

NLP meets Computational Social (Media) Science

Dirk Hovy

Computer Science Department
University of Copenhagen

`mail@dirkhovy.com`

 @dirk_hovy

Social Media

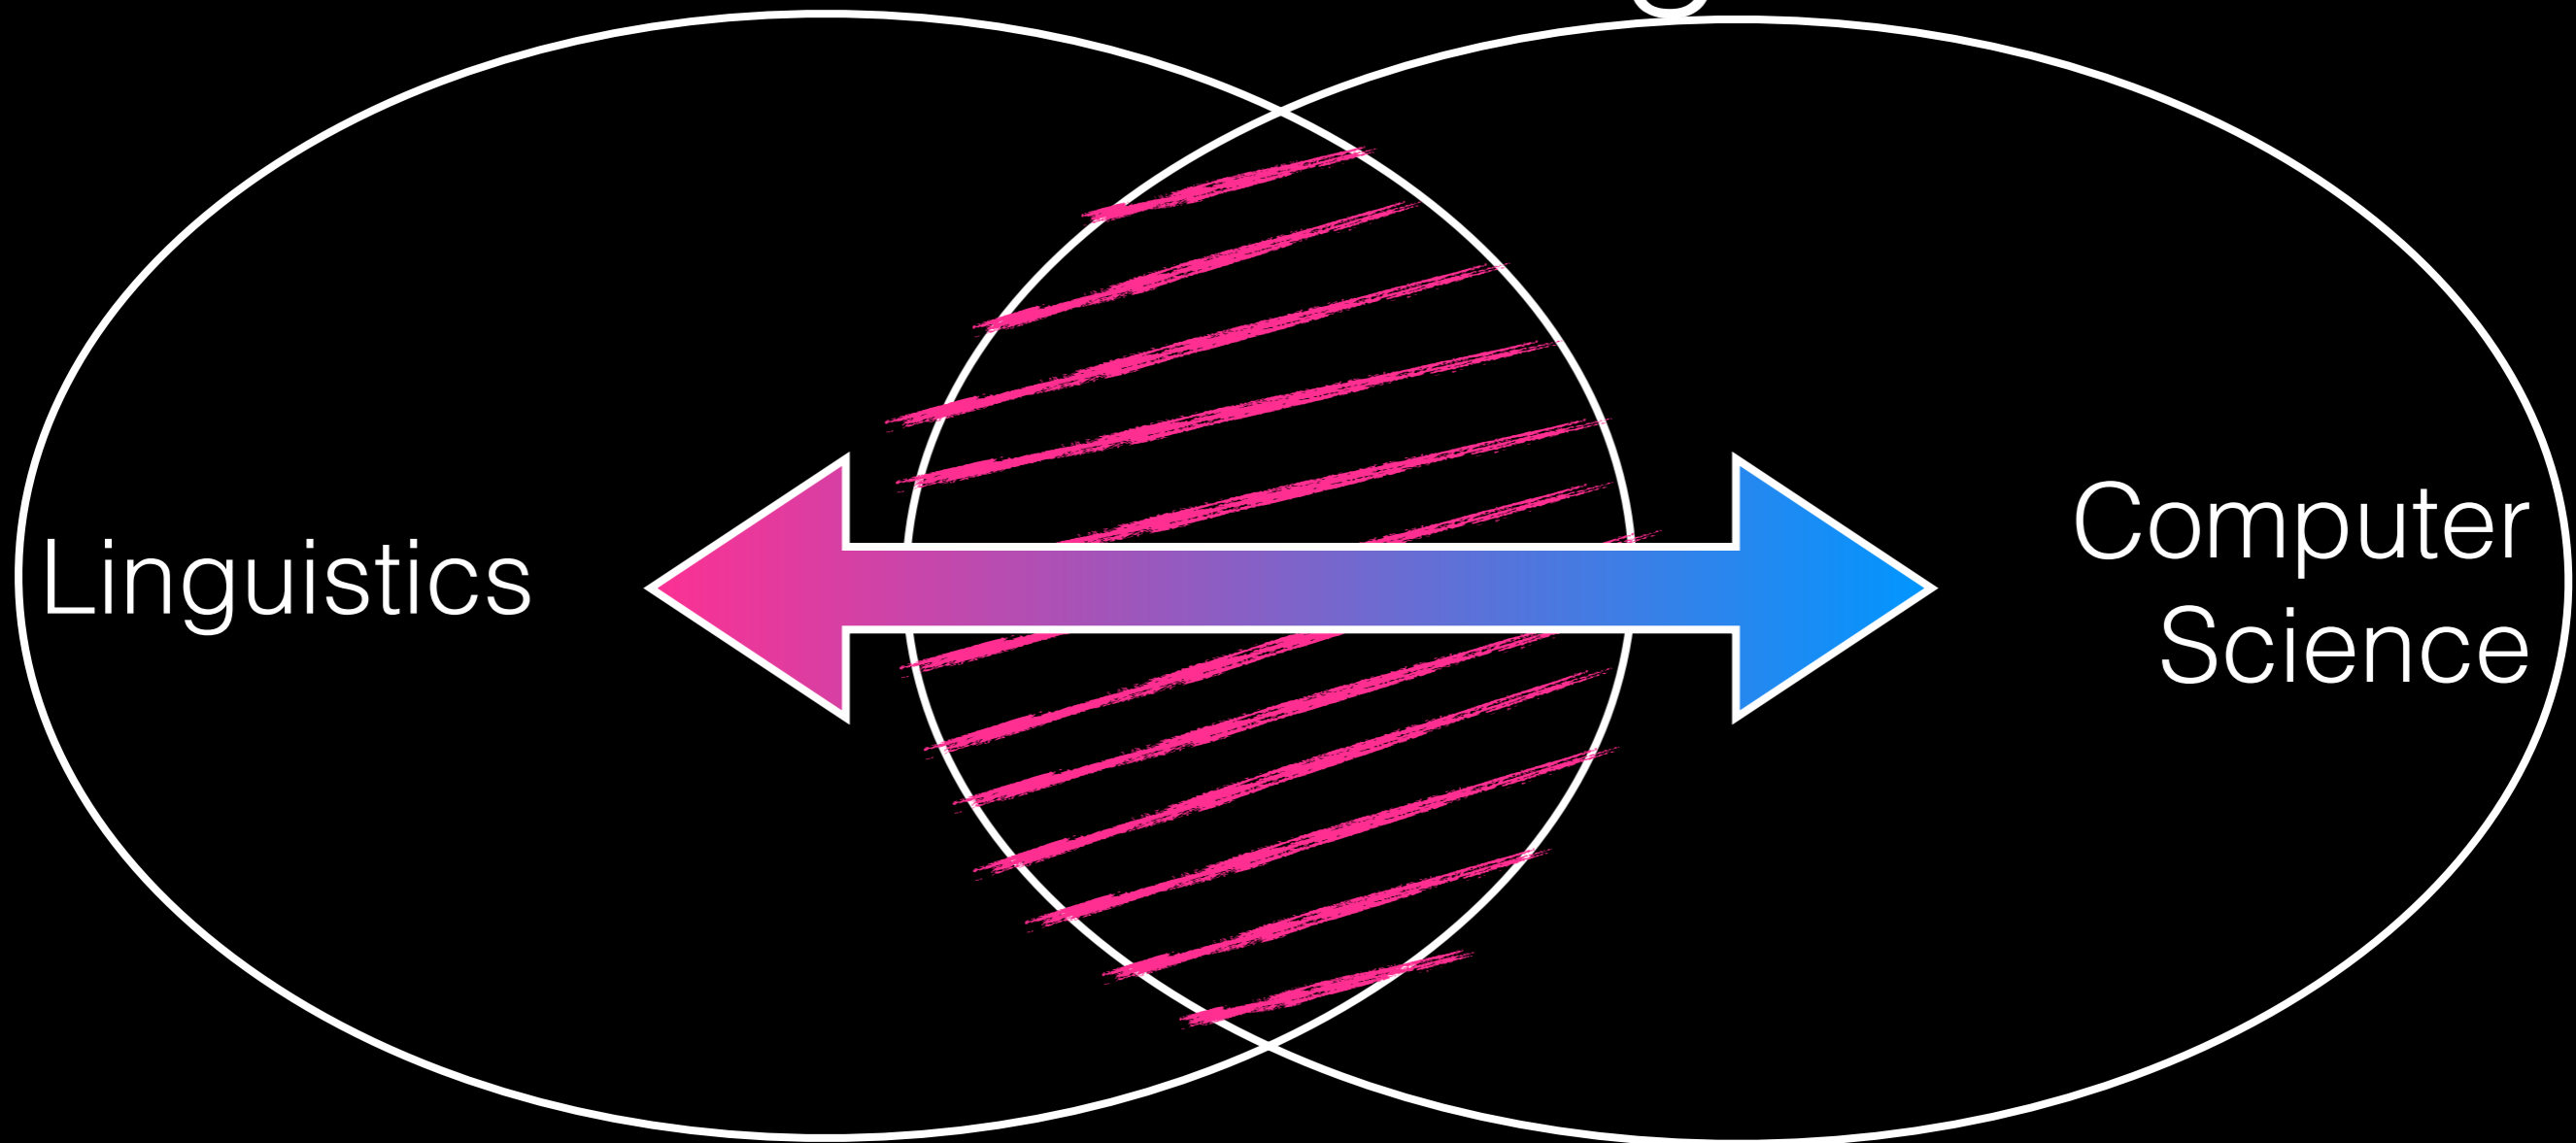
BUSINESS

PSYCHOLOGY POL. SCI.

SOCIOLOGY

SOCIOLOGISTS

Natural Language Processing



informed linguistic hypotheses large-scale statistical analysis

Natural Language Processing

SENT. ANALYSIS

positive

NER

O

O

PERSON

PERSON

O

PARSING

POS

PRON

VERB

NAME

NAME

PUNCT

I

admire Rosa Parks

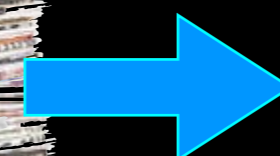
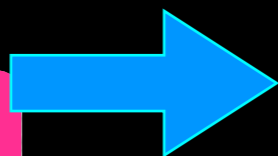
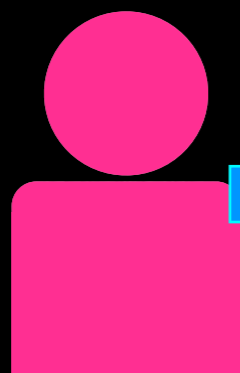
.

nsubj

dobj

nn

punct



MODEL



Who buys what?



Who buys what?



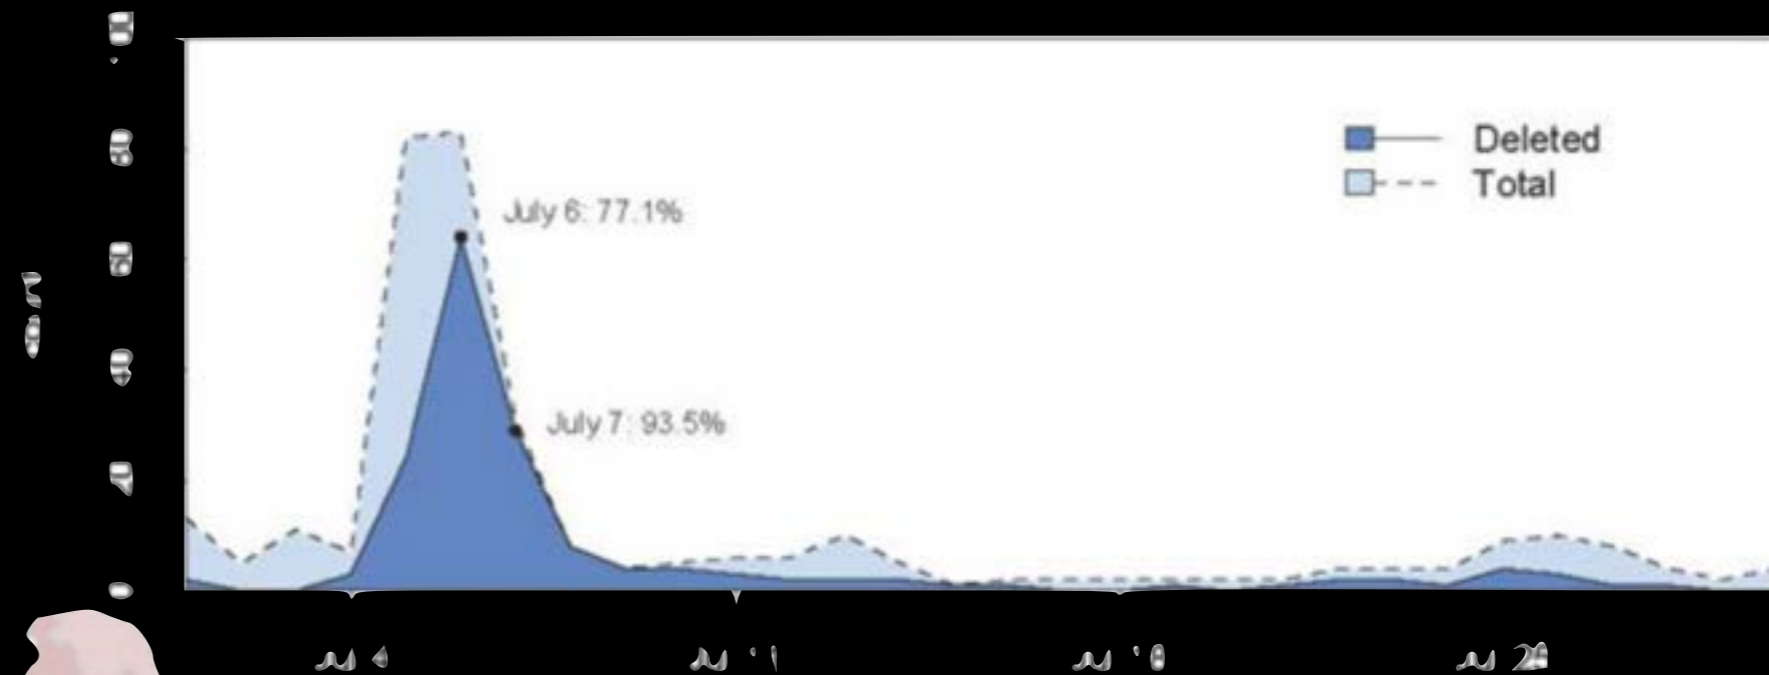
Sexism in Politics

Tone

- Positive
- Negative
- Neutral



Censorship



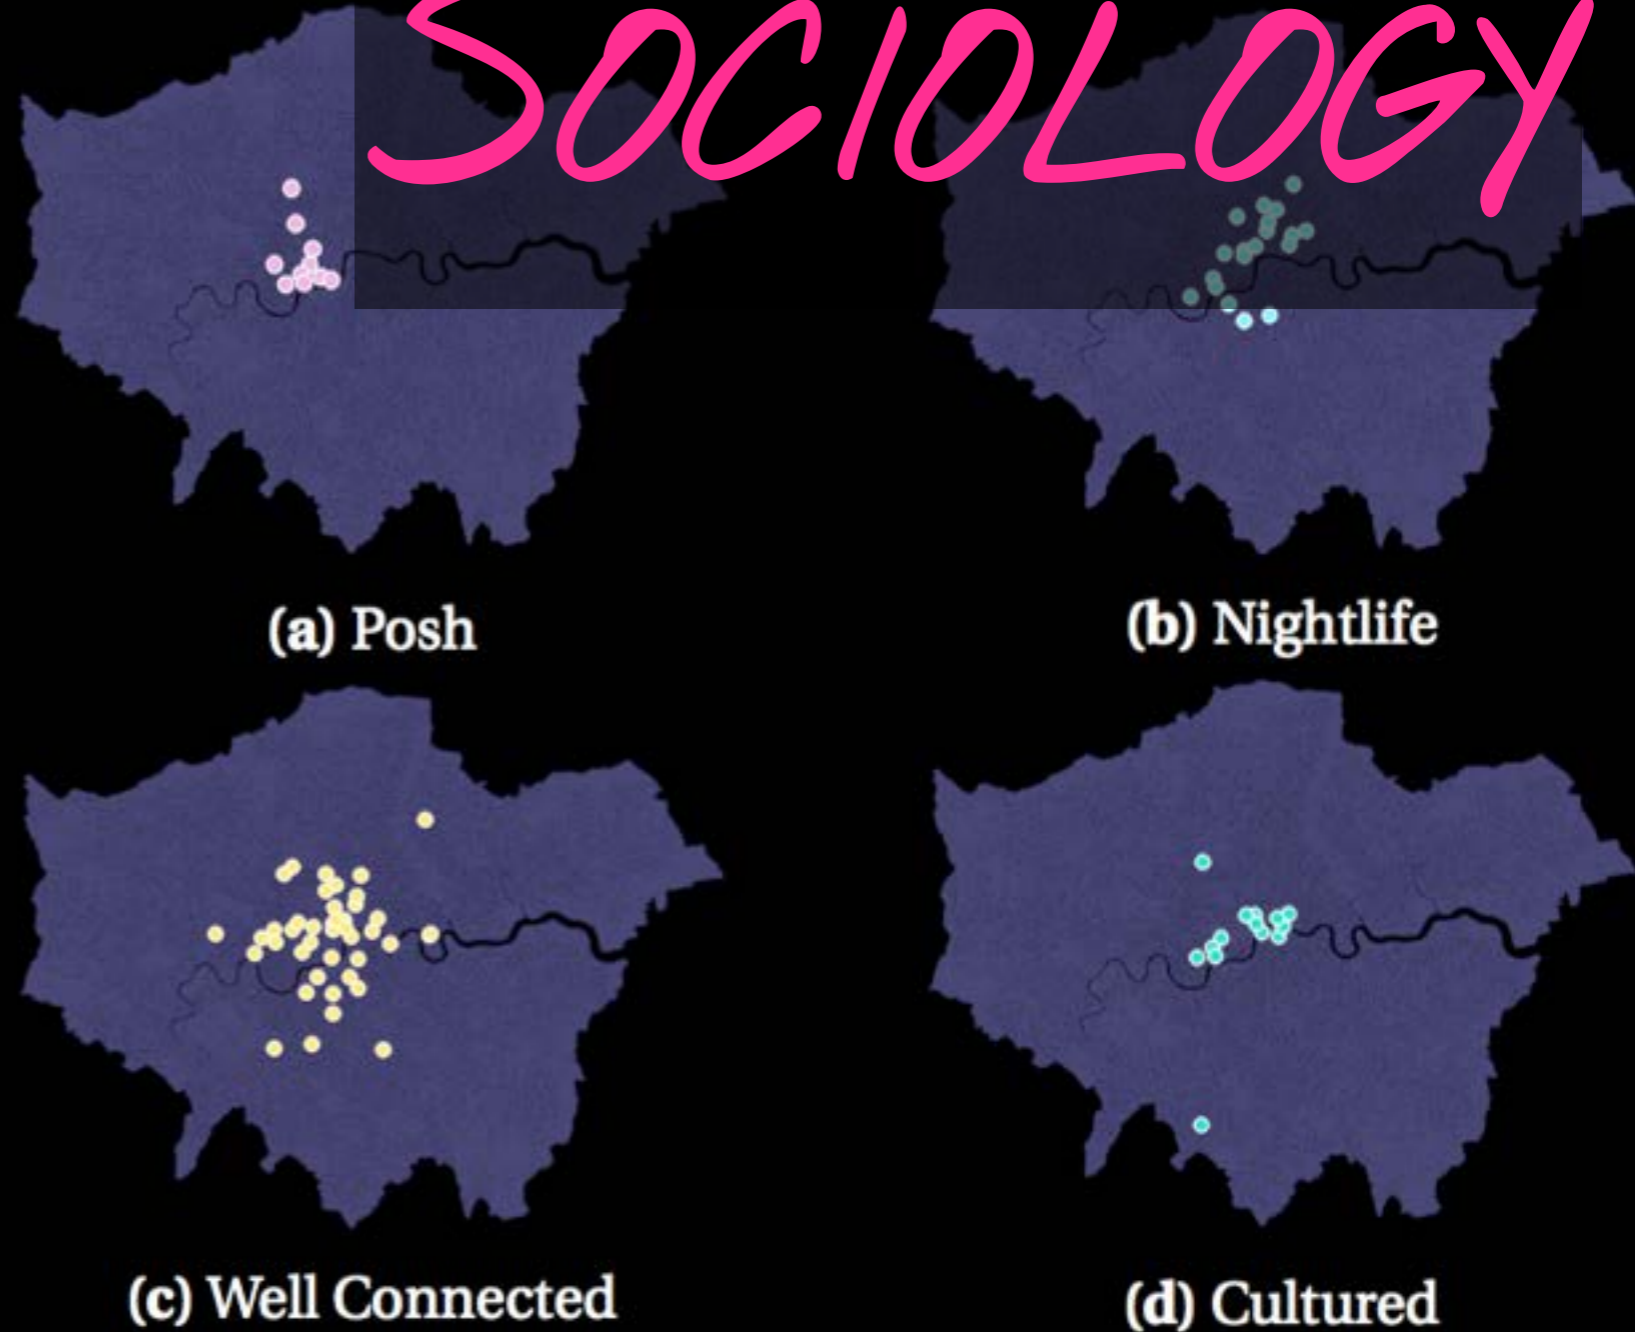
POL. SCI.

Neighborhoods

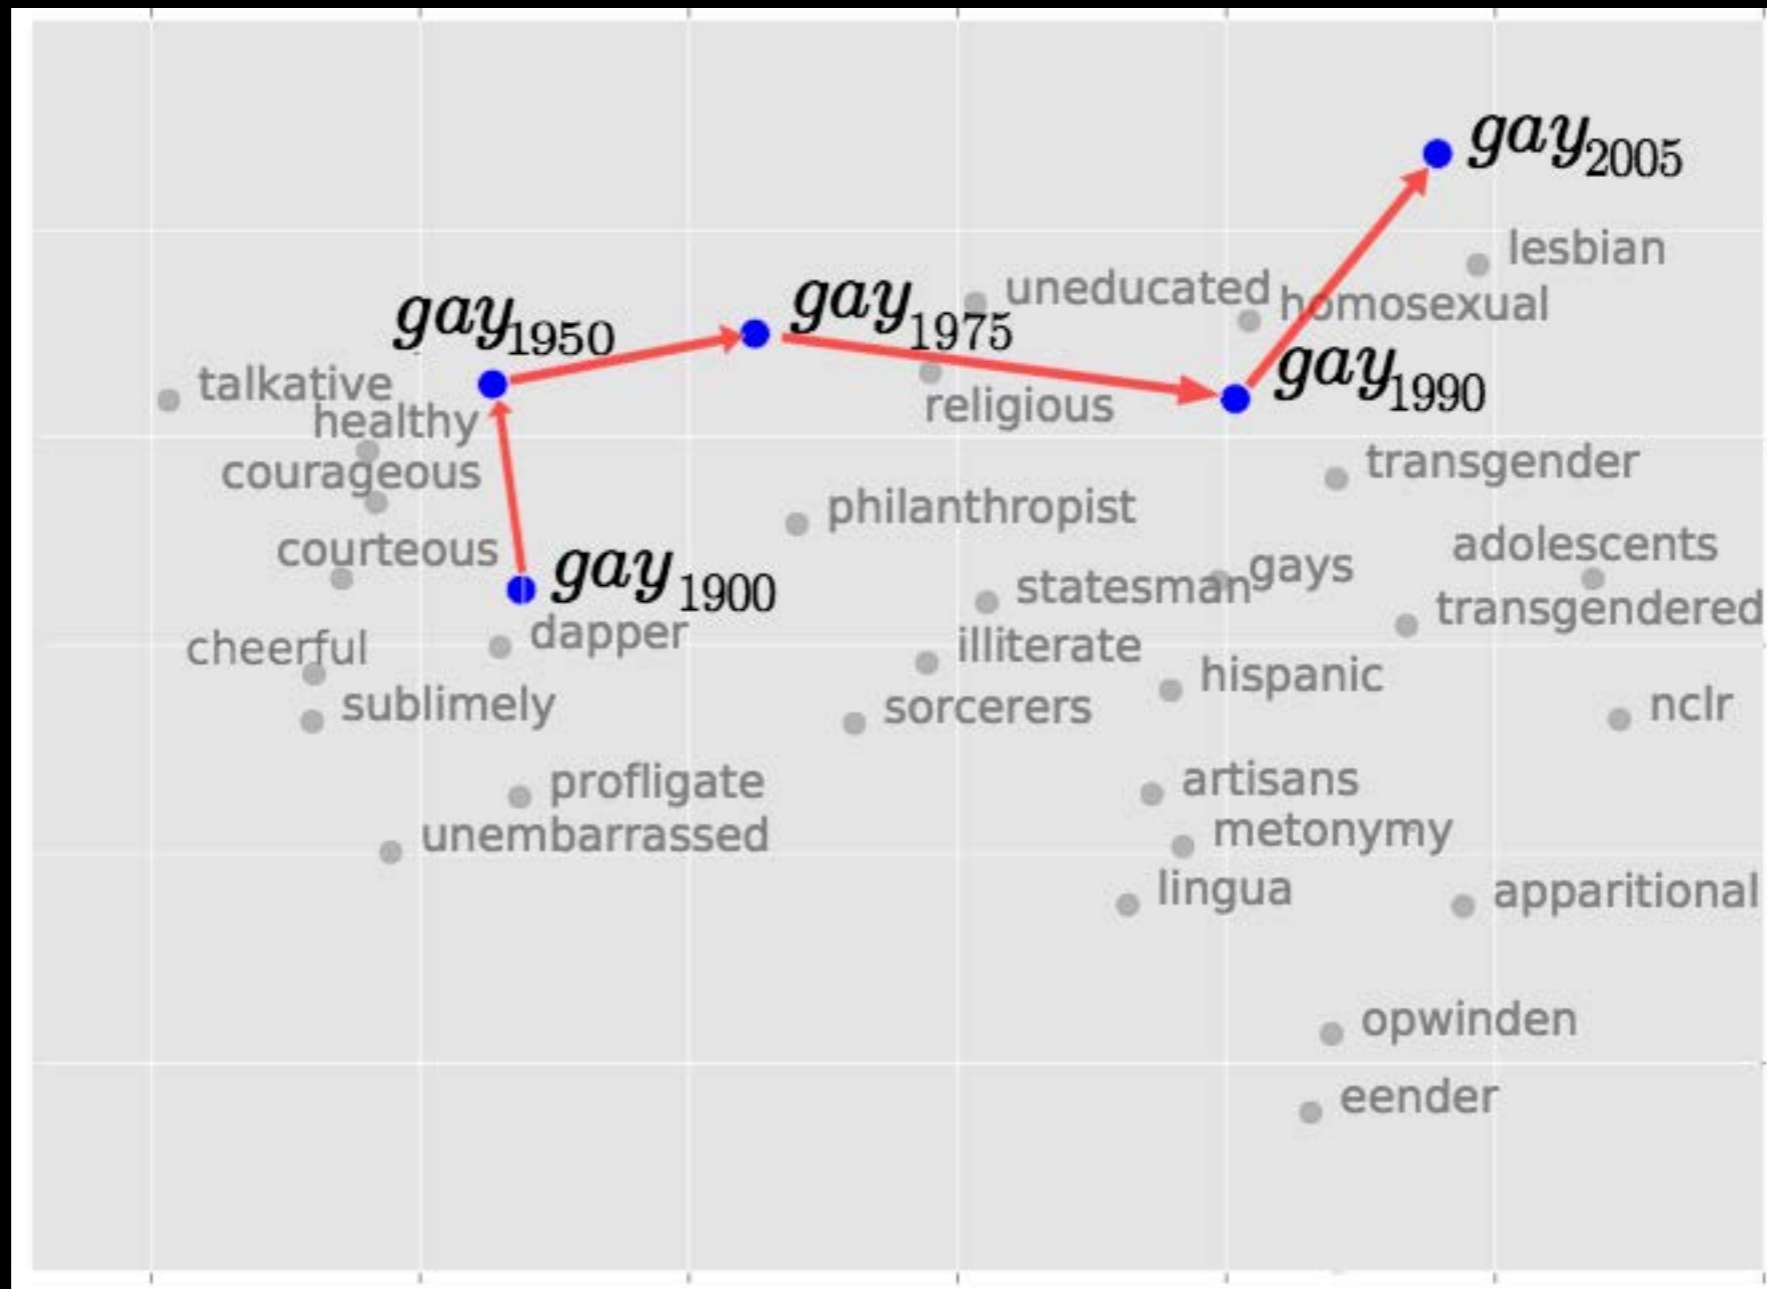
Q: What area of London should i live in?

