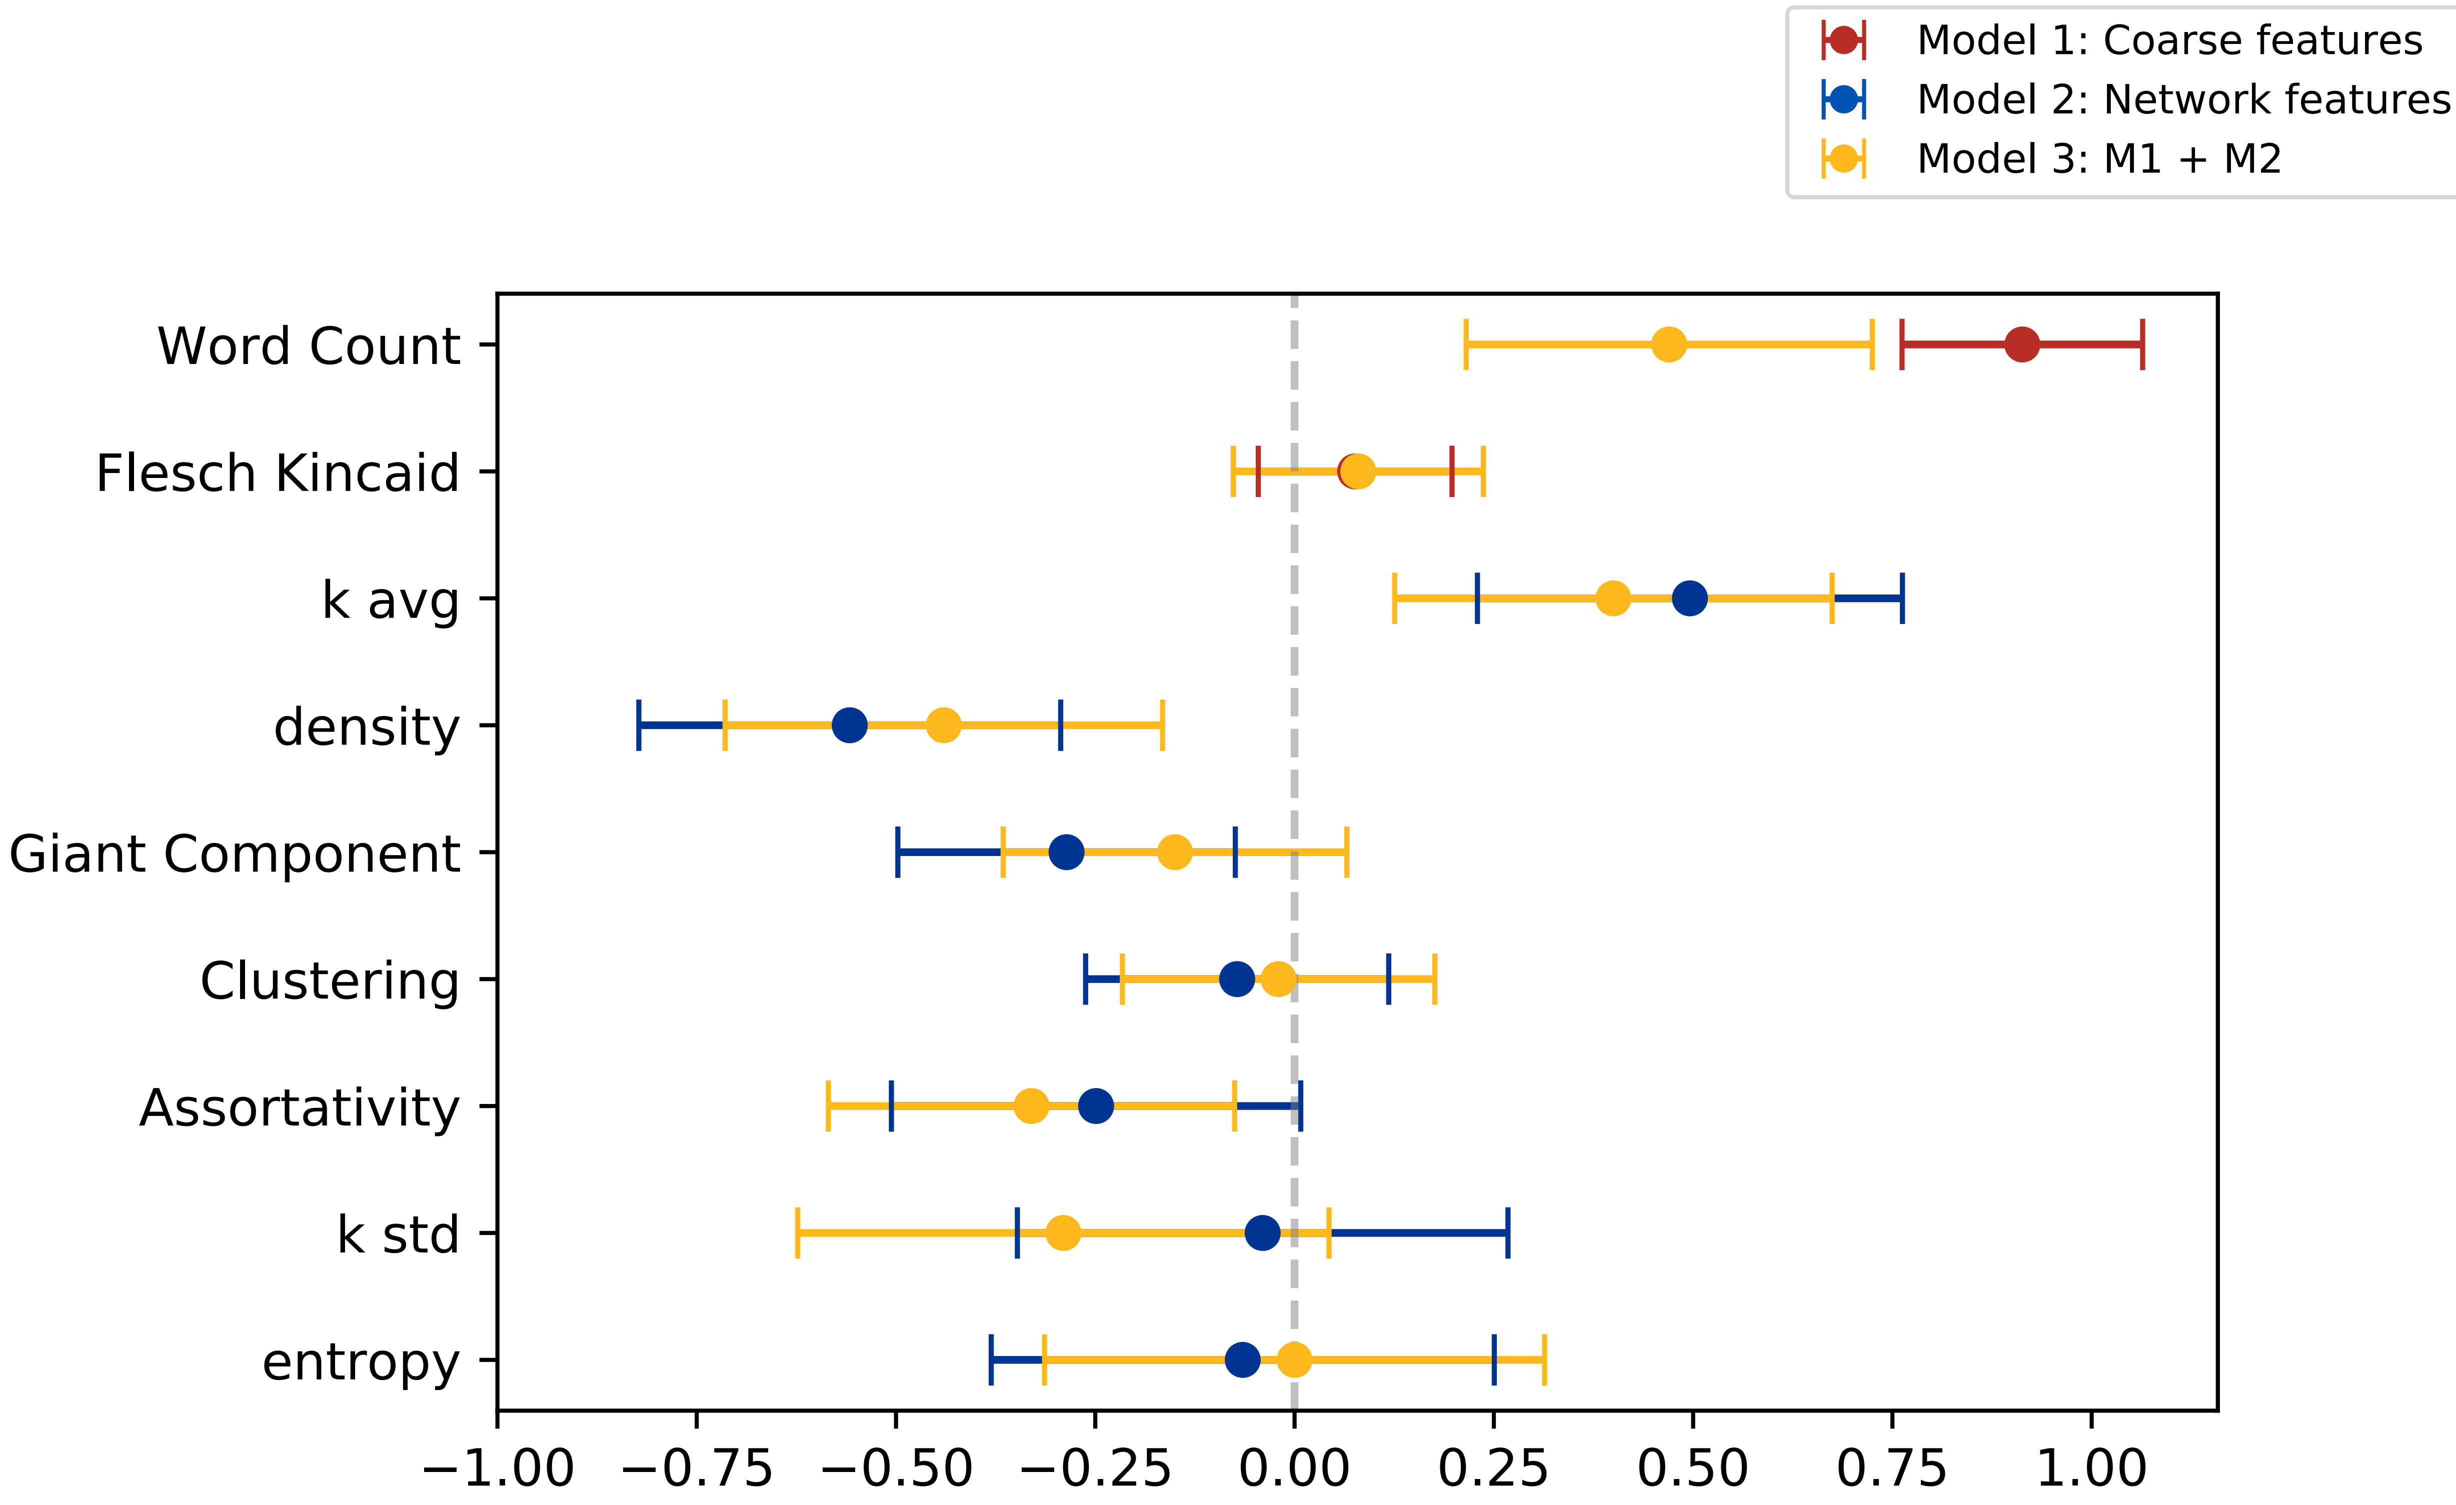

3. Potential for Behavioral Insights

2. Ideological "Turing test"

   - 1000 subjects, recruited by YouGov
   - Asked to provide "liberal" and "conservative" positions on one of three topics
   (1) abortion (2) minimum wage (3) national defense

Research Questions

- Is structure driven by ideology or by individual traits? **Individual traits.**
- Does structure suggest argument quality? **Yes.**

Hopkins and Noel, 2016

Summary

- New method for inferring structure of expressed reasons
- Reveals small but meaningful individual variation
- Correlated with known personality traits
- Potential for new insights into dynamics of public opinion