A: Cool areas to live in at the moment are: / Clapham / Balham / Battersea / Hoxton / Camden

SOCIOLOGY



Language Change



SOCIOLINGUISTICS

Measuring Language Variation

SOCIOLINGUISTICS

Your search for *træls* found 2,794 results.

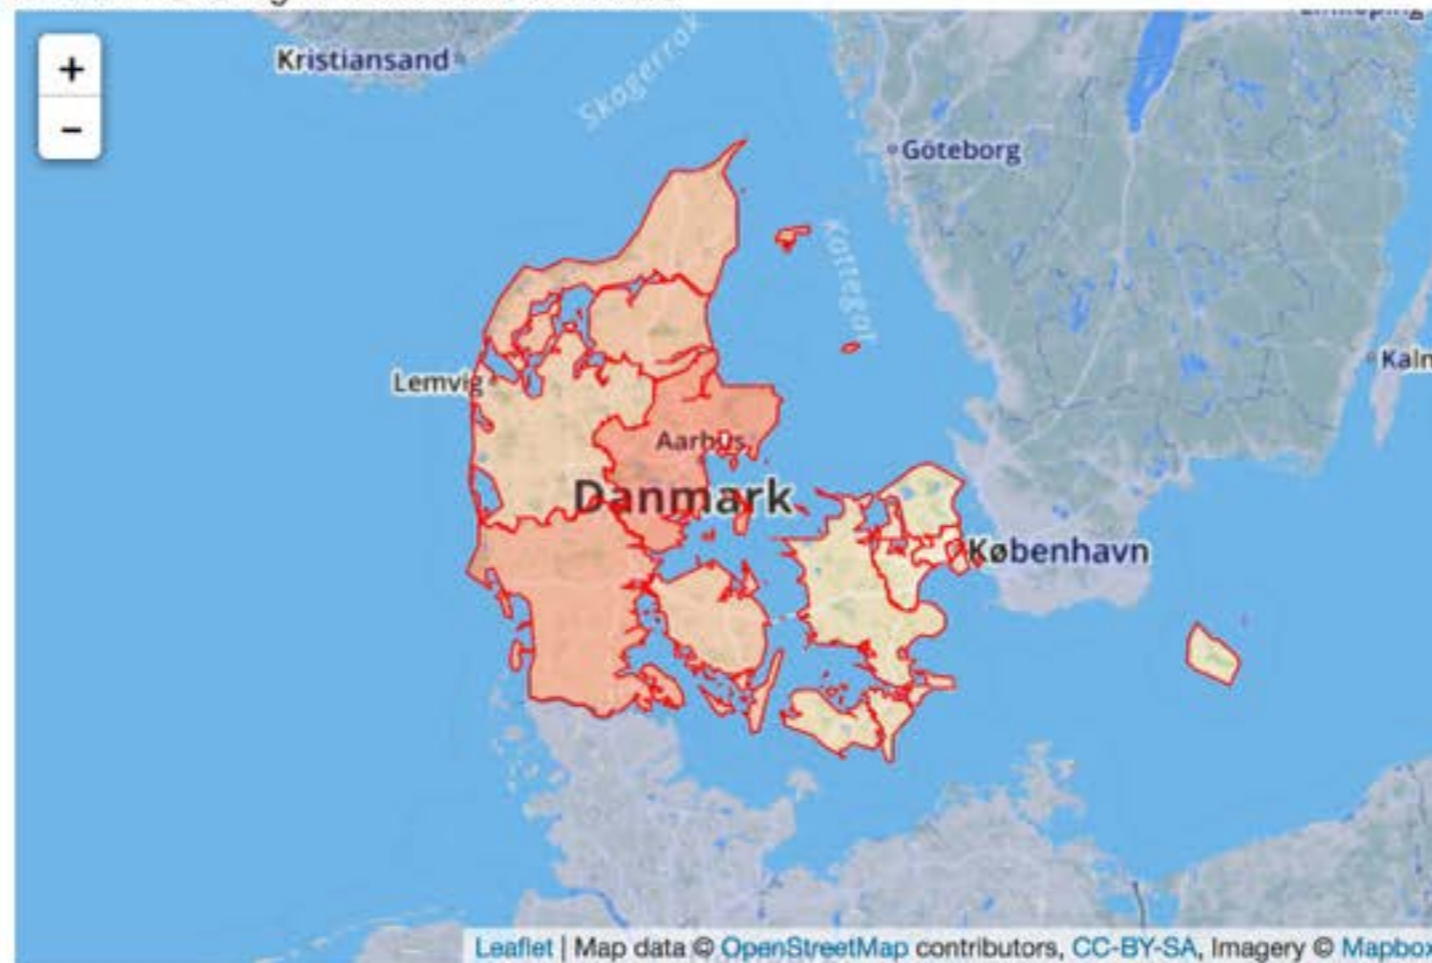
Note that we don't have the same amount of information for every demographic factor, so even though there might be enough data to display, e.g., a gender distribution, there might not be enough information to compute an age distribution.

Age and gender



Total records with NUTS information: 1,048

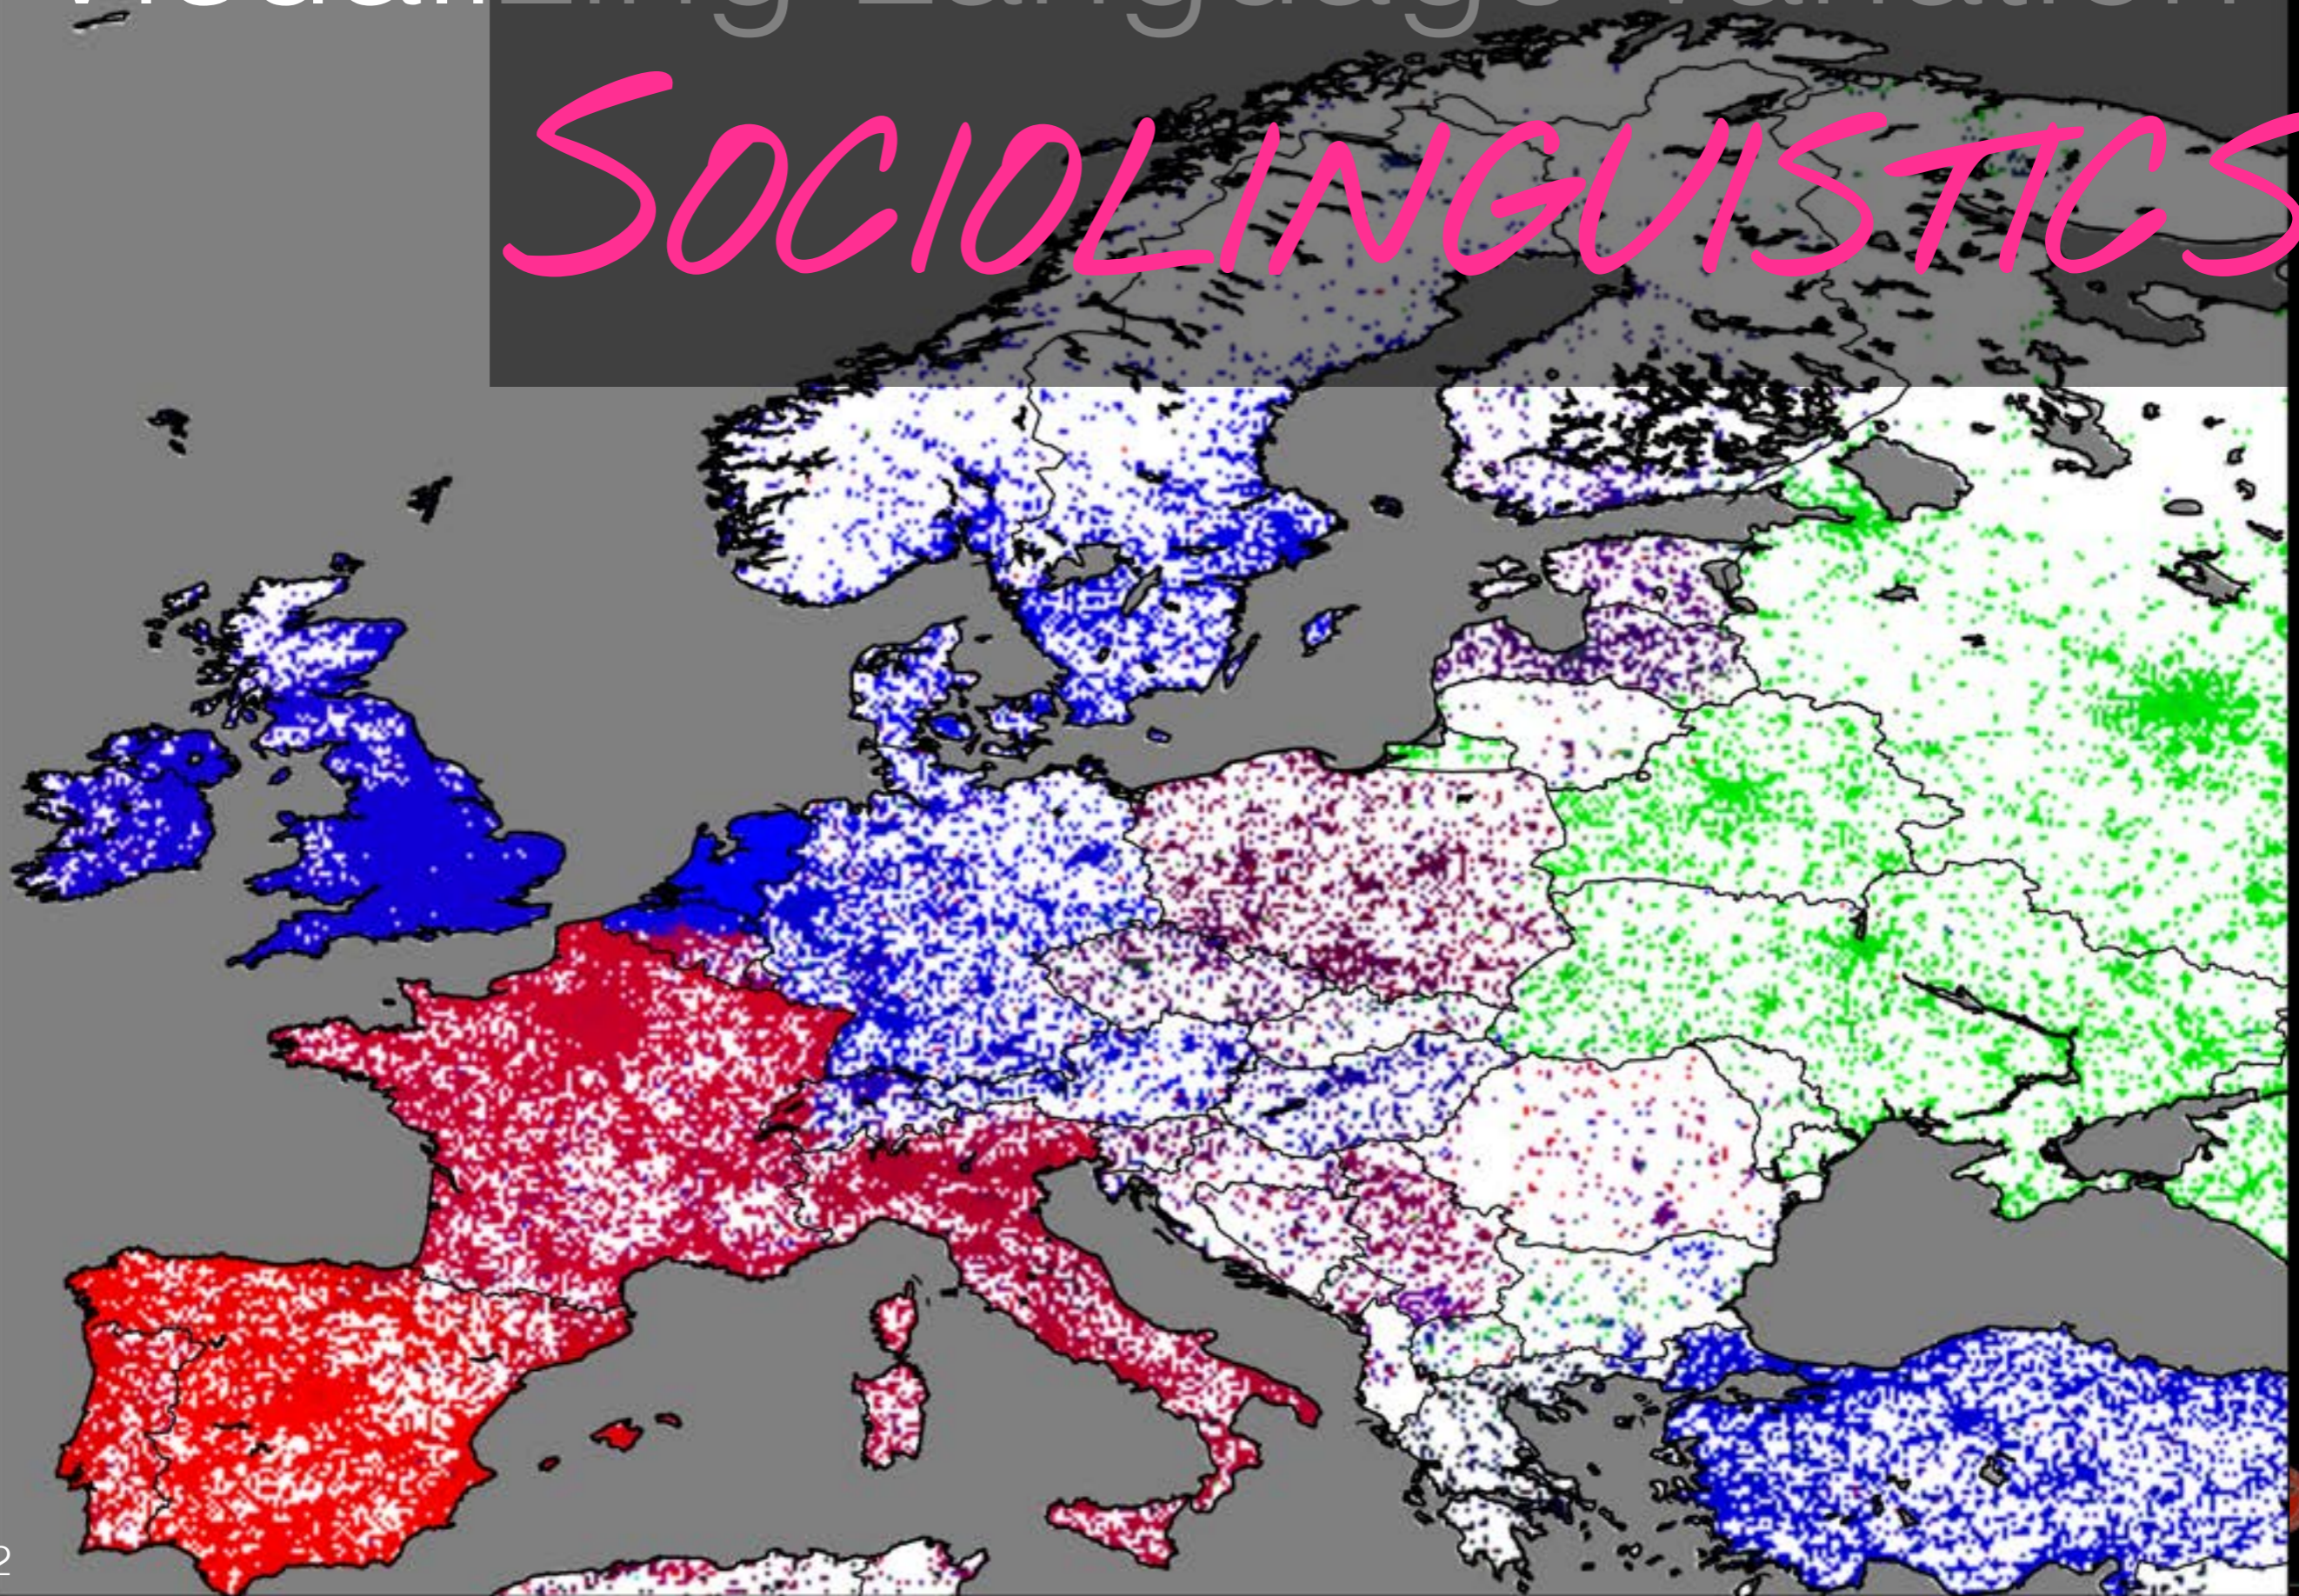
The map shows the percentage of records for each NUTS region that use the term *træls*.



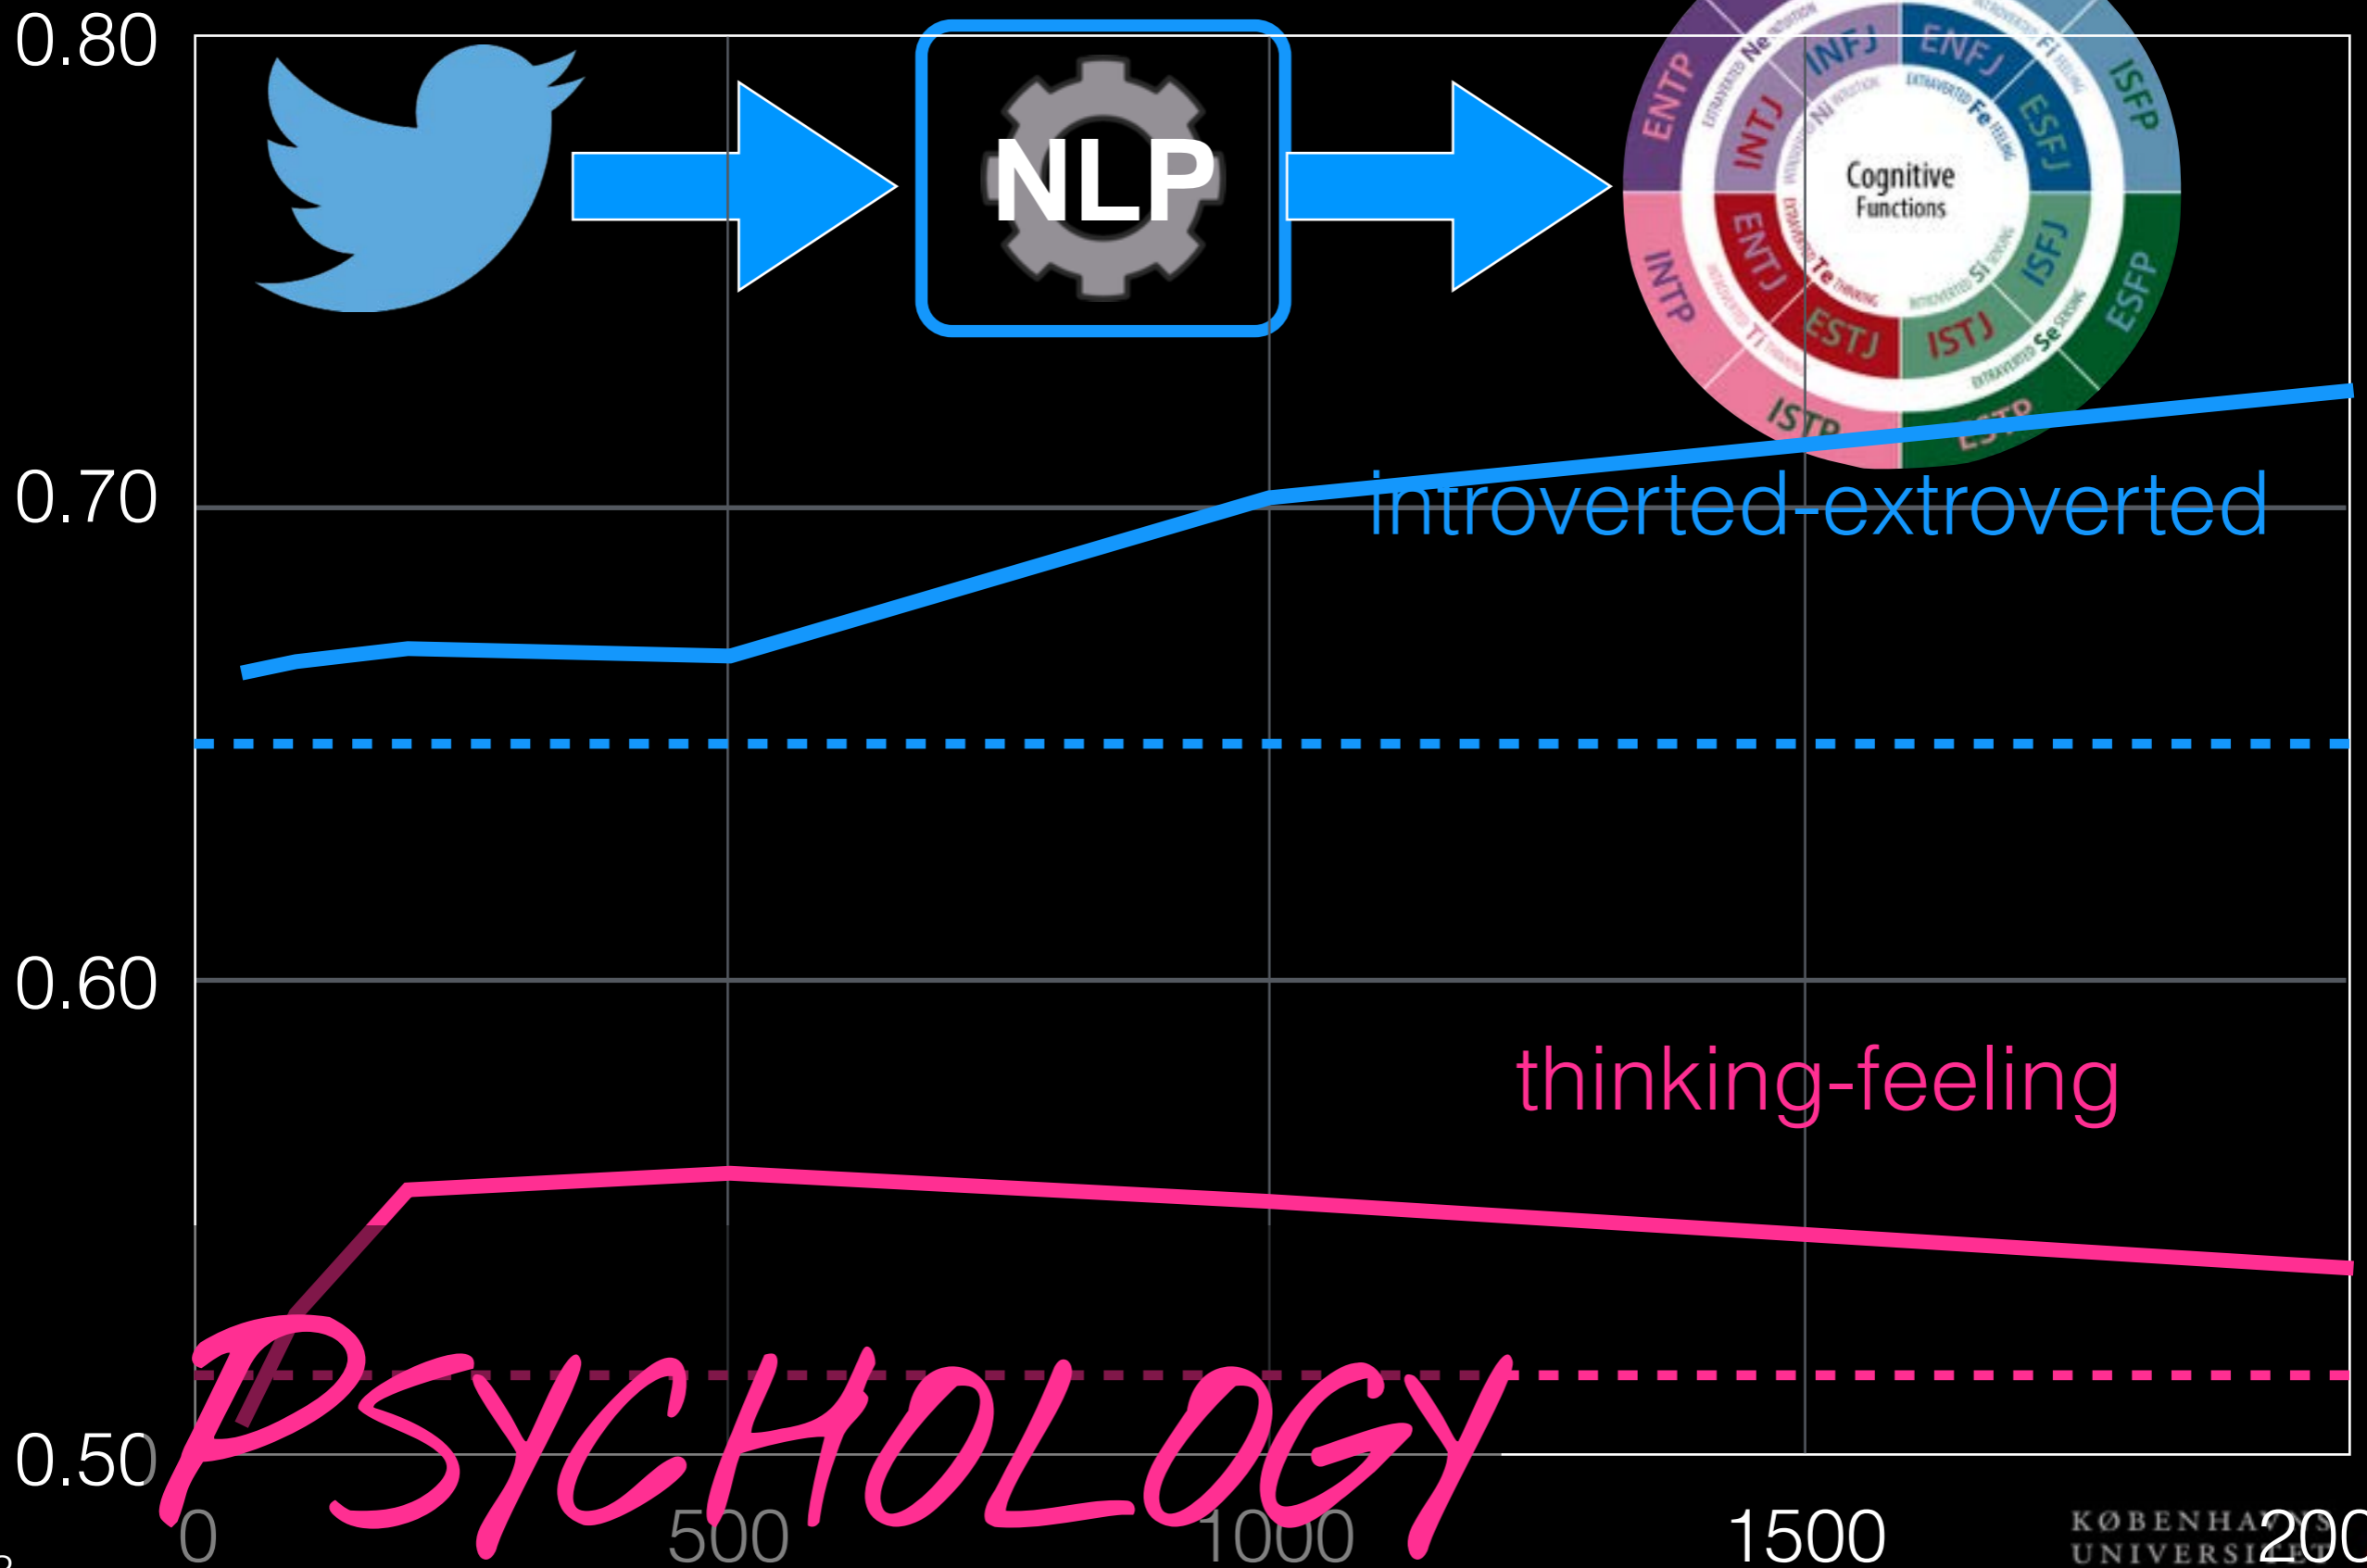
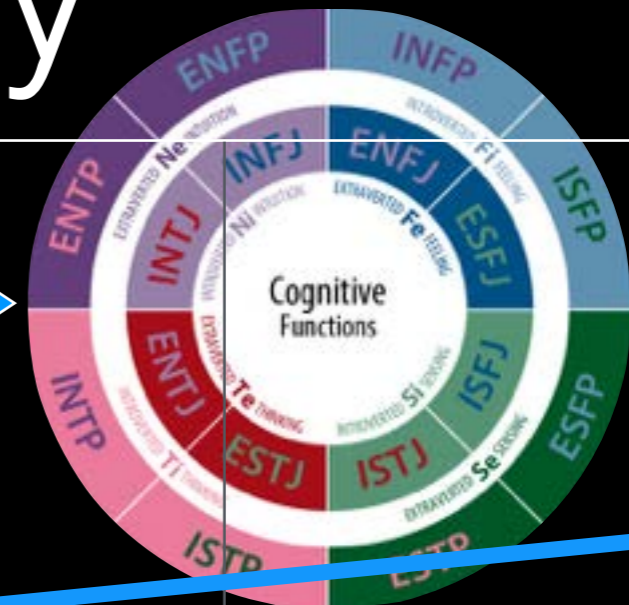
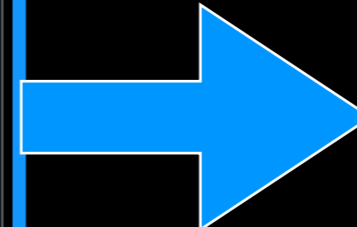
The term is used more frequently than average in following regions: Byen København (DK011), Fyn (DK031), Nordjylland (DK050), Sydjylland (DK032), Vestjylland (DK041), Østjylland (DK042) (the ratio of term in the vocabulary of these regions is larger than the average ratio across regions).

Visualizing Language Variation

SOCIOLINGUISTICS



Personality



The Transparent User



age	Rosenthal and McKeown, 2011; Nguyen et al., 2011
gender	Alowibdi et al., 2013; Ciot et al., 2013; Liu and Ruths, 2013; Volkova et
profession	Preoțiu-Pietro et al., 2015
income	Preoțiu-Pietro et al., 2015b
personality	Plank and Hovy, 2015
political views	Volkova et al., 2014
location	Bergsma et al., 2013; Rahimi et al., 2016

PSYCHOLOGY



Workshop on NLP and Computational Social Sciences

@ACL

Aug 3, 2017

Vancouver, Canada

<https://sites.google.com/site/nlpandcss/>



@nlpandcss

That's it, NLP can help
social science and business

All is great!

...no,

really...

Ok, so maybe
there's *one* thing

The Assumption

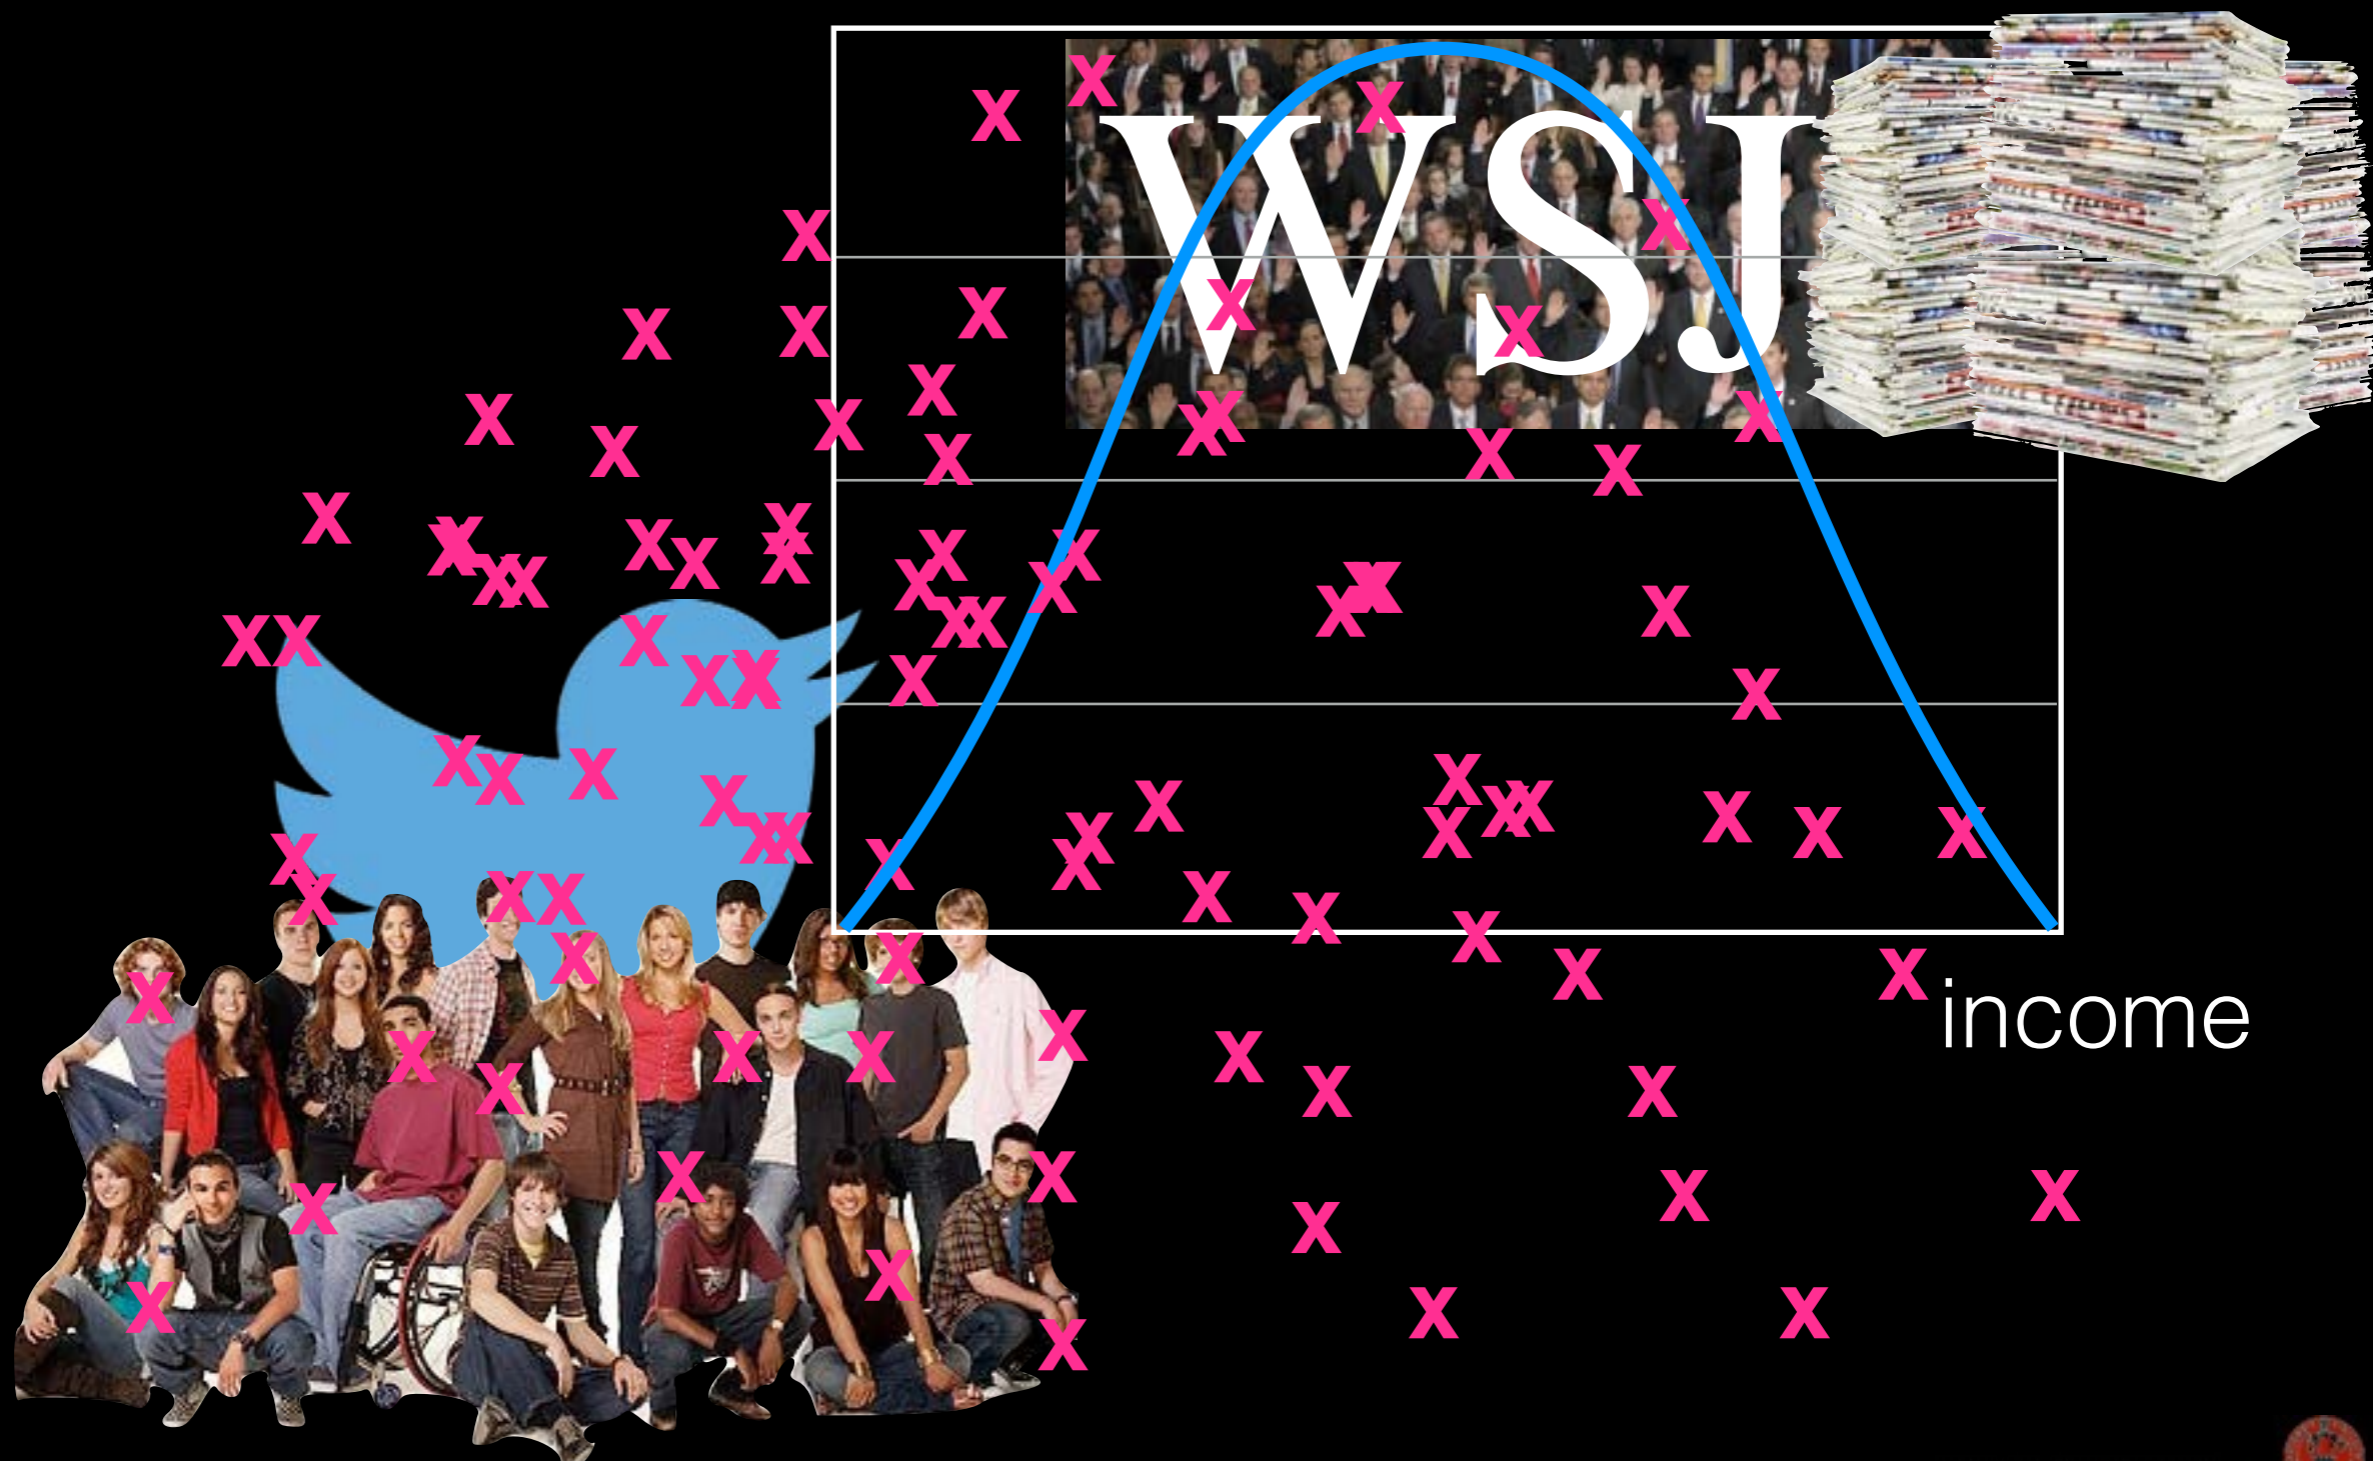
=

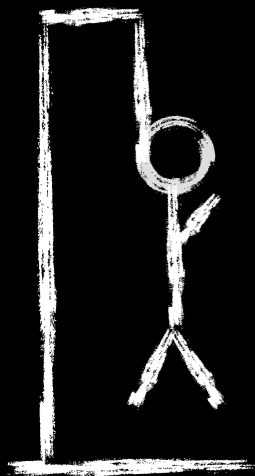


All language is i.i.d.

Language Distributions

age





Shannon Game



WHAT'S THE
NEXT WORD?

The house
A friend
Then dog
If car
When water
My hovercraft
He pants
You God
I word
...



water ?

eels .