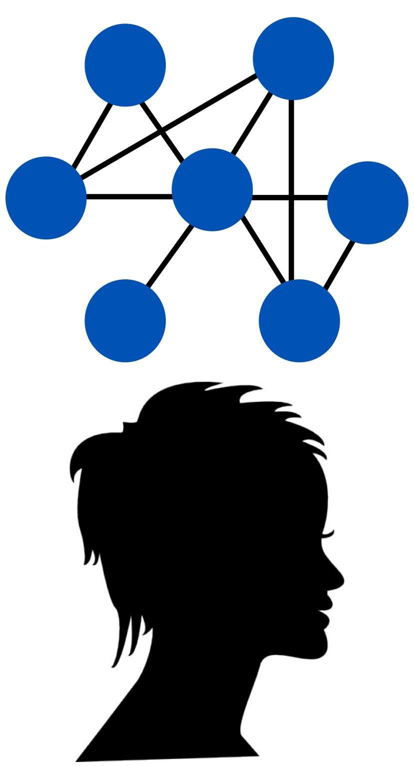

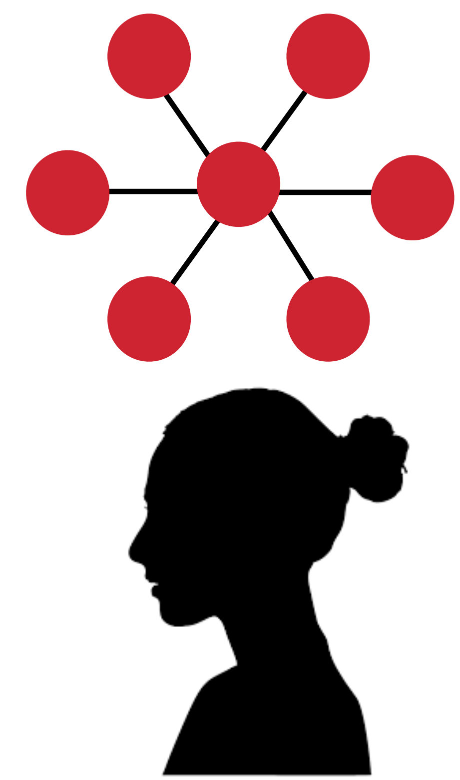

Final Thoughts

- NLP methods can be used to address a range of questions
- The key is to figure out (1) how to operationalize your question and (2) what features are of interest
- Remember: language is high dimensional