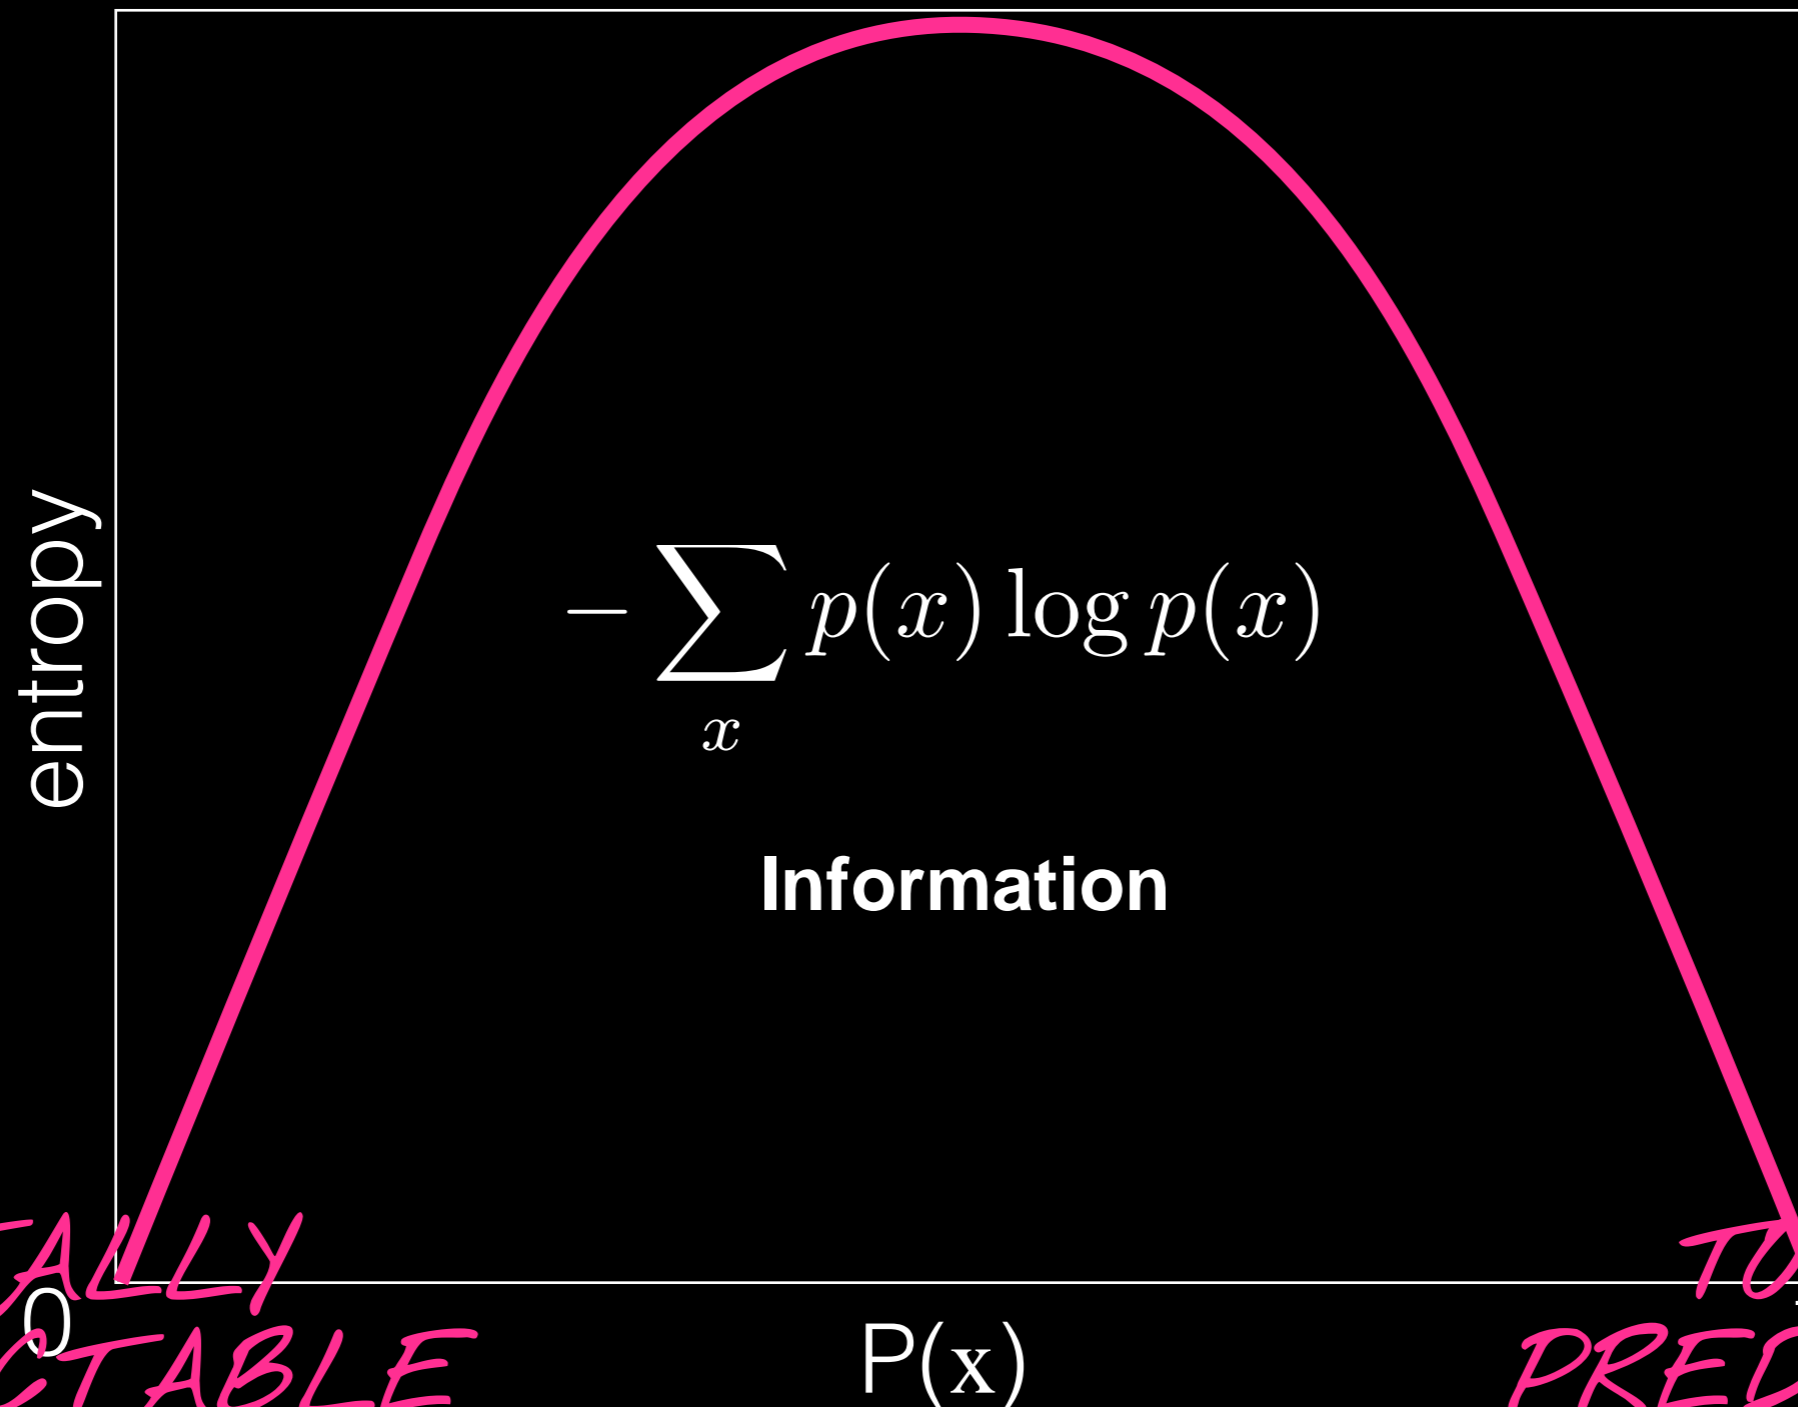
air

HOW MUCH
LONGER?



Shannon Game

VERY SURPRISING



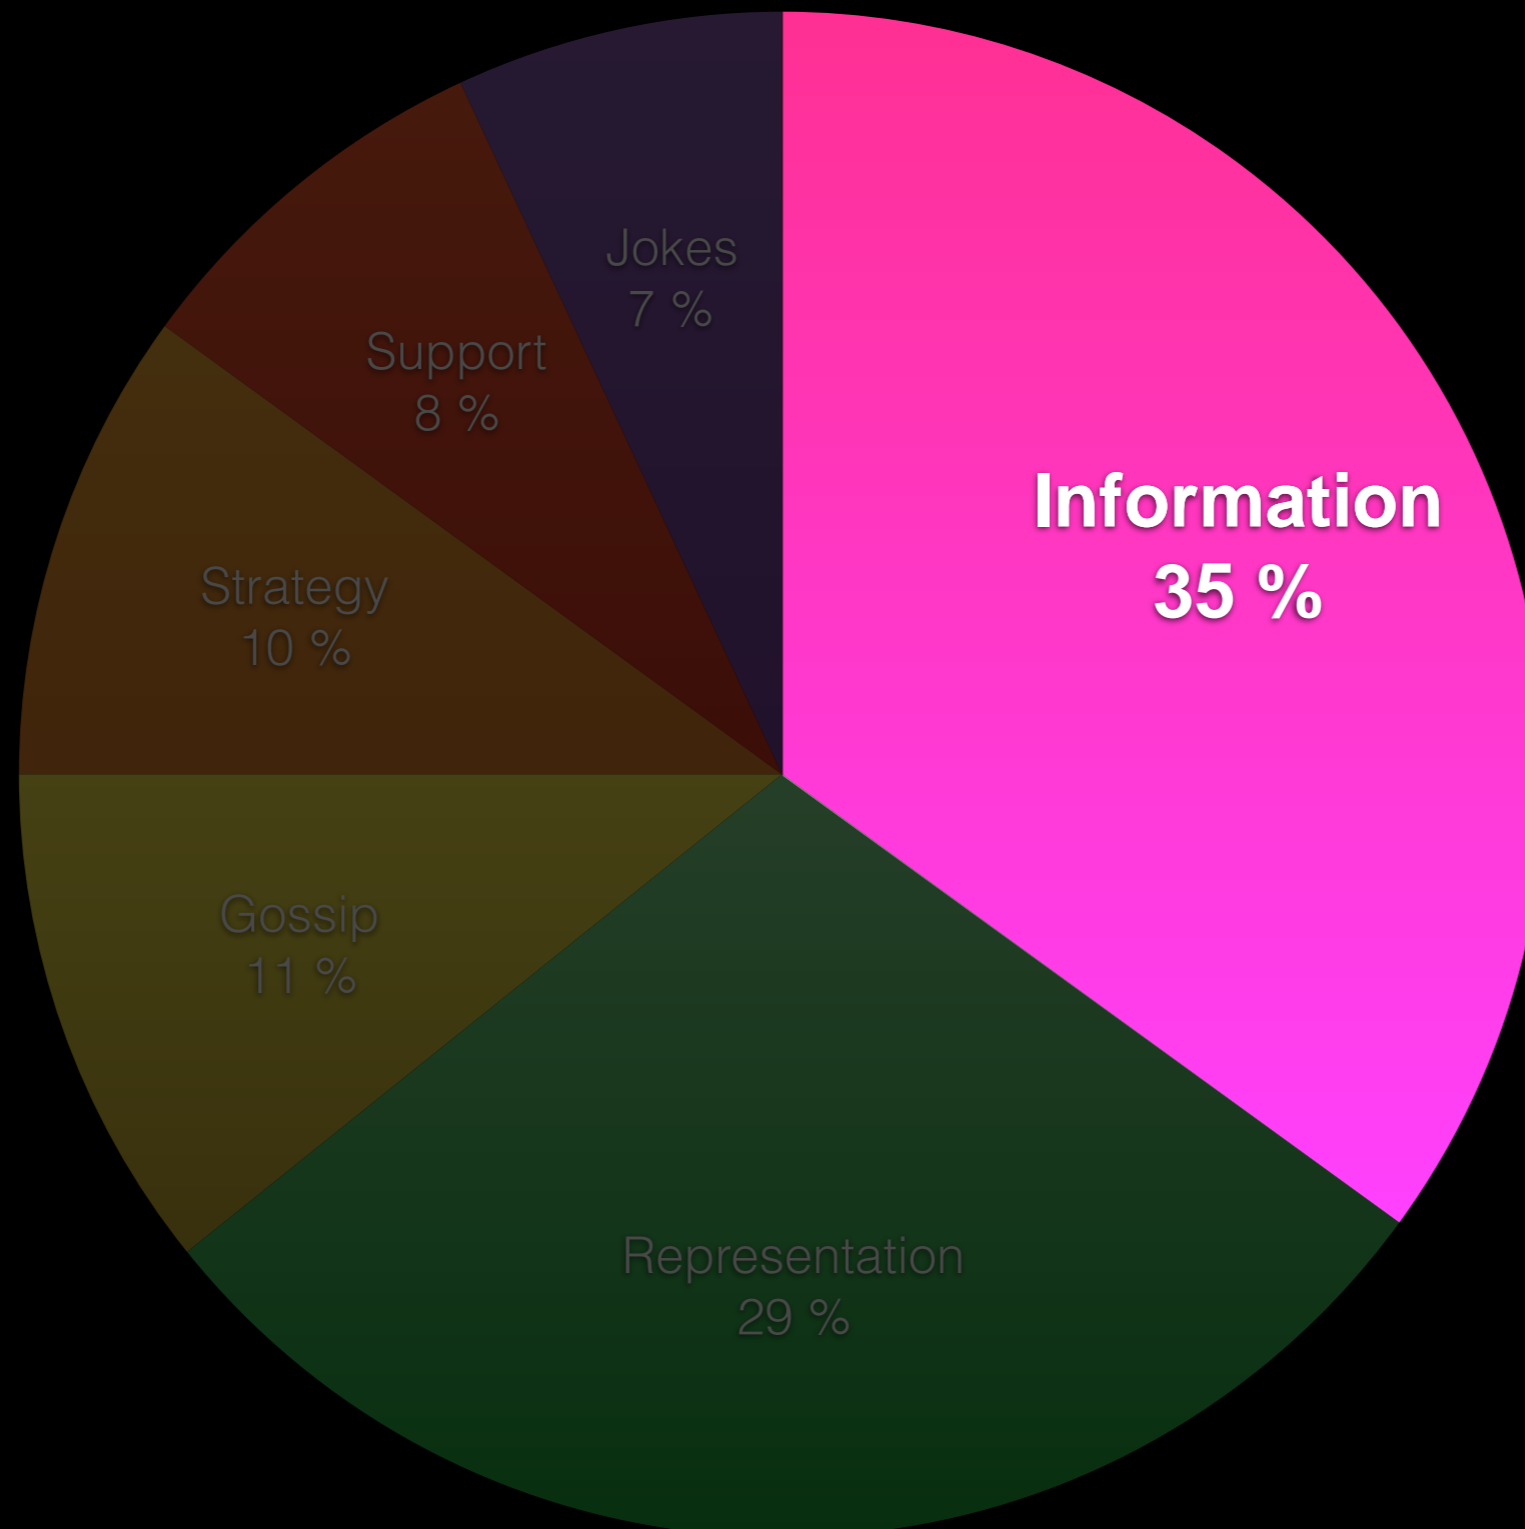
*TOTALLY
PREDICTABLE*

*TOTALLY
PREDICTABLE*

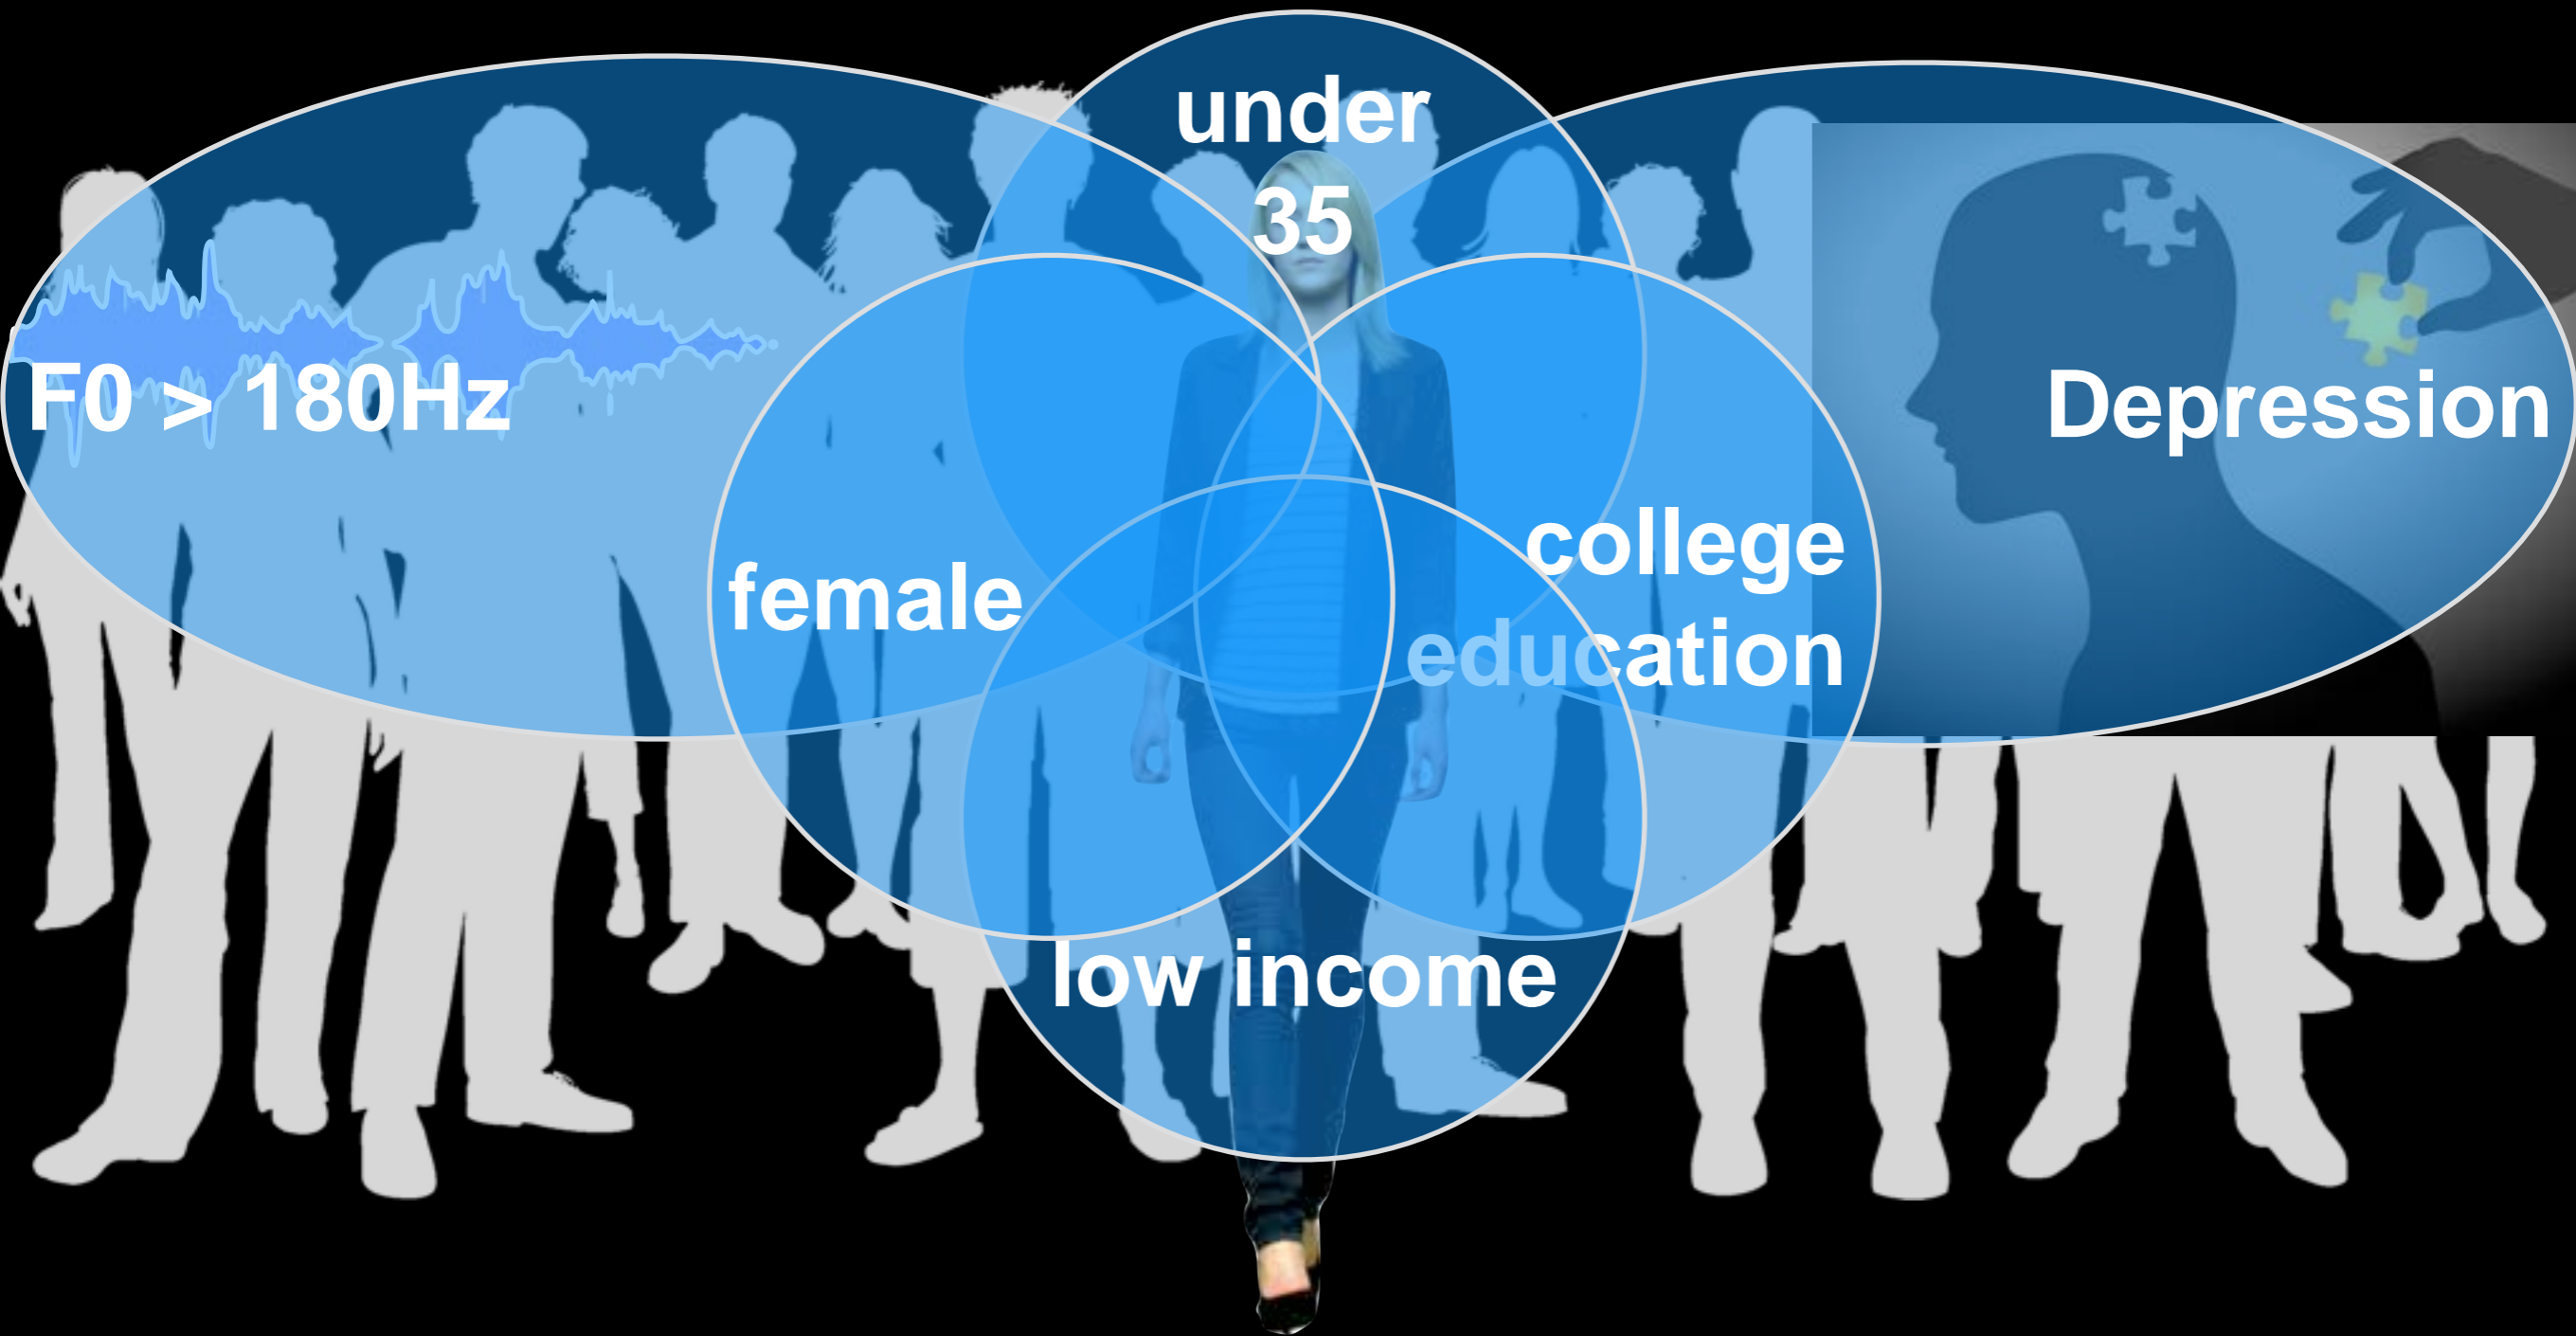
COPENHAGEN
UNIVERSITY



Language

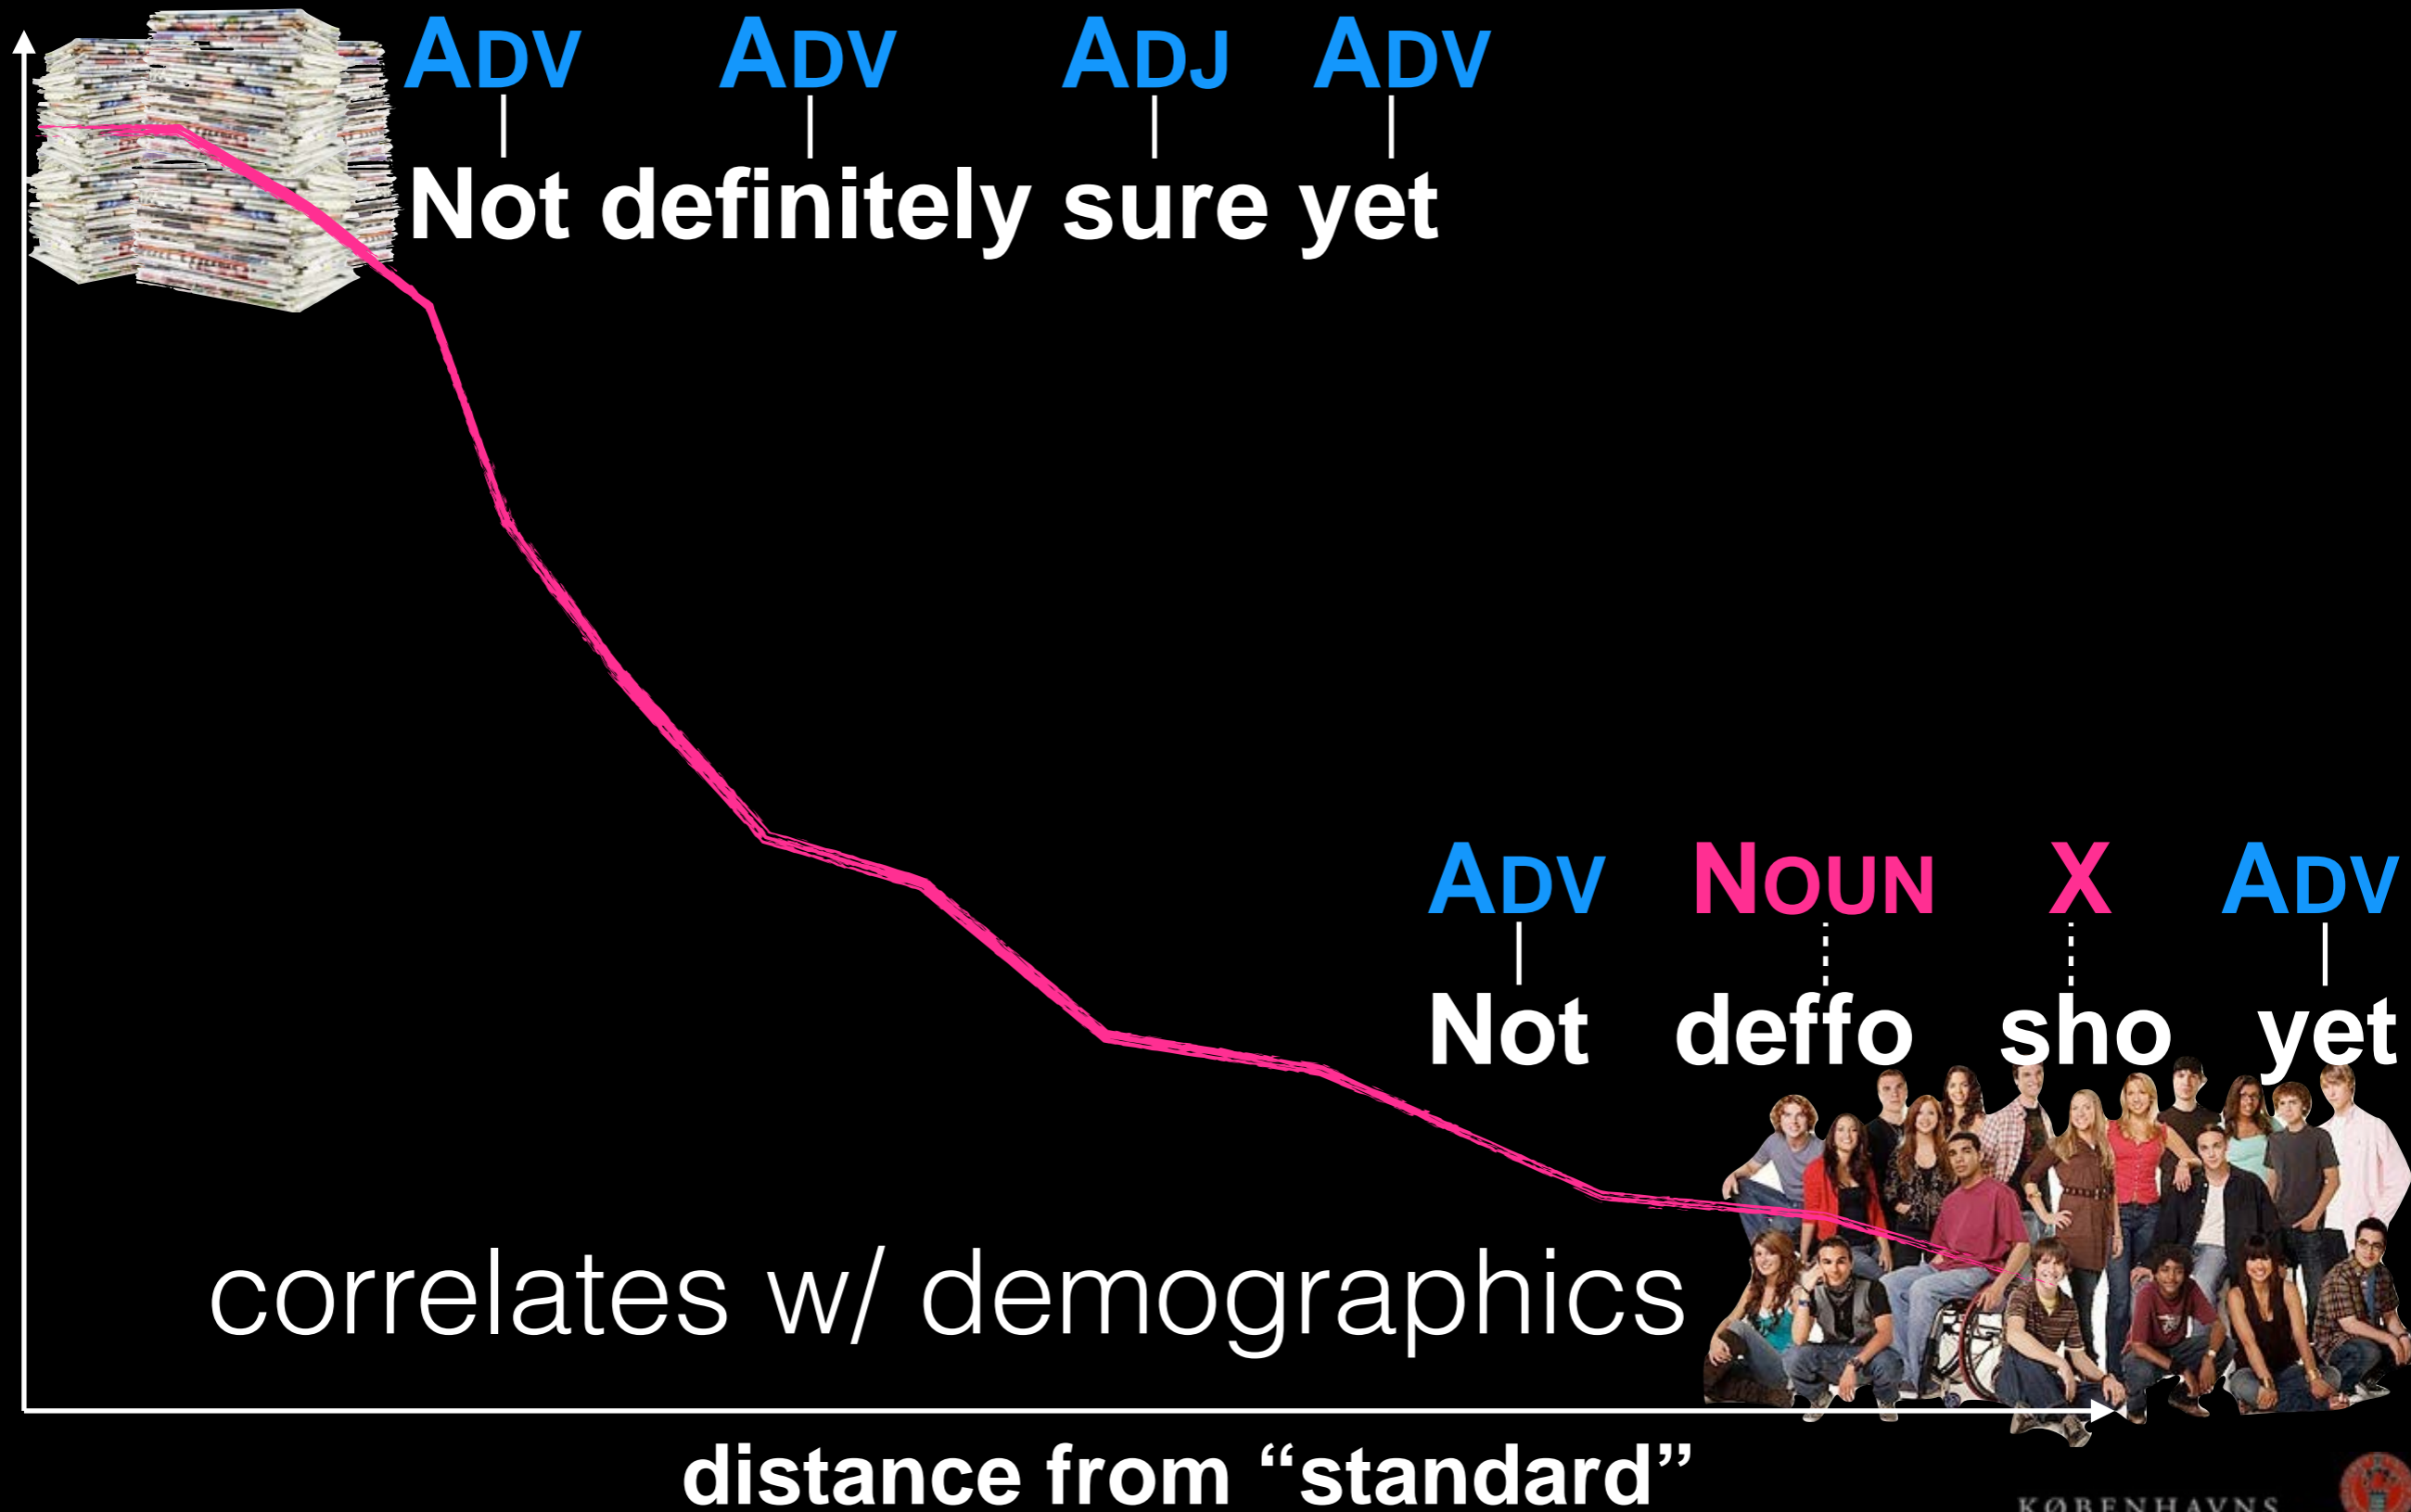


Demographics

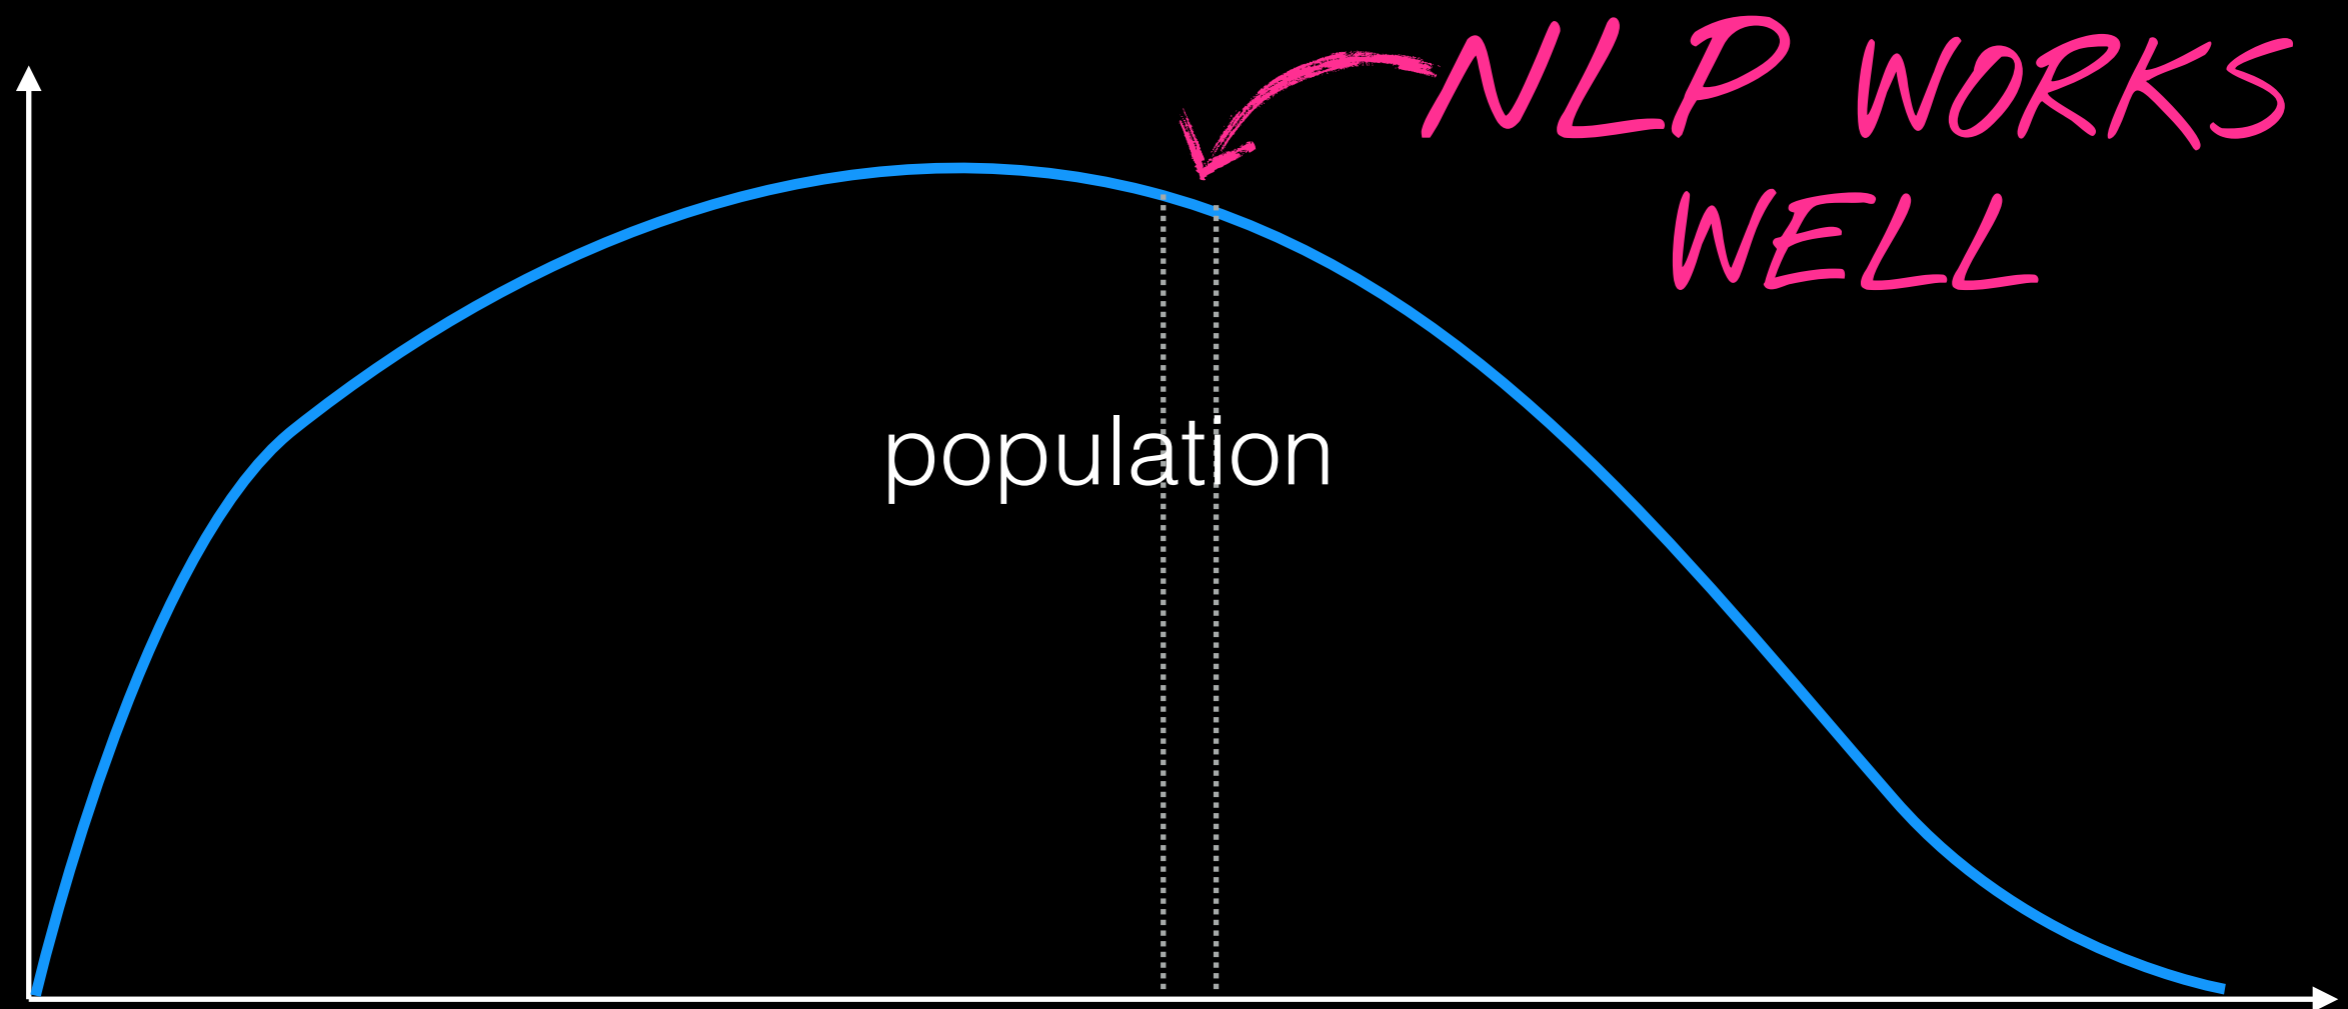


NLP
performance

The Problem



The Consequences

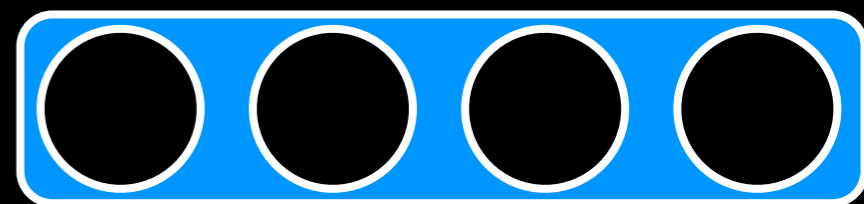
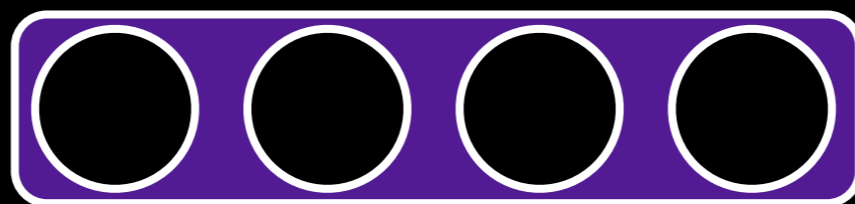
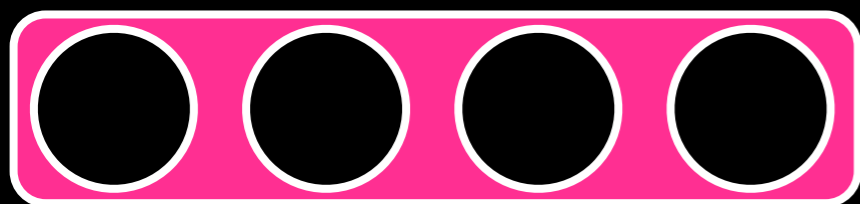


More D

HOW CAN WE ADD
EXTRA-LINGUISTIC INFO TO
NLP MODELS?

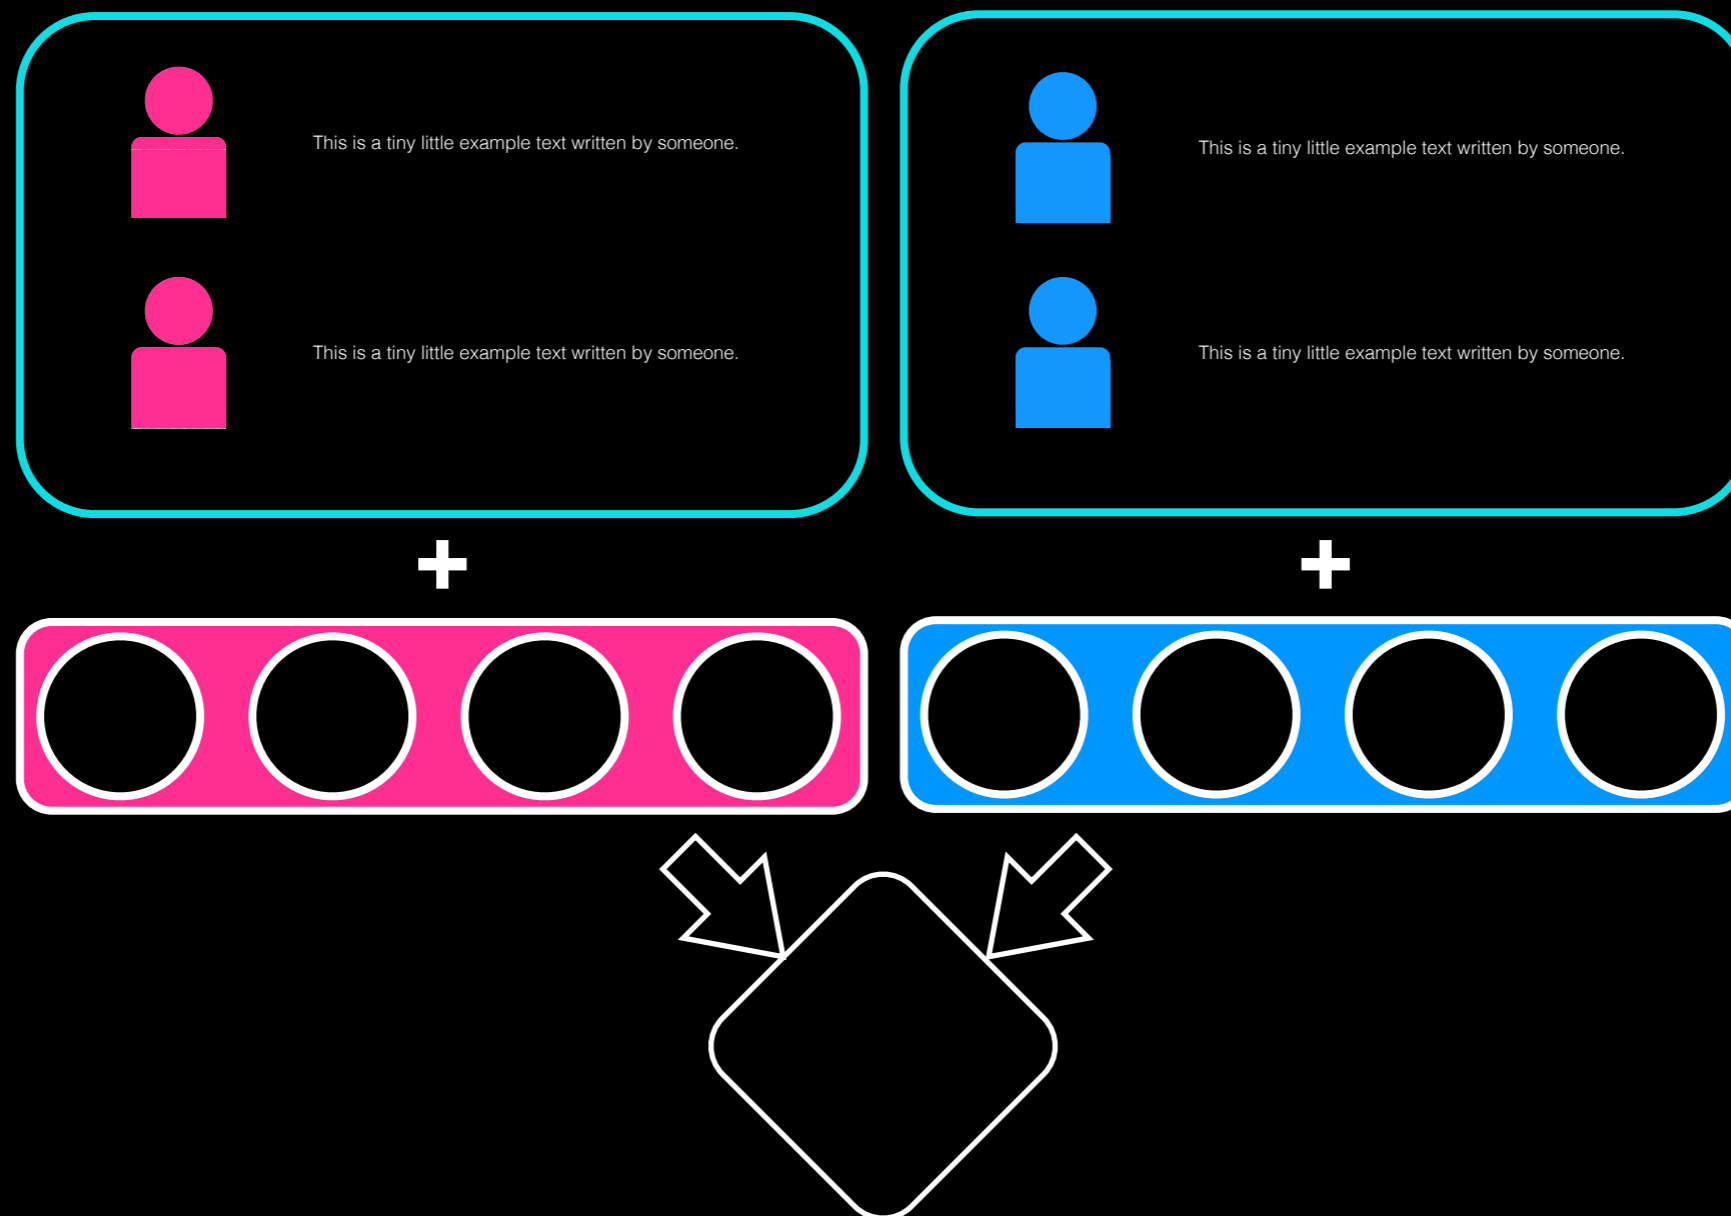


Improving NLP with Demographic Information



Example 1:

Text Classification



Data

    , 2 genders, 2 age groups

task	labels	train	test
topic classification	5	739k 126k	493k 84k
sentiment analysis	3	345k 72k	230k 48k
age/gender classification	2	301k 301k	201k 201k

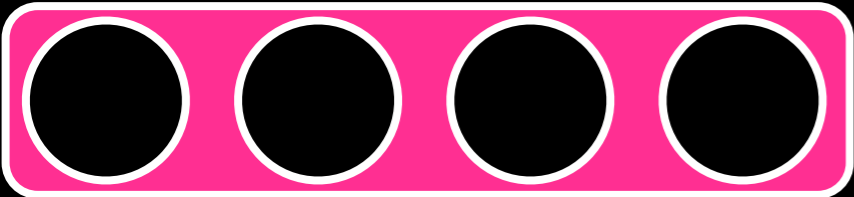



Embeddings




This is a tiny little example text written by someone.