Sarah Shugars

Northeastern University

shugars.s@northeastern.edu

@Shugars she/her

Appendix

Emotional Measures

Data: Number of Users

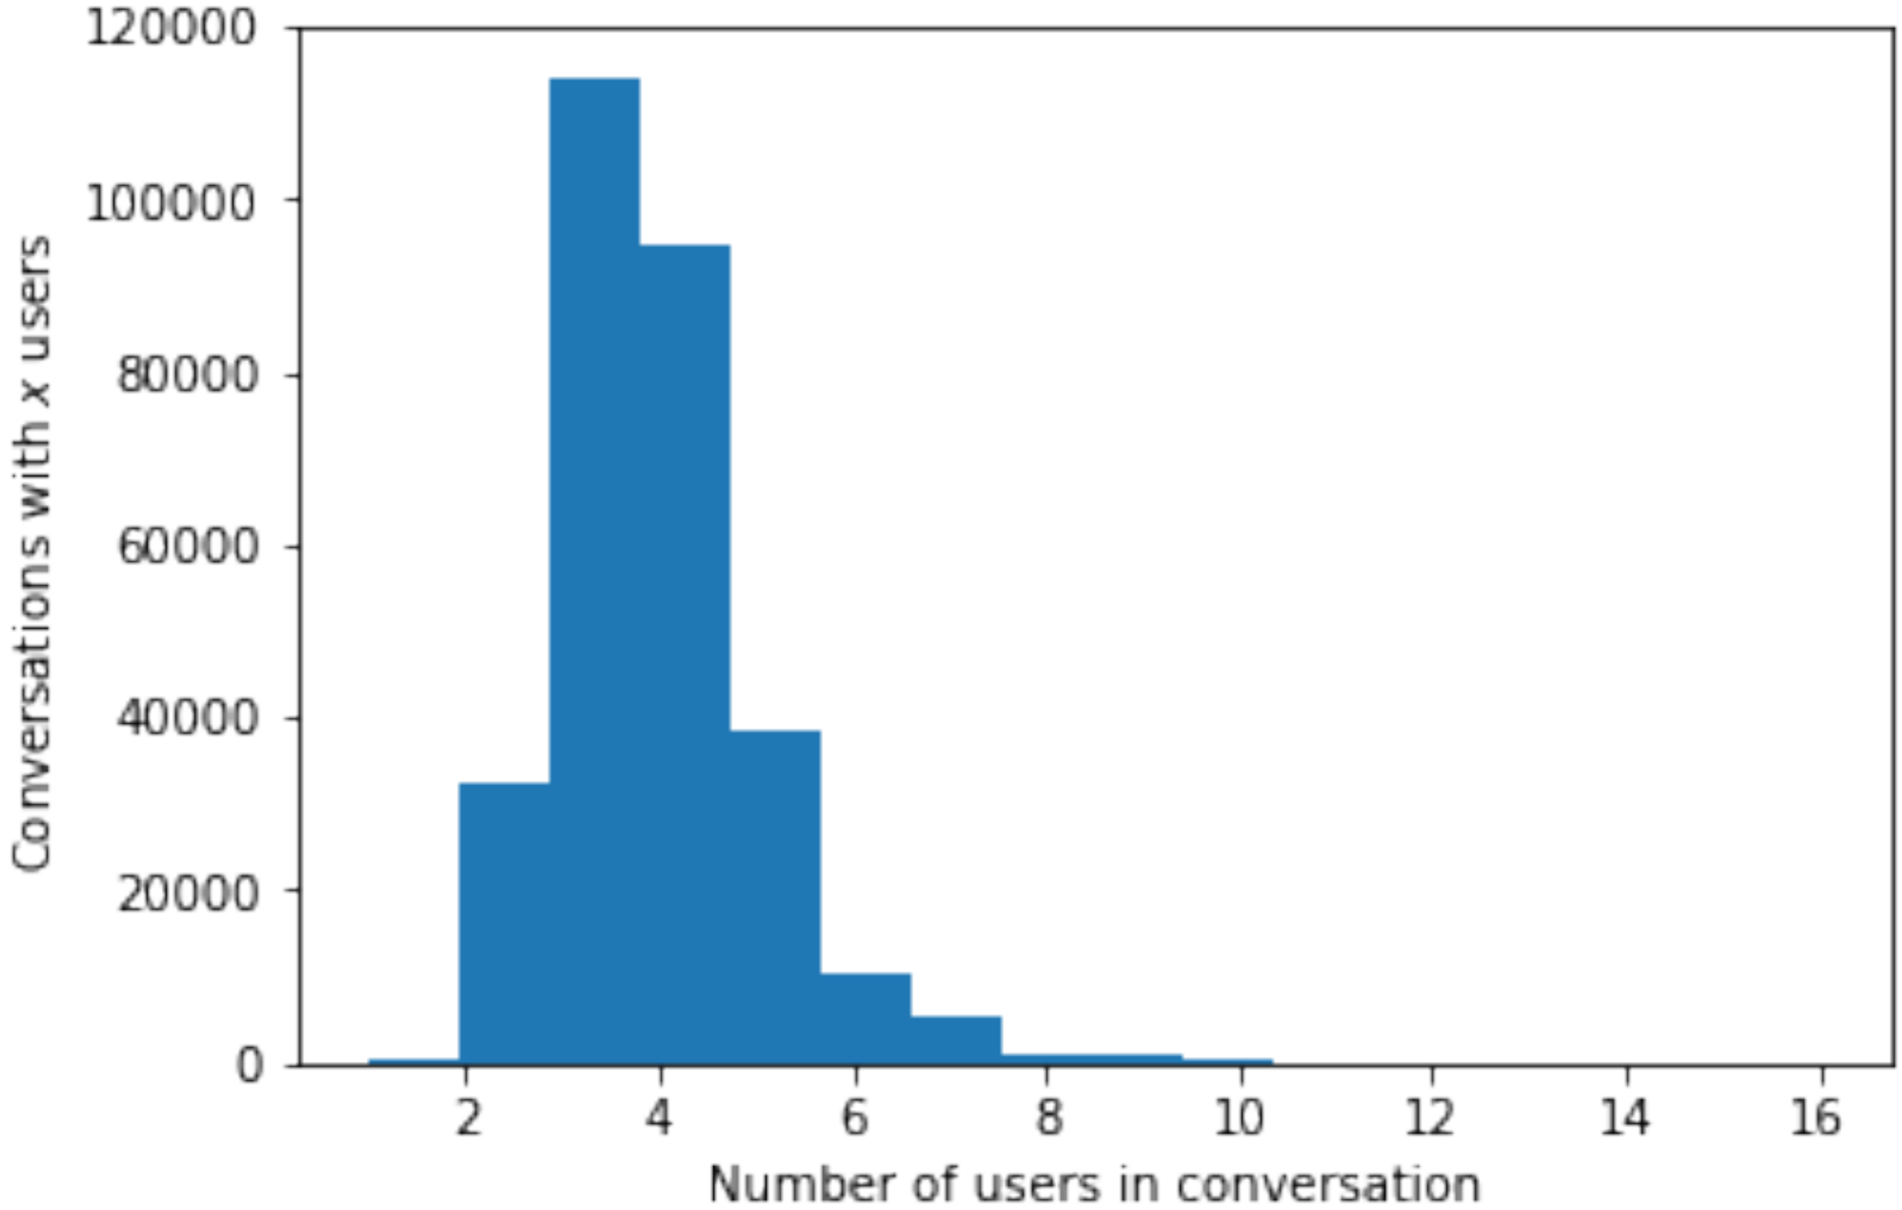

Data: Conversation Length

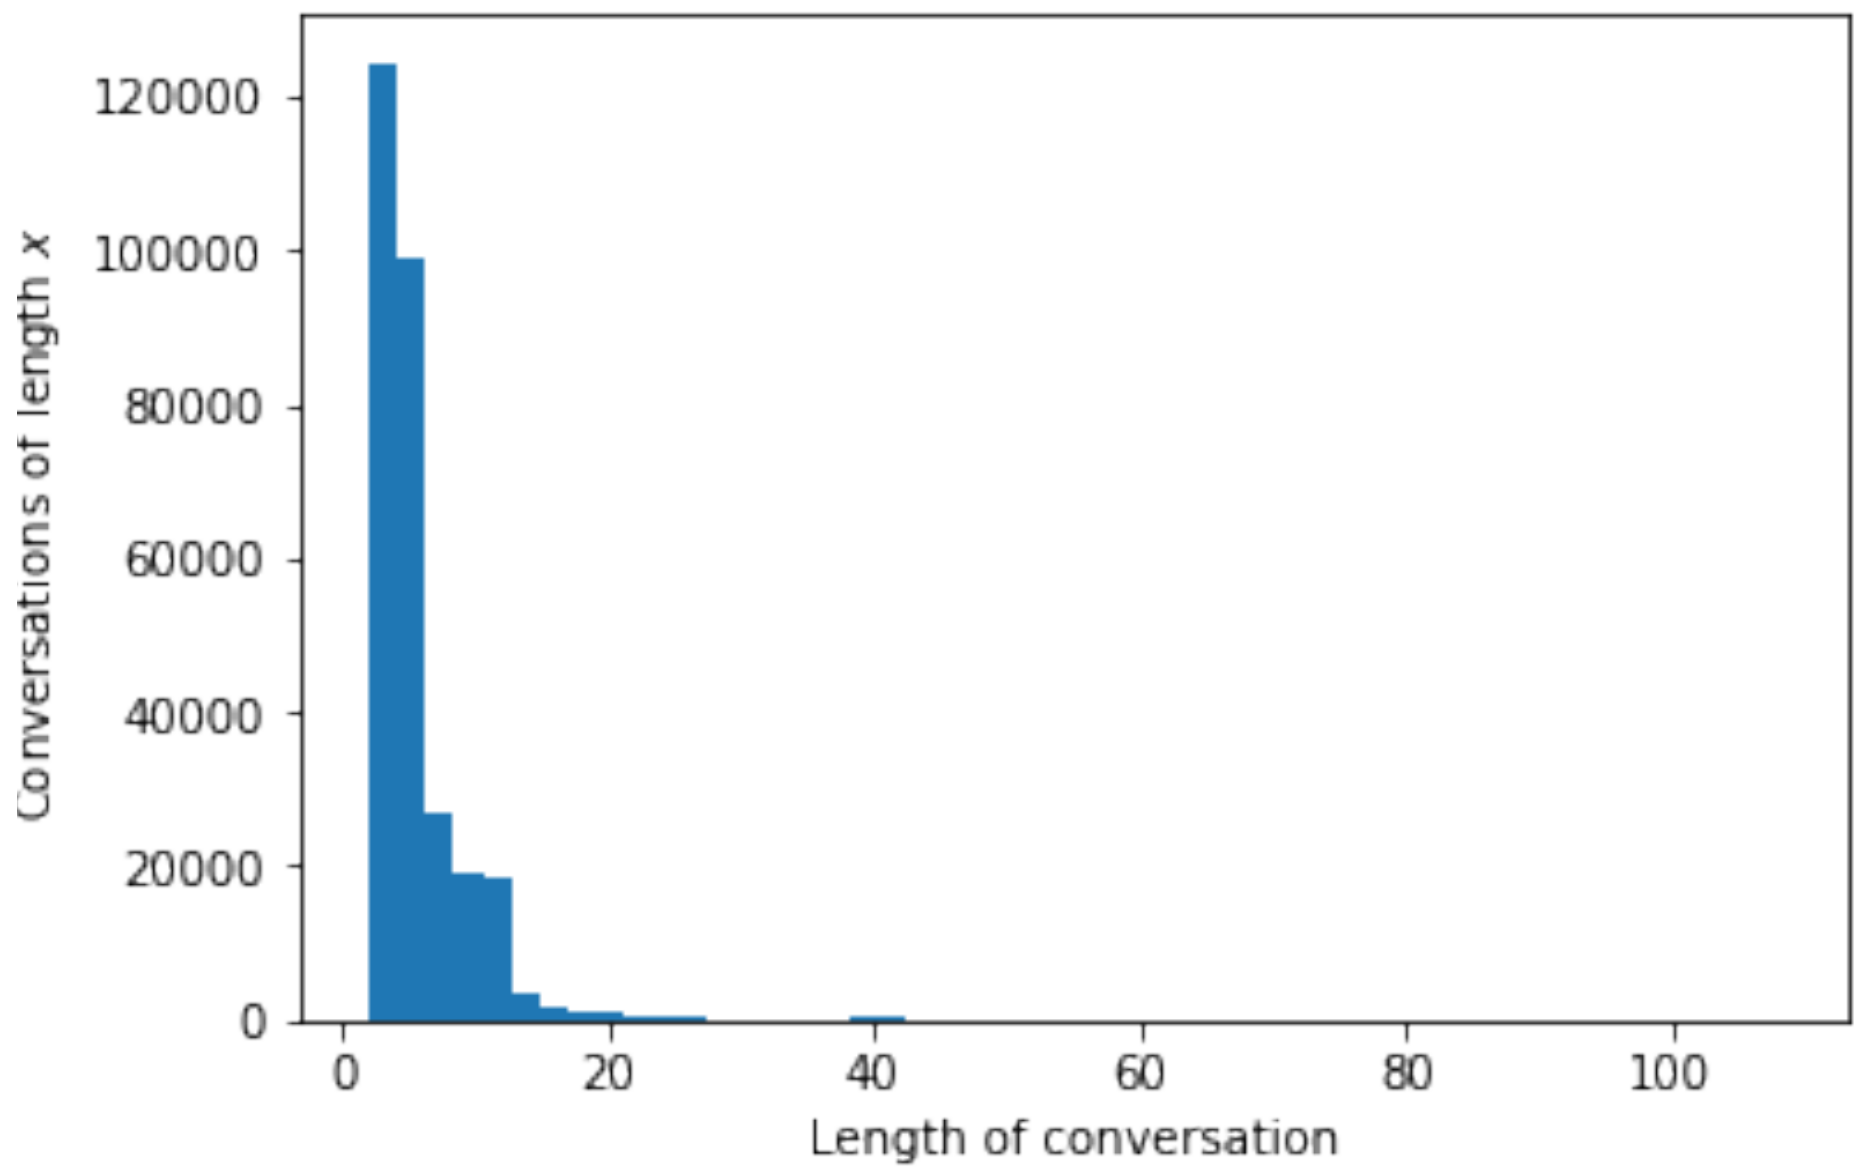

Model 1: Findings

MCMC: 6 chains of 50k

Parameters: $2\,\beta + 4000\,\gamma + 4000\,\theta$

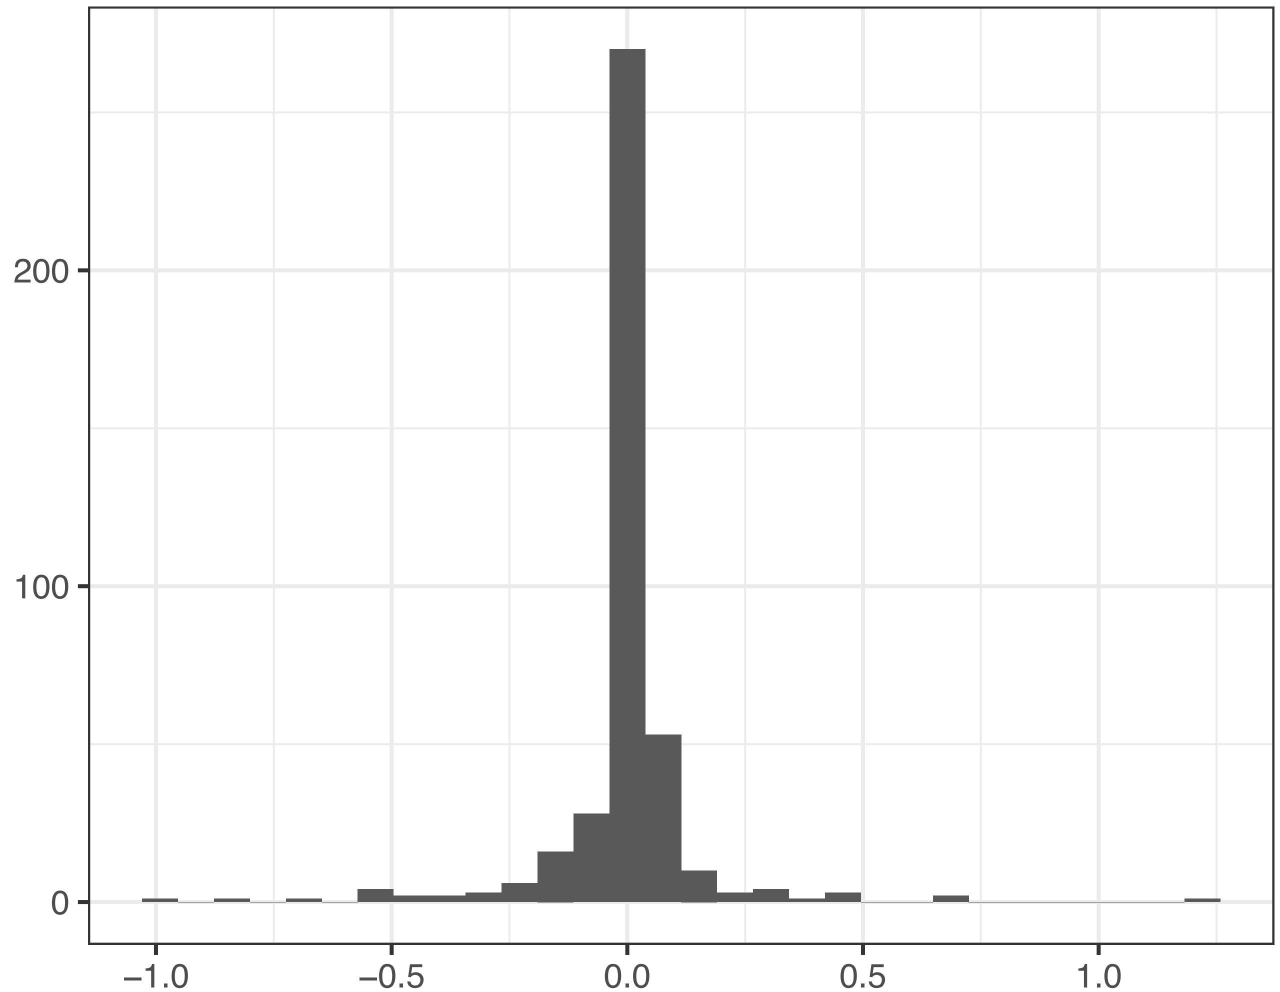

Inferred ideological positions

Confusion Matrix

		Truth	
		0	1
Prediction	0	0.898	0.012
	1	0.008	0.082

Topical correlation

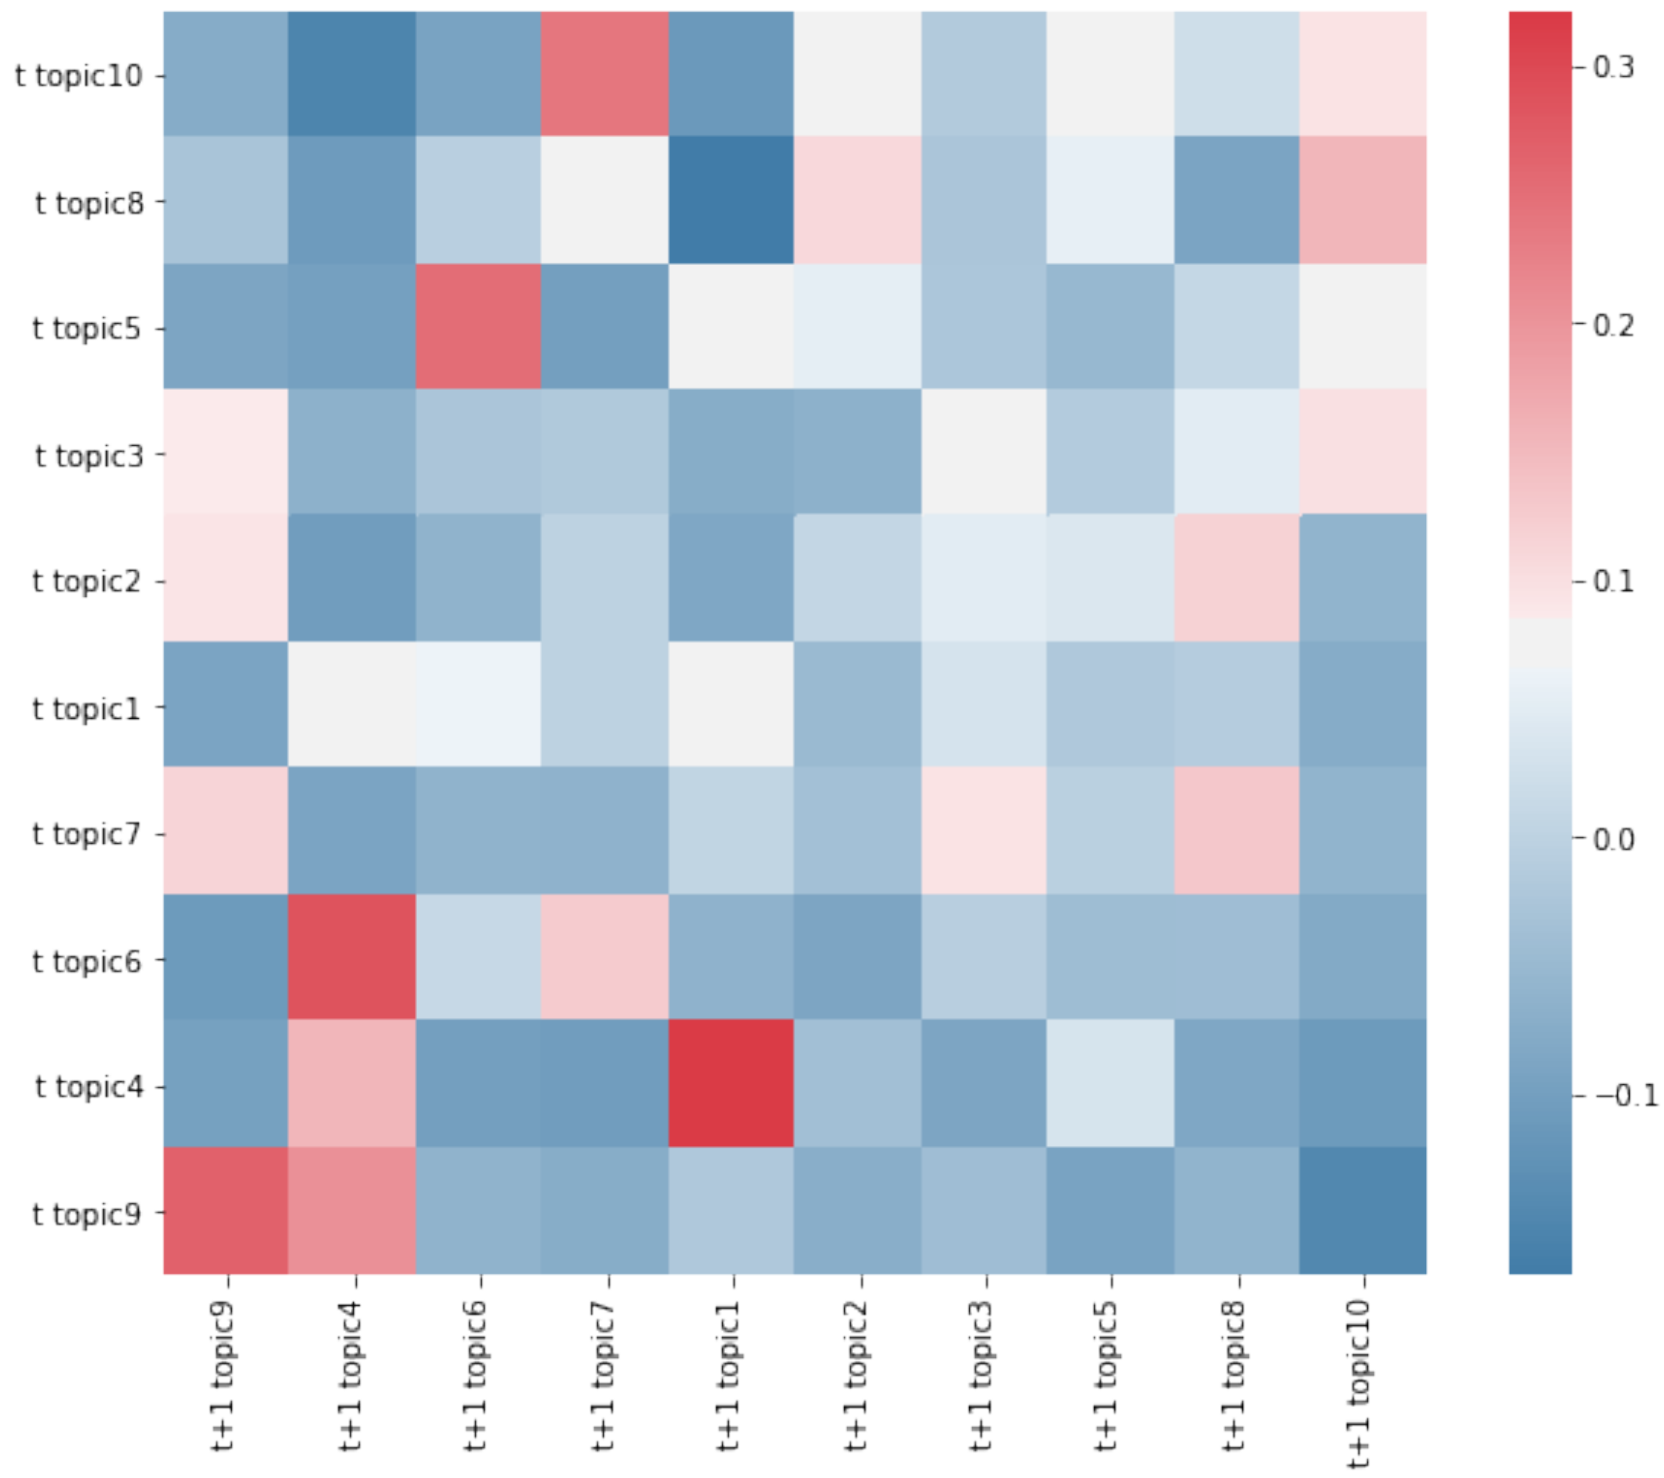