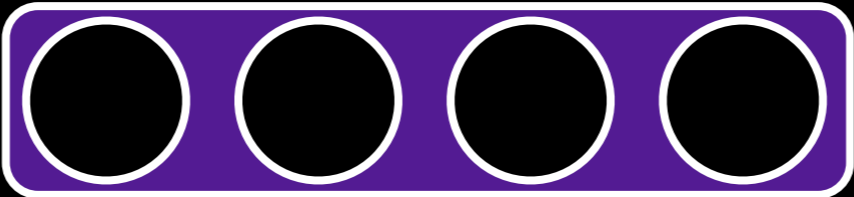


This is a tiny little example text written by someone.



This is a tiny little example text written by someone.

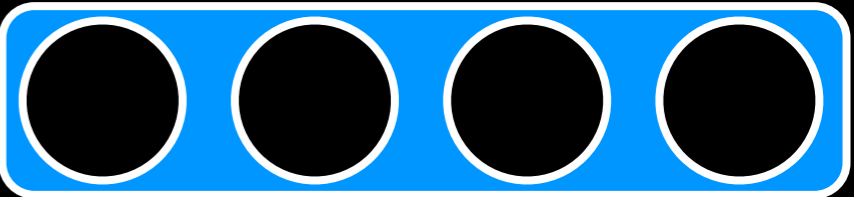






This is a tiny little example text written by someone.



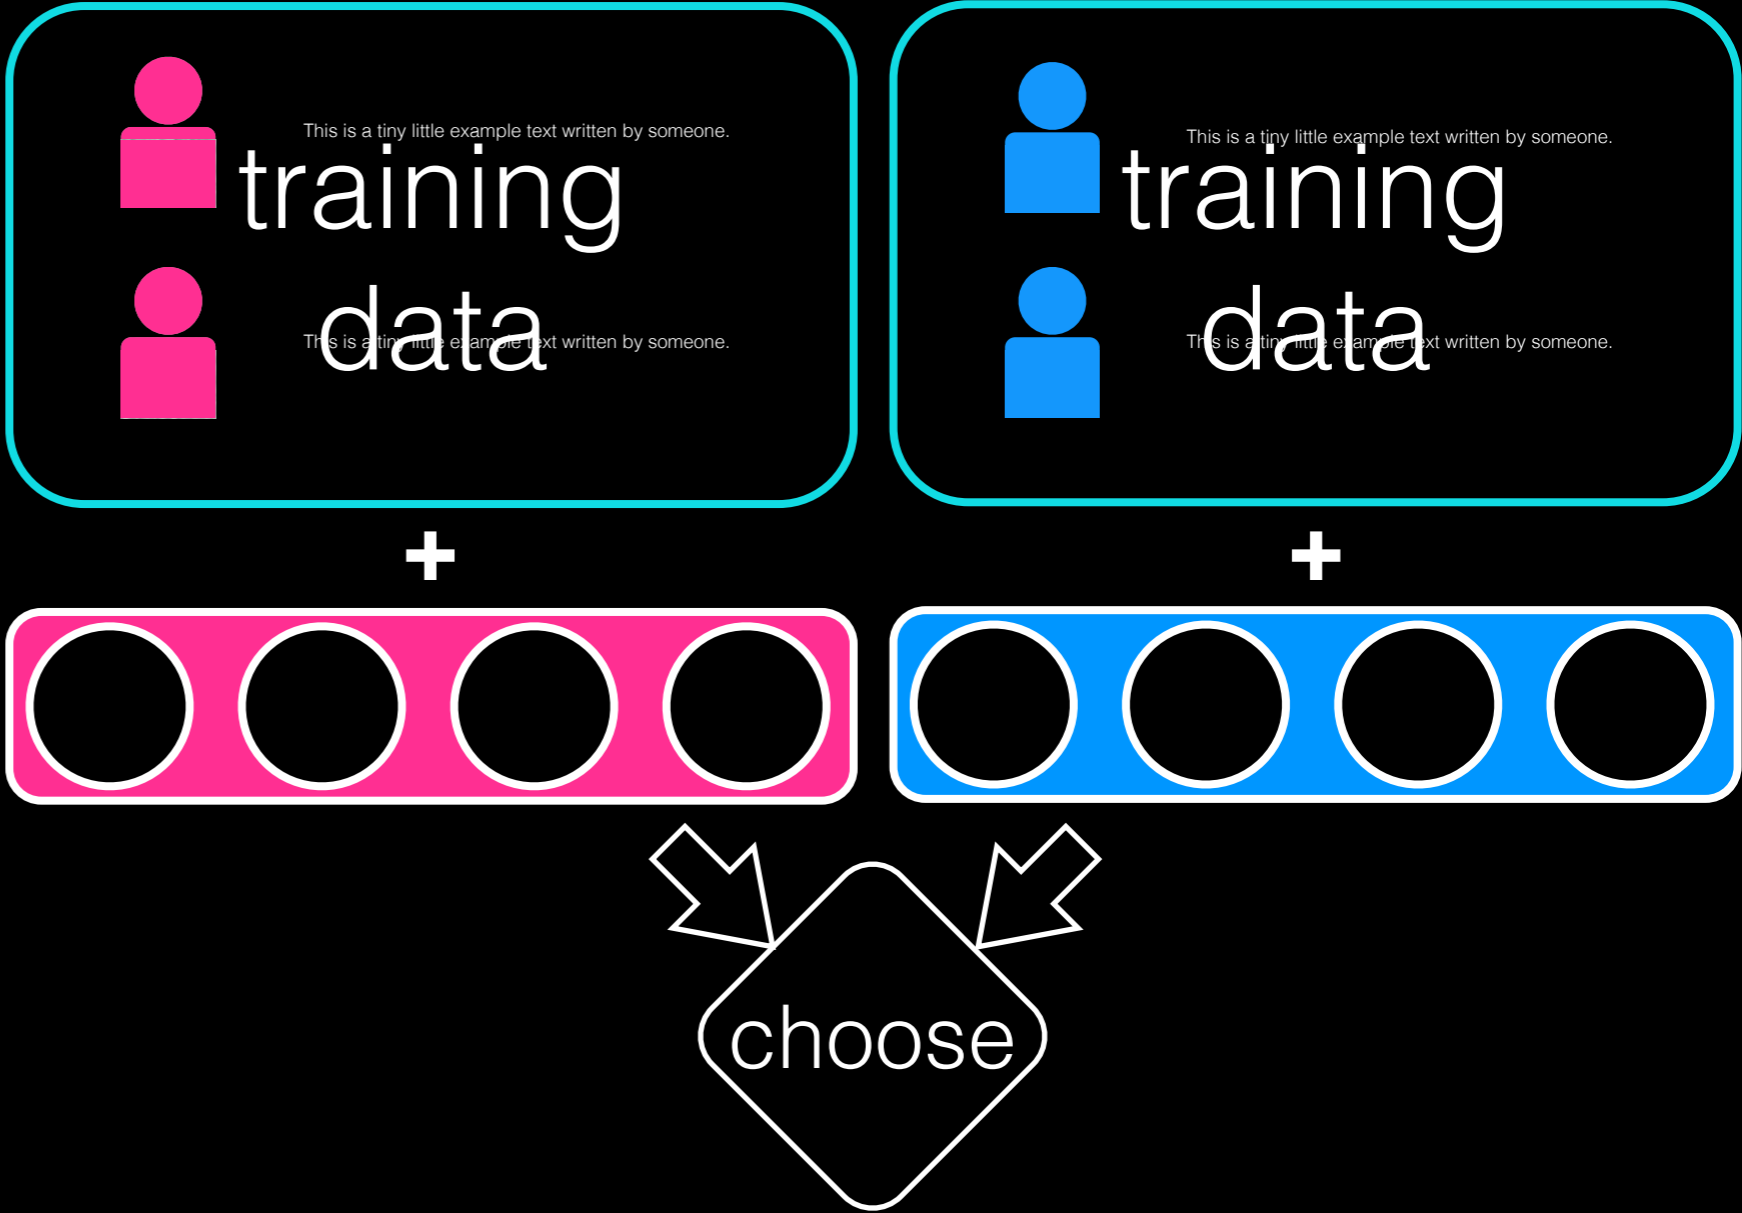


Systems

AGNOSTIC

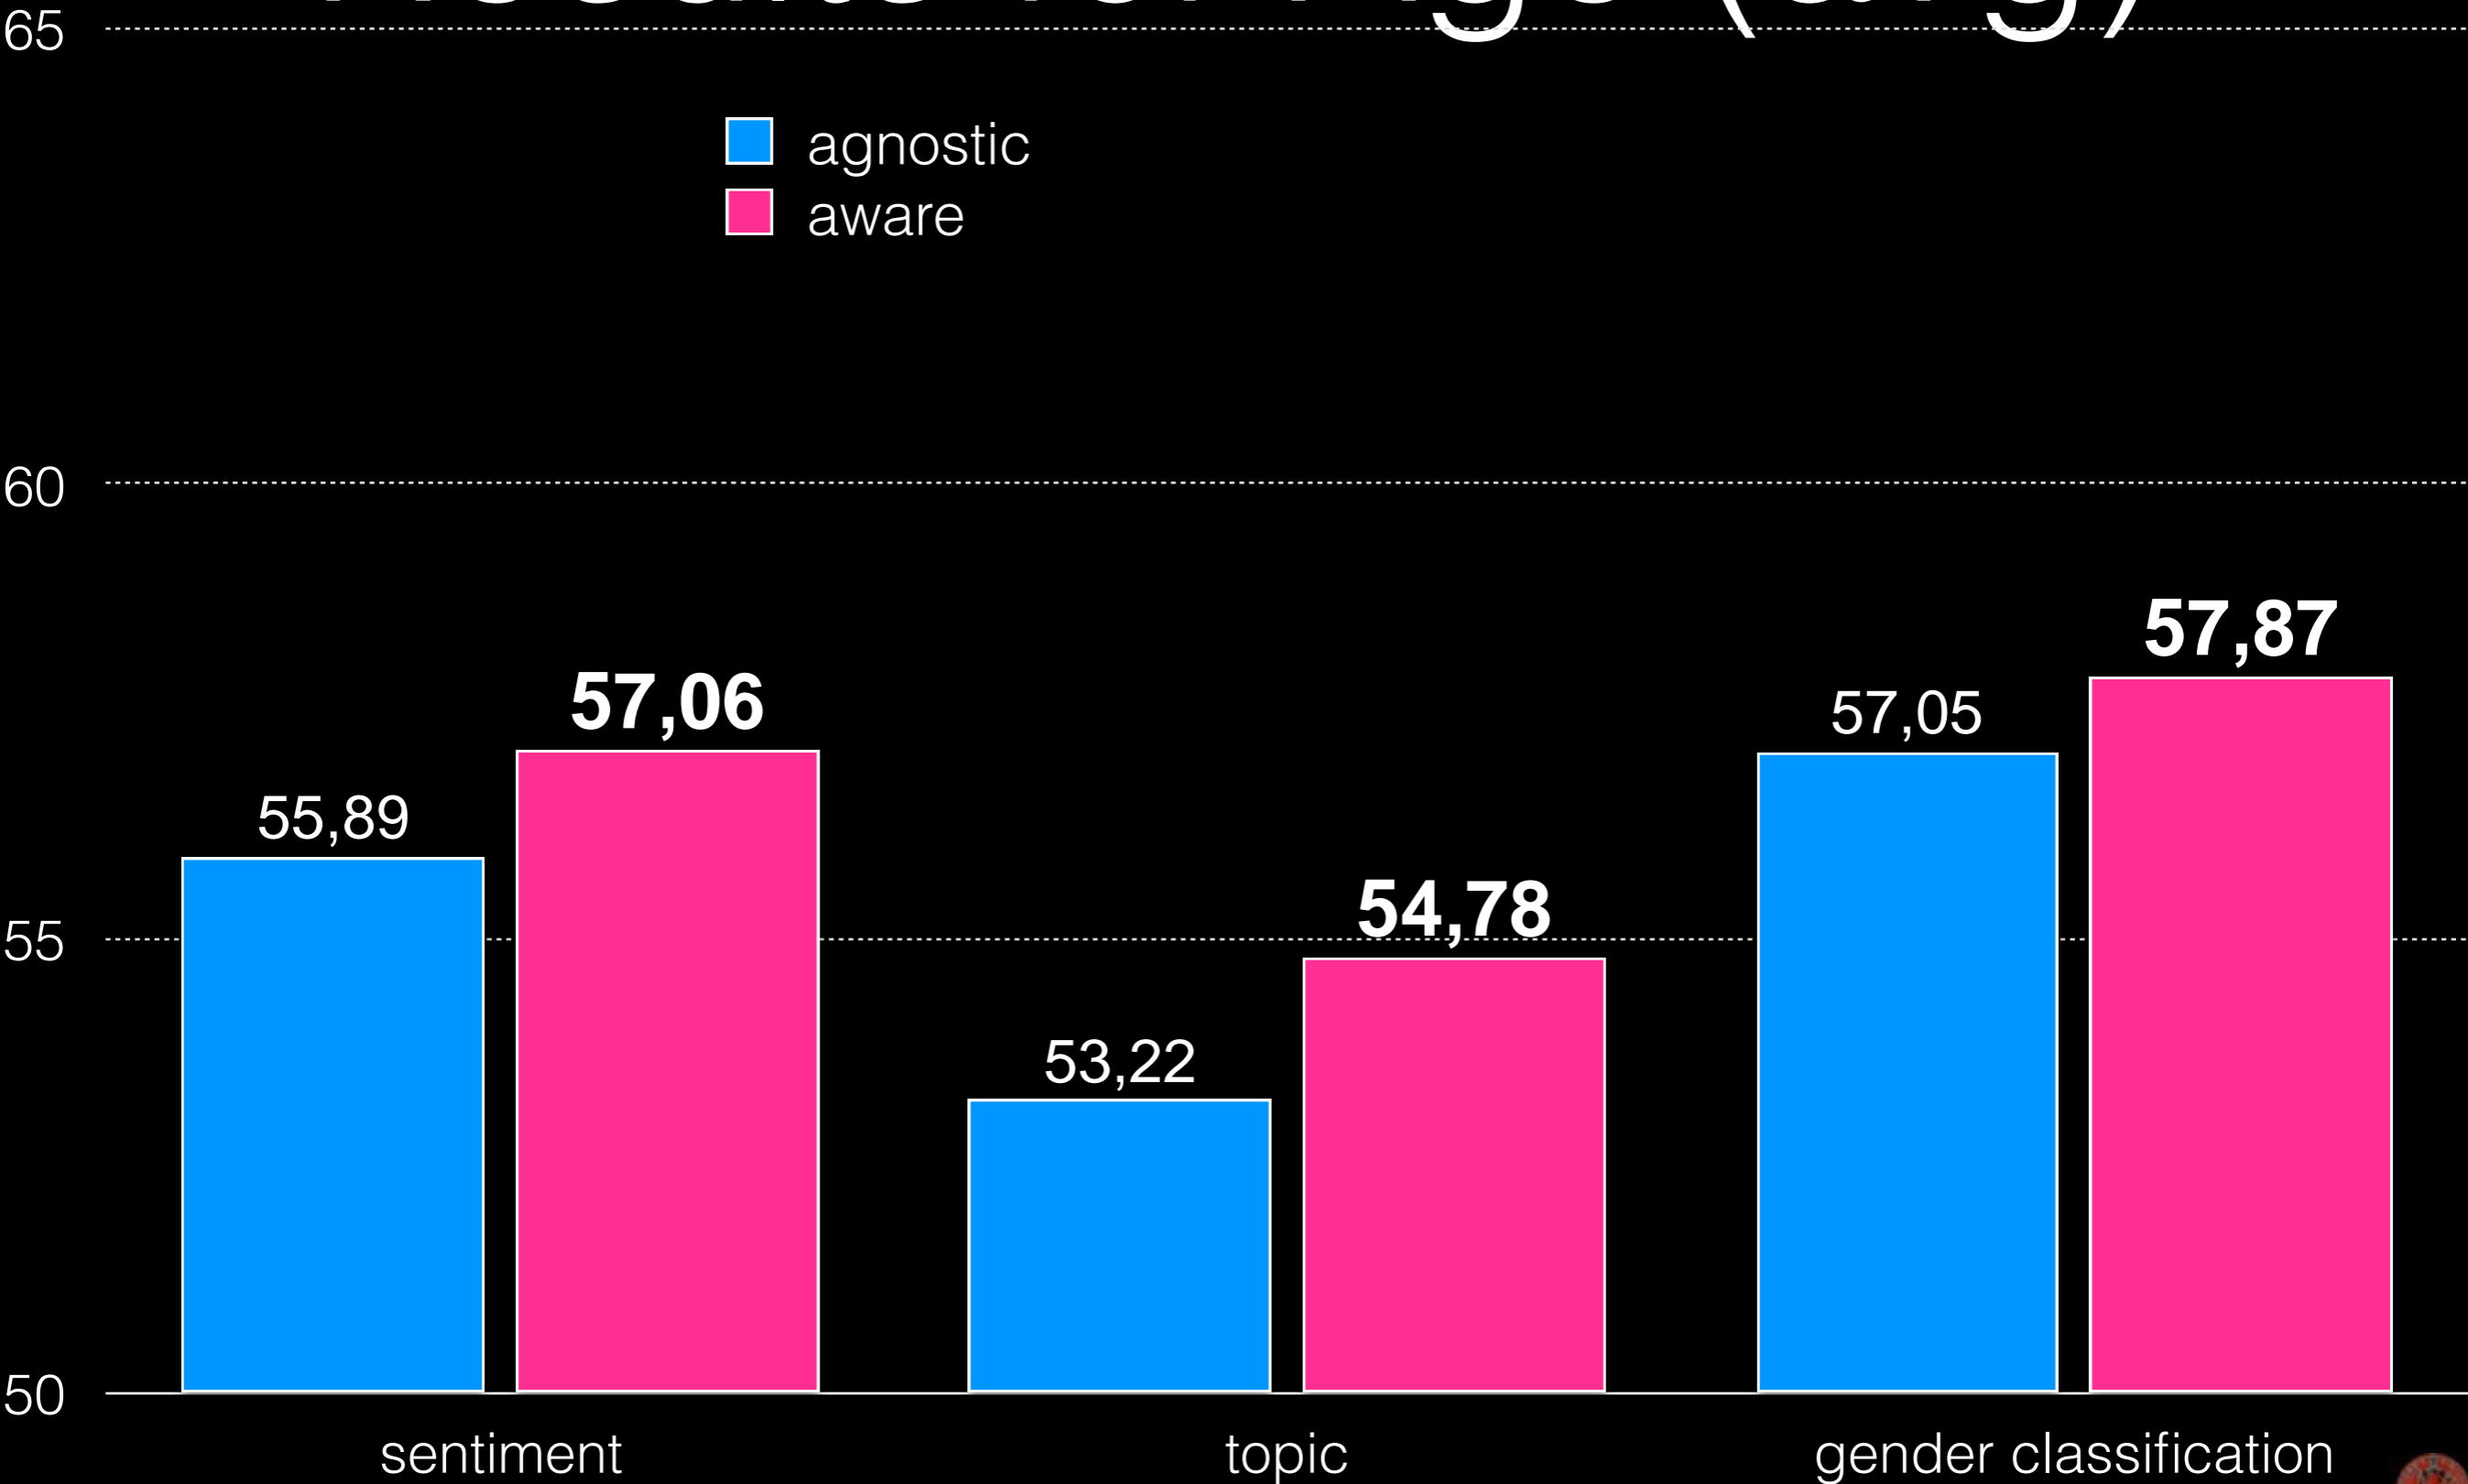


INFORMED



Results for Age (avg)

agnostic
aware



Results for Gender (avg)





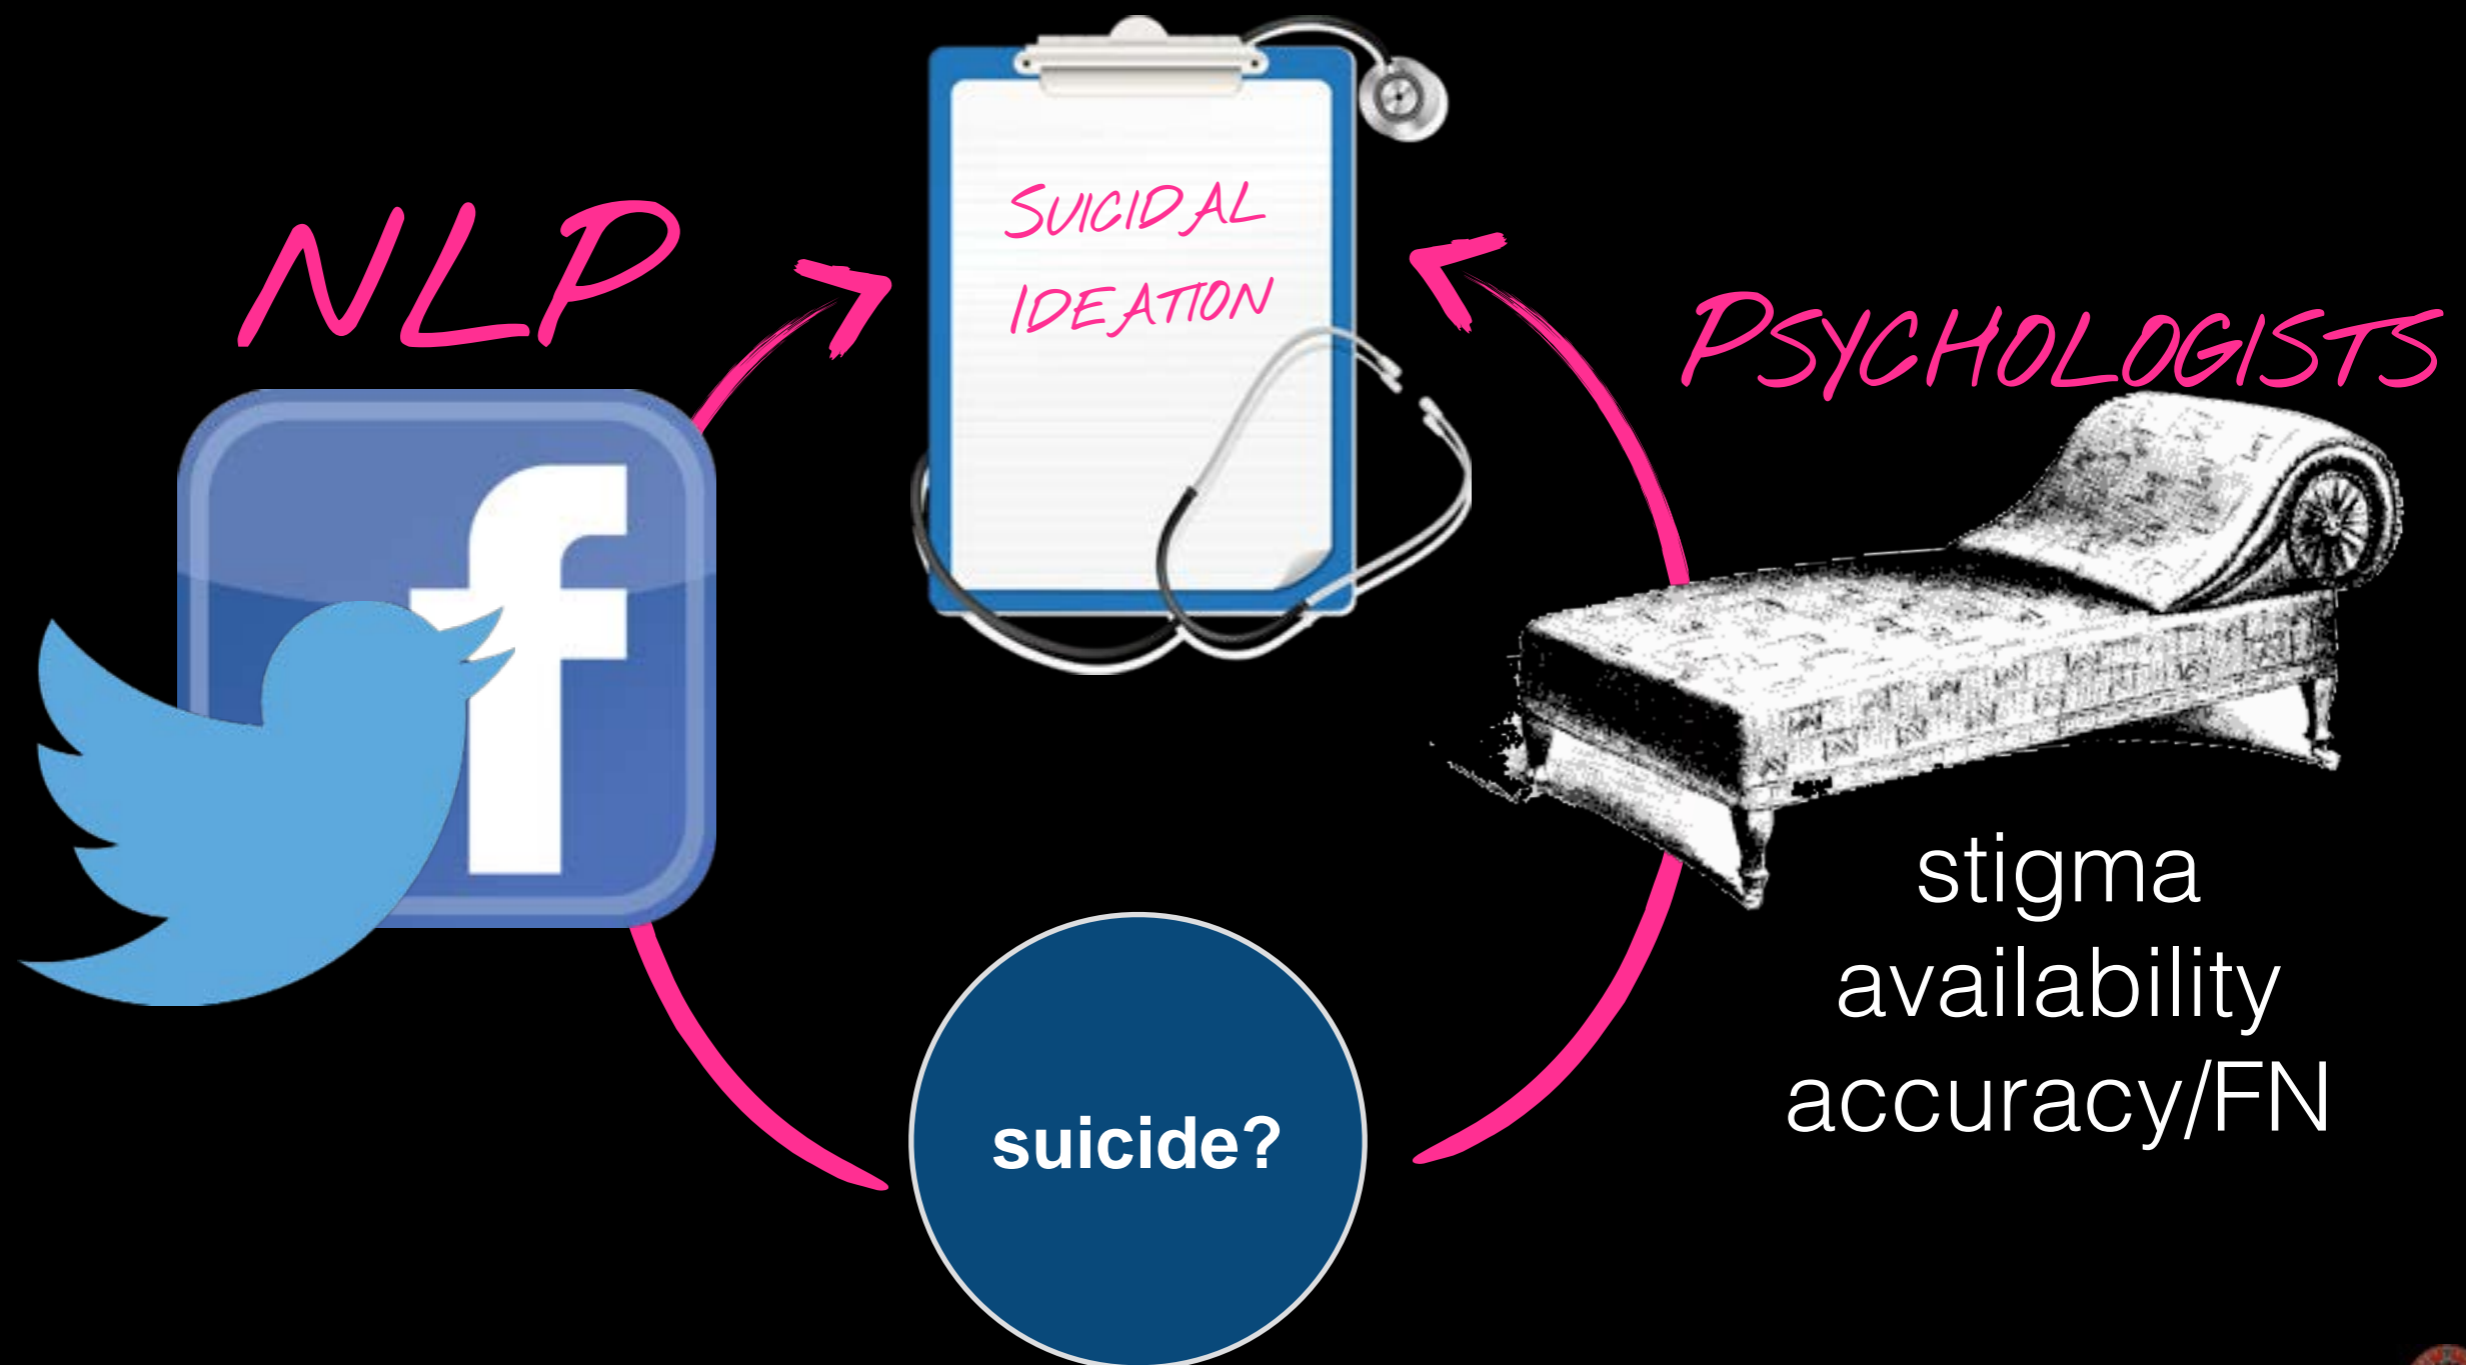
Example 2:



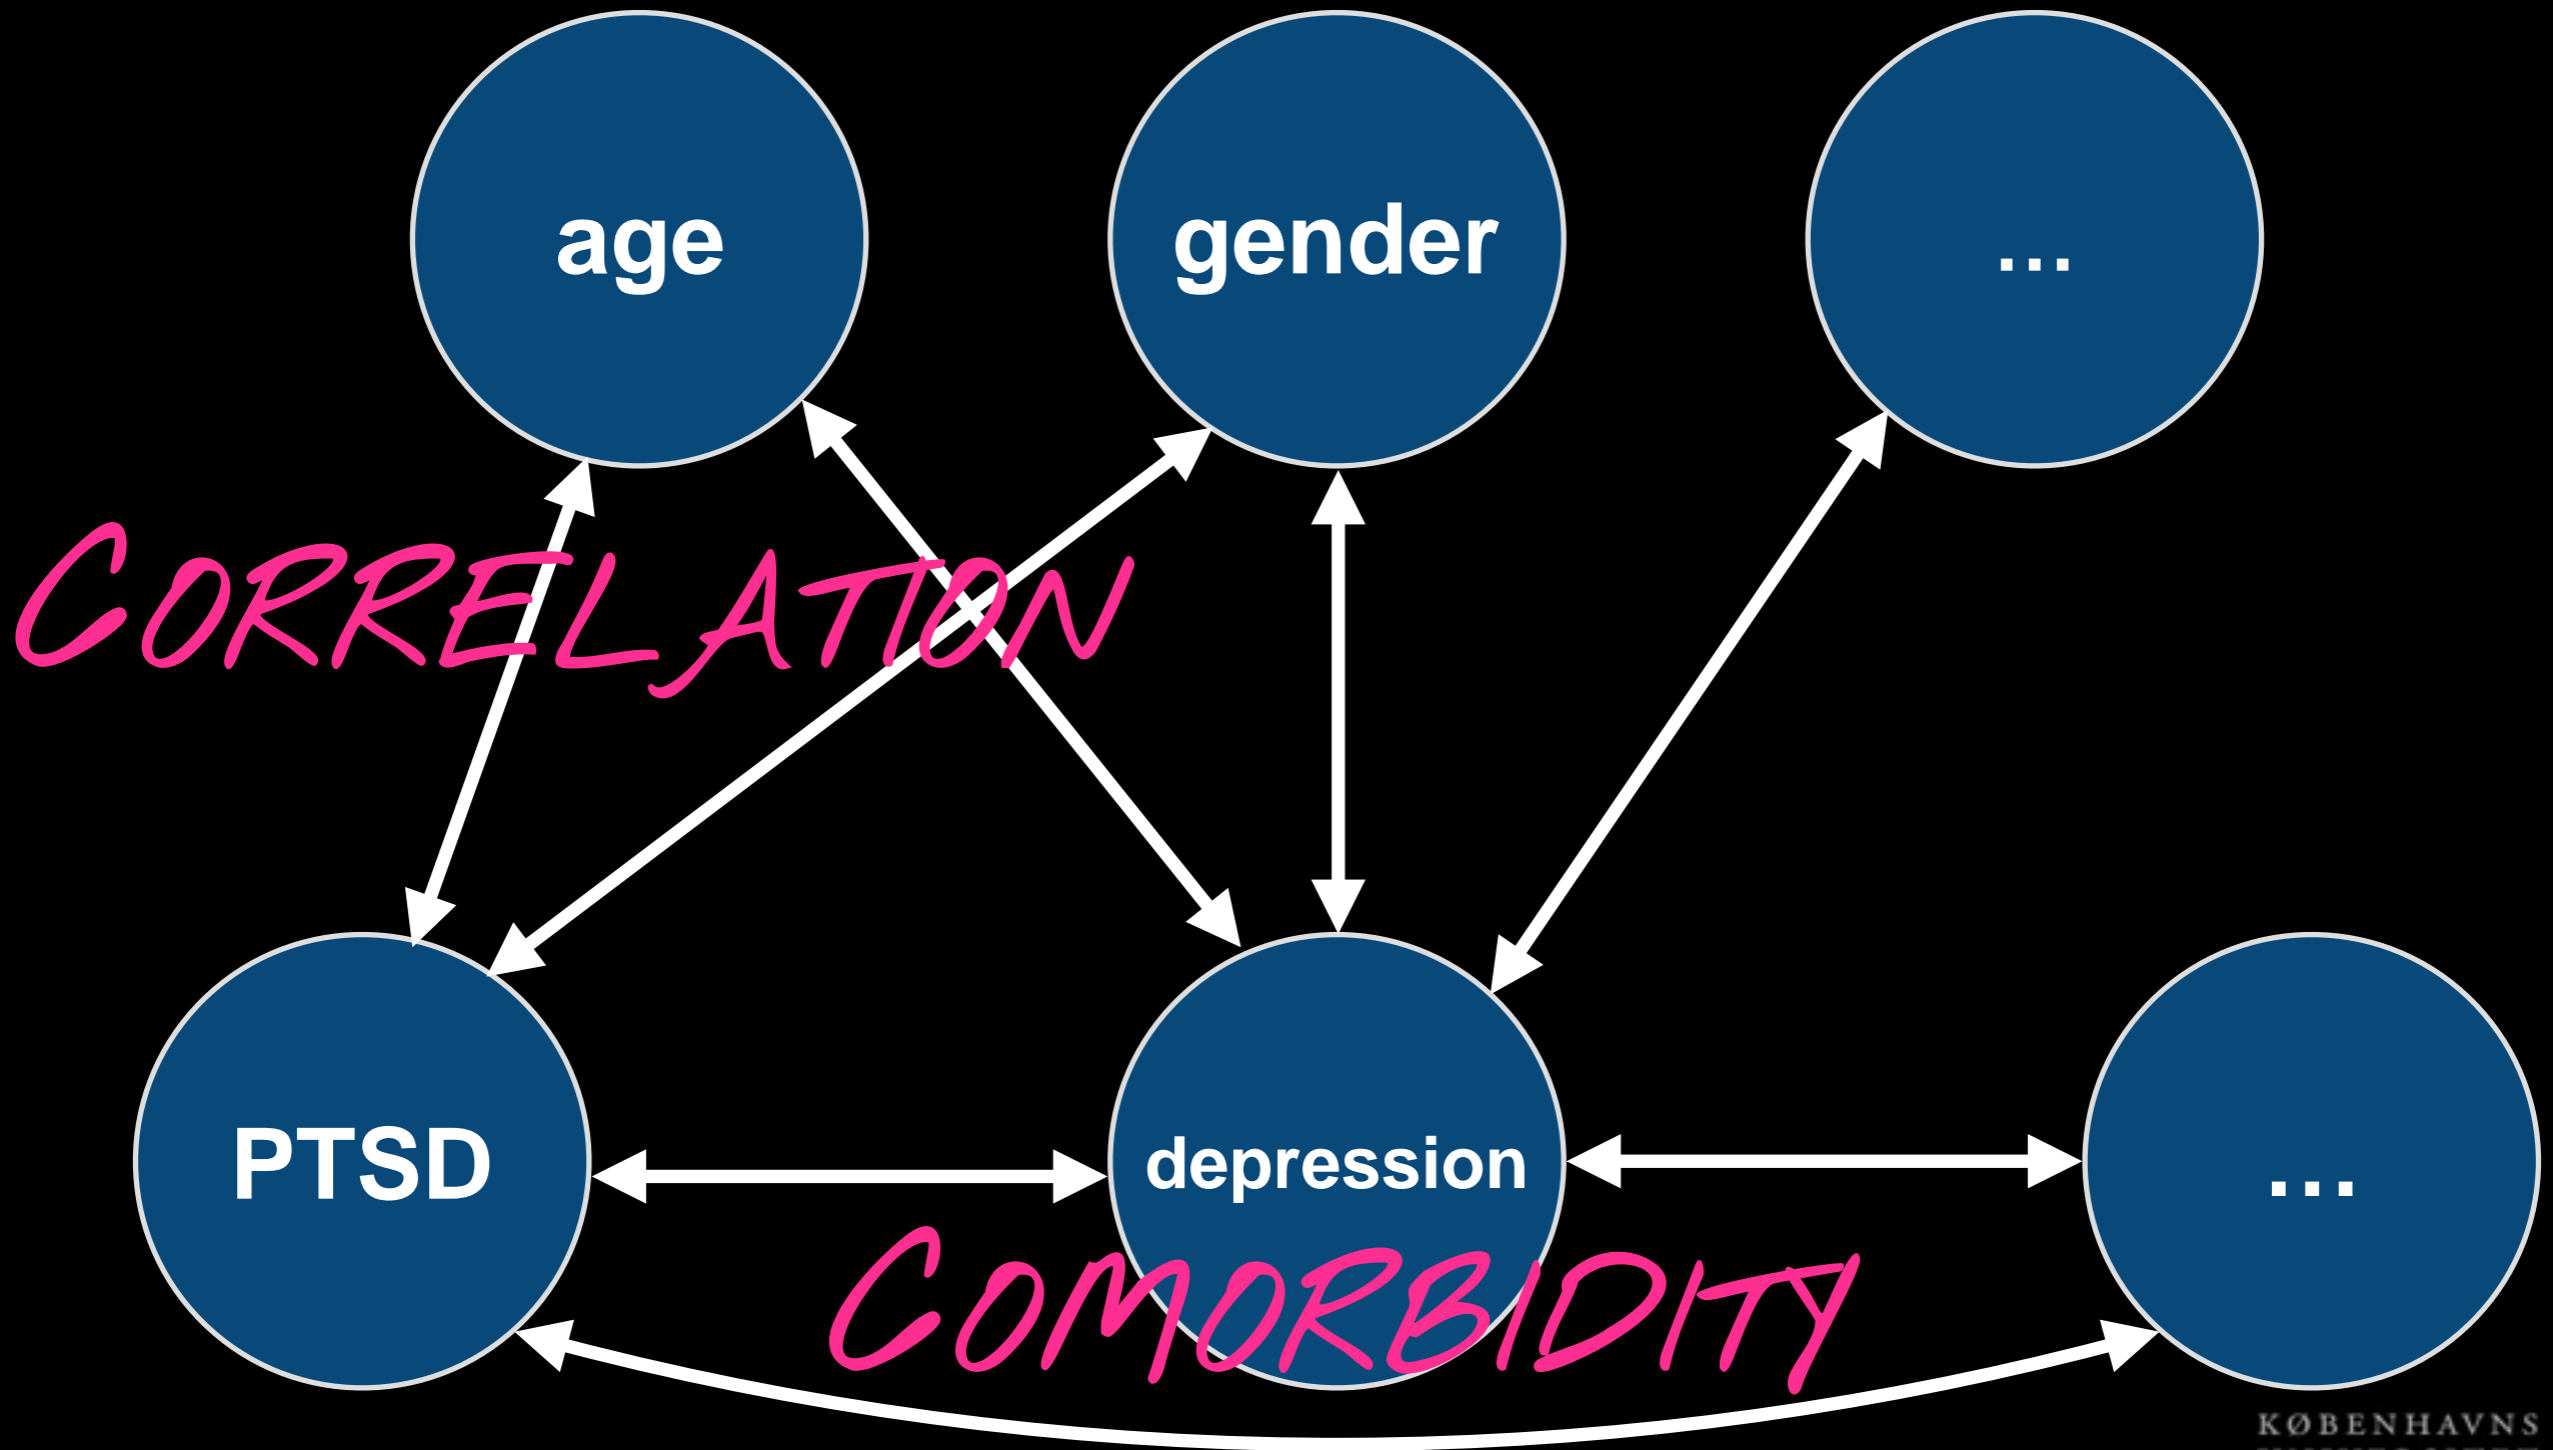
Multitask Learning for Mental Health Conditions



Mental Health Risk Prediction



Comorbidity and Correlation



Motivation from “The Karate Kid”



Single-task
Learners
(STL)



Multitask Learner
(MTL)



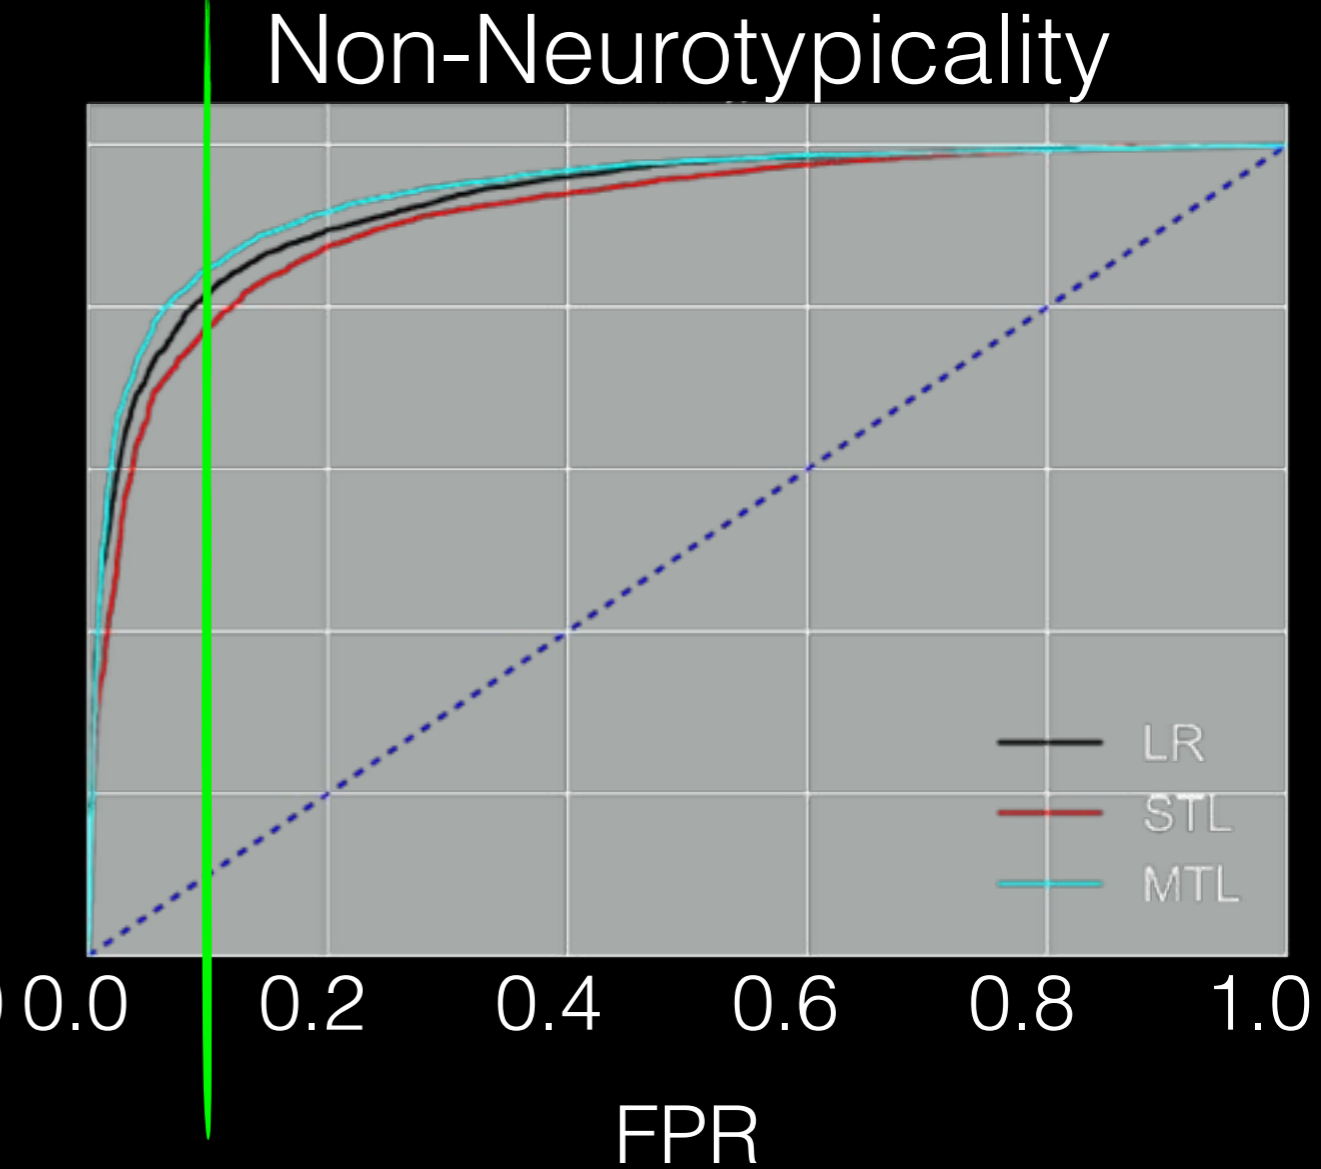
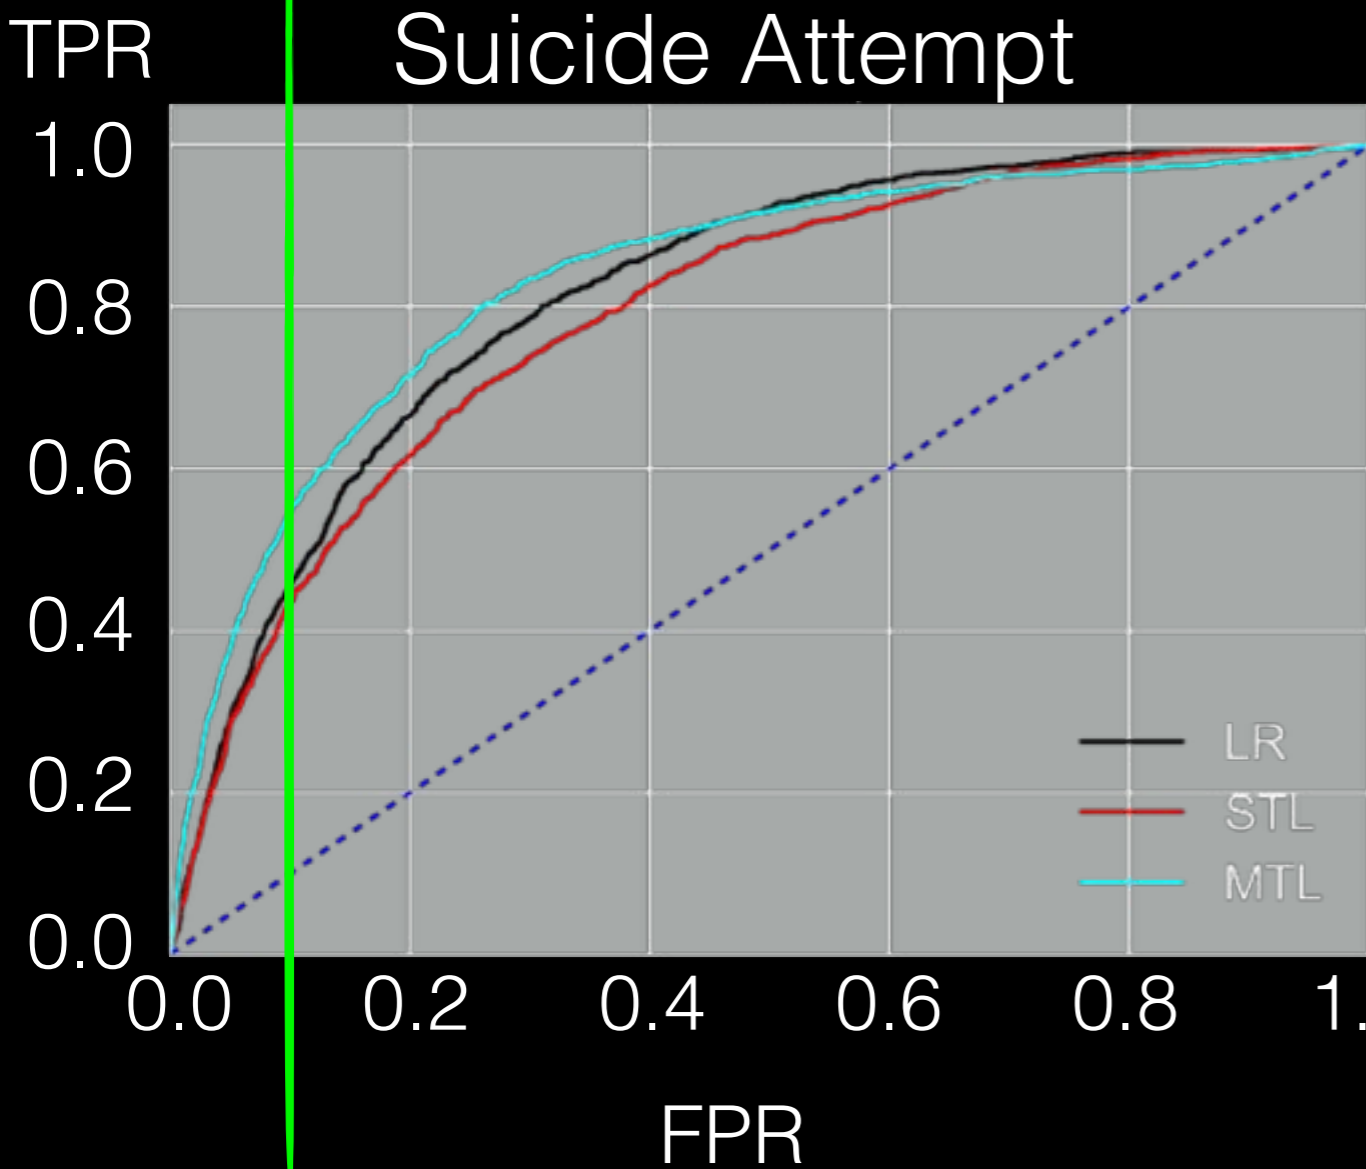
Data