Emotional correlation

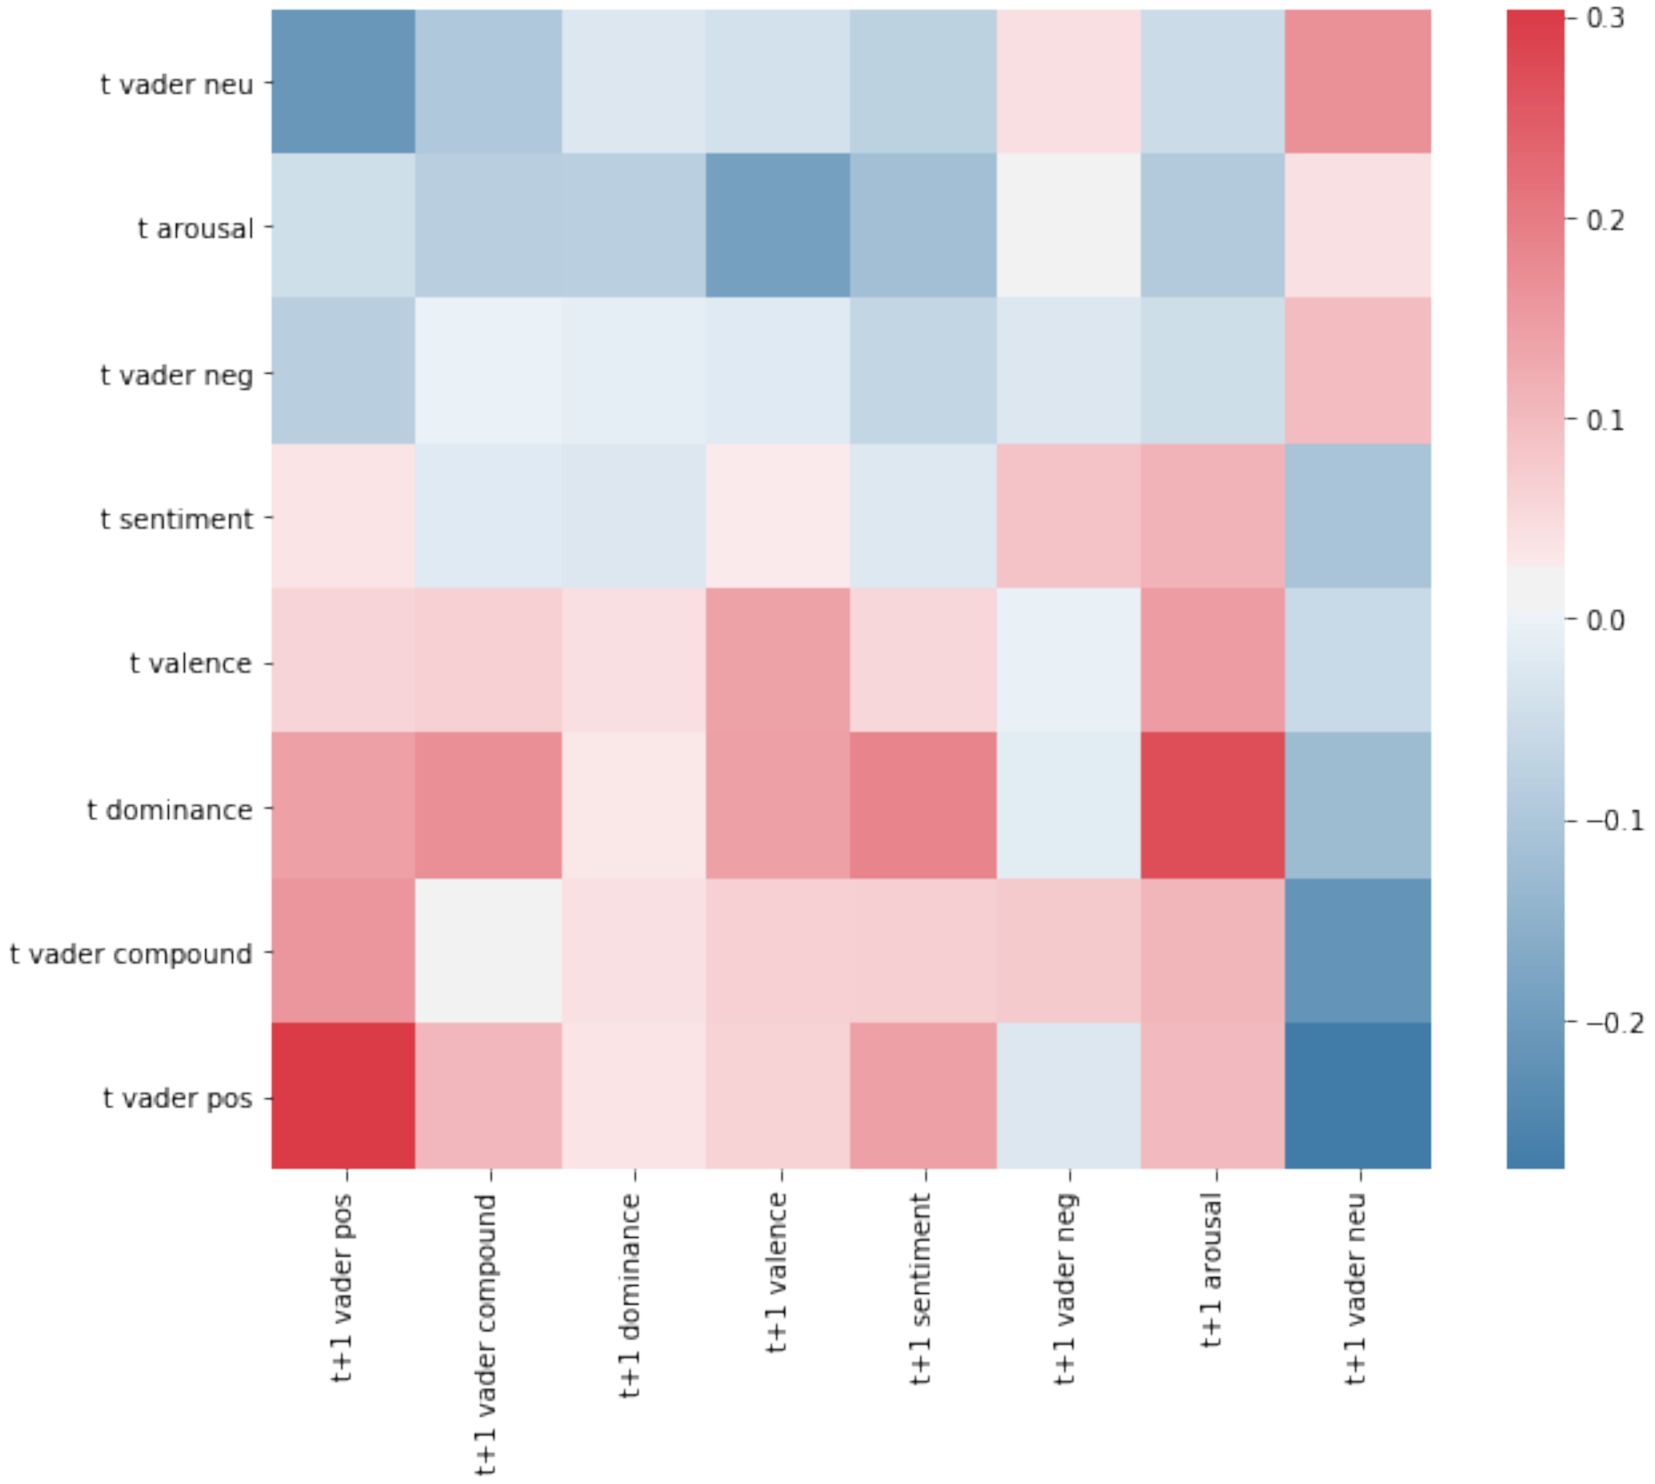

Response predictors: Candidate's prev. tweet

	Coef	Significance after p correction FDR	Clust	FDR+Cl
prev response	0.883	***	***	***
favorite count	−0.311	***		
retweet count	−262.234	*		
reply count	0.141	***		
quality	262.523	*		
source	0.037	***		
xday	0.169	***	**	
yday	0.239	***	*	
xhour	0.048	***		
yhour	0.193	***	*	
chars	0.367	***	***	**
has url	0.037	***		
mentions	0.155	***		
hashtags	−0.078	***	**	*
sentiment	0.362	***	*	
vader neg	0.641	***	***	**
vader pos	−0.313	***	**	*
valence	−0.084	***		
arousal	0.151	***		
dominance	−0.174	***		
time since prev	−0.658	***	***	**
topic 2	1.853	***	**	
topic 3	−0.037			
topic 4	−0.364	***		
topic 5	0.246	***		
topic 6	−0.536	***		
topic 7	−1.153	***		
topic 8	−2.787	***	***	**
topic 9	−0.573	***		
topic 10	2.404	***	**	*

Note: *p<0.1; **p<0.05; ***p<0.01

Response predictors: Current tweet

	Coef	Significance after p correction		
		FDR	Clust	FDR+Cl
favorite count	1.657	***		
retweet count	9.171	***		
reply count	−10.055		***	***
quality	−43.260			
source	−0.348	***	***	***
xday	−0.354	***	***	**
yday	−0.345	***	***	***
xhour	0.146	***		
yhour	−0.044	***		
chars	0.649	***	***	***
has url	0.075	***		
mentions	−0.412	***	**	
hashtags	−0.083	***	**	
sentiment	−0.135	***		
vader neg	−0.111	***		
vader pos	0.152	***		
valence	−0.524	***	**	*
arousal	−0.116	***		
dominance	0.343	***	*	
topic 2	0.815	***		
topic 3	2.140	***	**	*
topic 4	1.541	***		
topic 5	2.913	***	***	**
topic 6	1.024	***		
topic 7	2.669	***	**	*
topic 8	1.537	***		
topic 9	−0.148	***		
topic 10	2.043	***	*	
difference	−0.036	***		
difference^2	0.188	***	**	

Note: *p<0.1; **p<0.05; ***p<0.01

Response predictors: Potential respondent

		Significance after p correction		
	Coef	FDR	Clust	FDR+Cl
verified	−0.625	***	**	*
followers count	−54.294	***		
following count	−0.187	***		
statuses count	0.026	***		
favourites count	0.005			
comments count	−0.170	***		
Note:		*p<0.1; **p<0.05; ***p<0.01		

Response predictors: Conversation features

		Significance after p correction		
	Coef	FDR	Clust	FDR+Cl
participants	−0.179	***		
thread length	0.105	***		
thread length^2	−0.026	***	***	**

Note: *p<0.1; **p<0.05; ***p<0.01