Task	N
Neurotypicality	4791
Anxiety	2407
Depression	1400
Suicide attempt	1208
Eating disorder	749
Schizophrenia	349
Panic disorder	263
PTSD	248
Bipolar disorder	191
All	9611

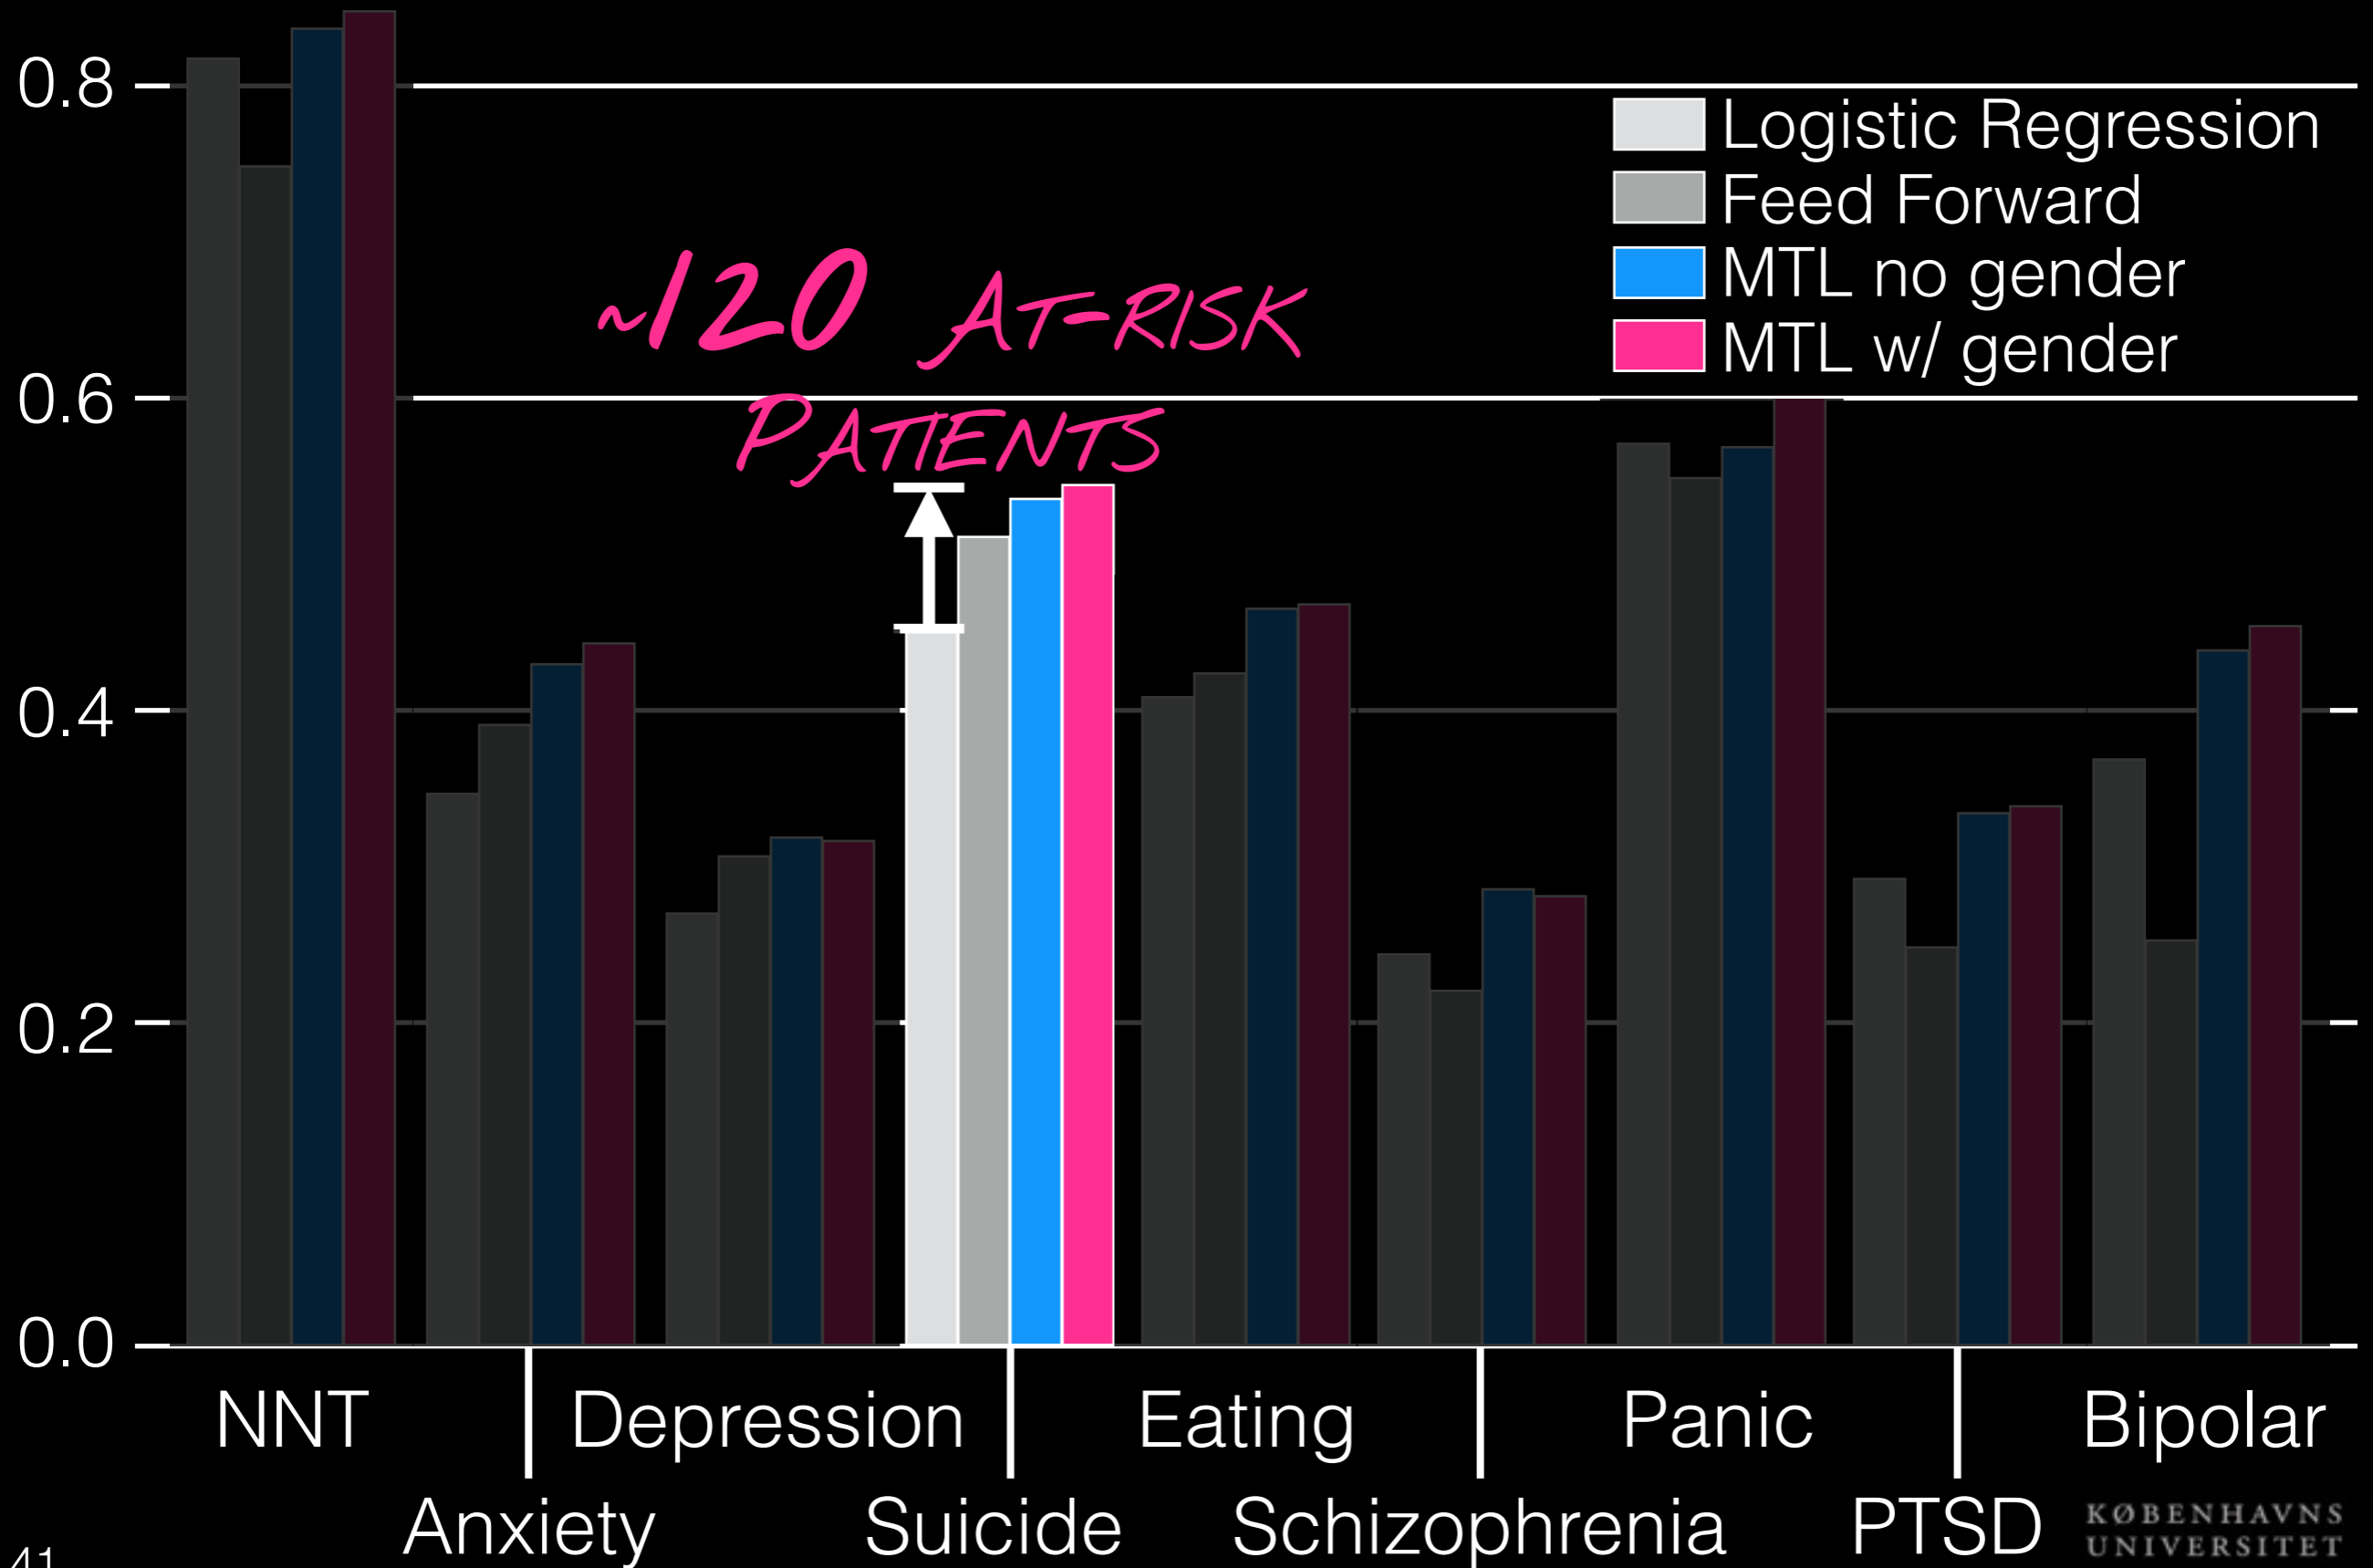
+GENDER

1101

Results ROC



Results: TPR@FPR=0.1



So then all is good, right?

...right?

Challenges



Ethics and NLP



35 male **???** German **???** college education



Exclusion

Exposure

Overgeneralization

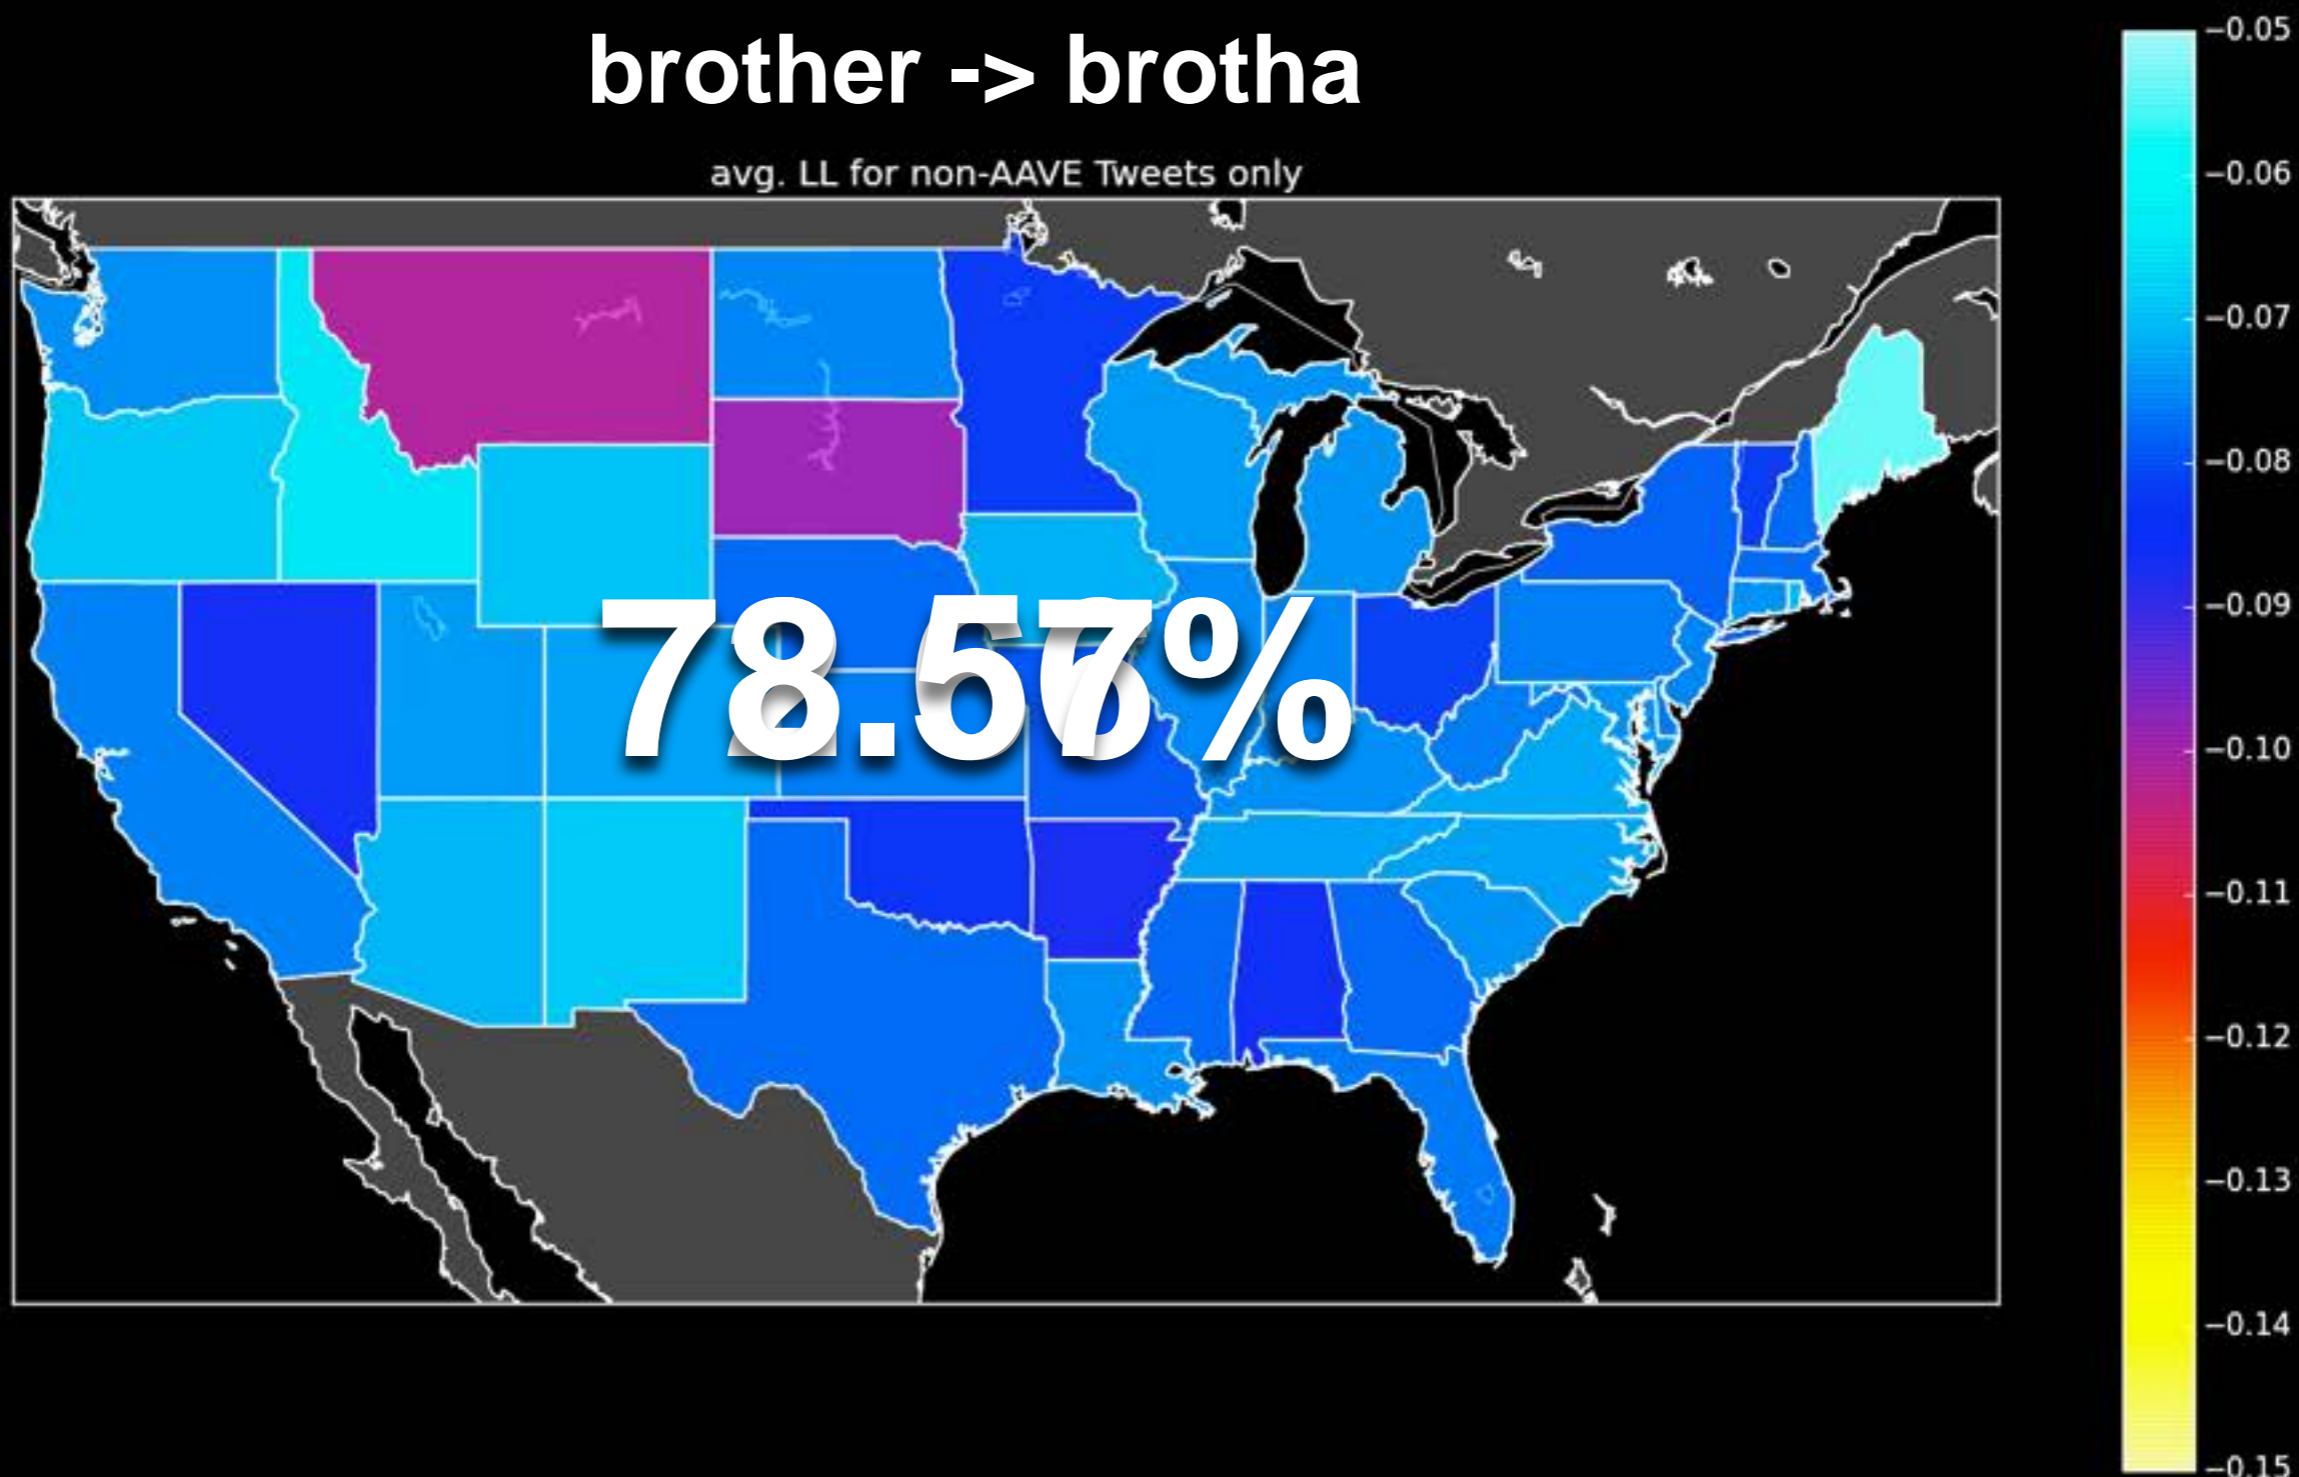
Dual Use



Exclusion

street -> skreet

brother -> brotha



Exclusion

accuracy

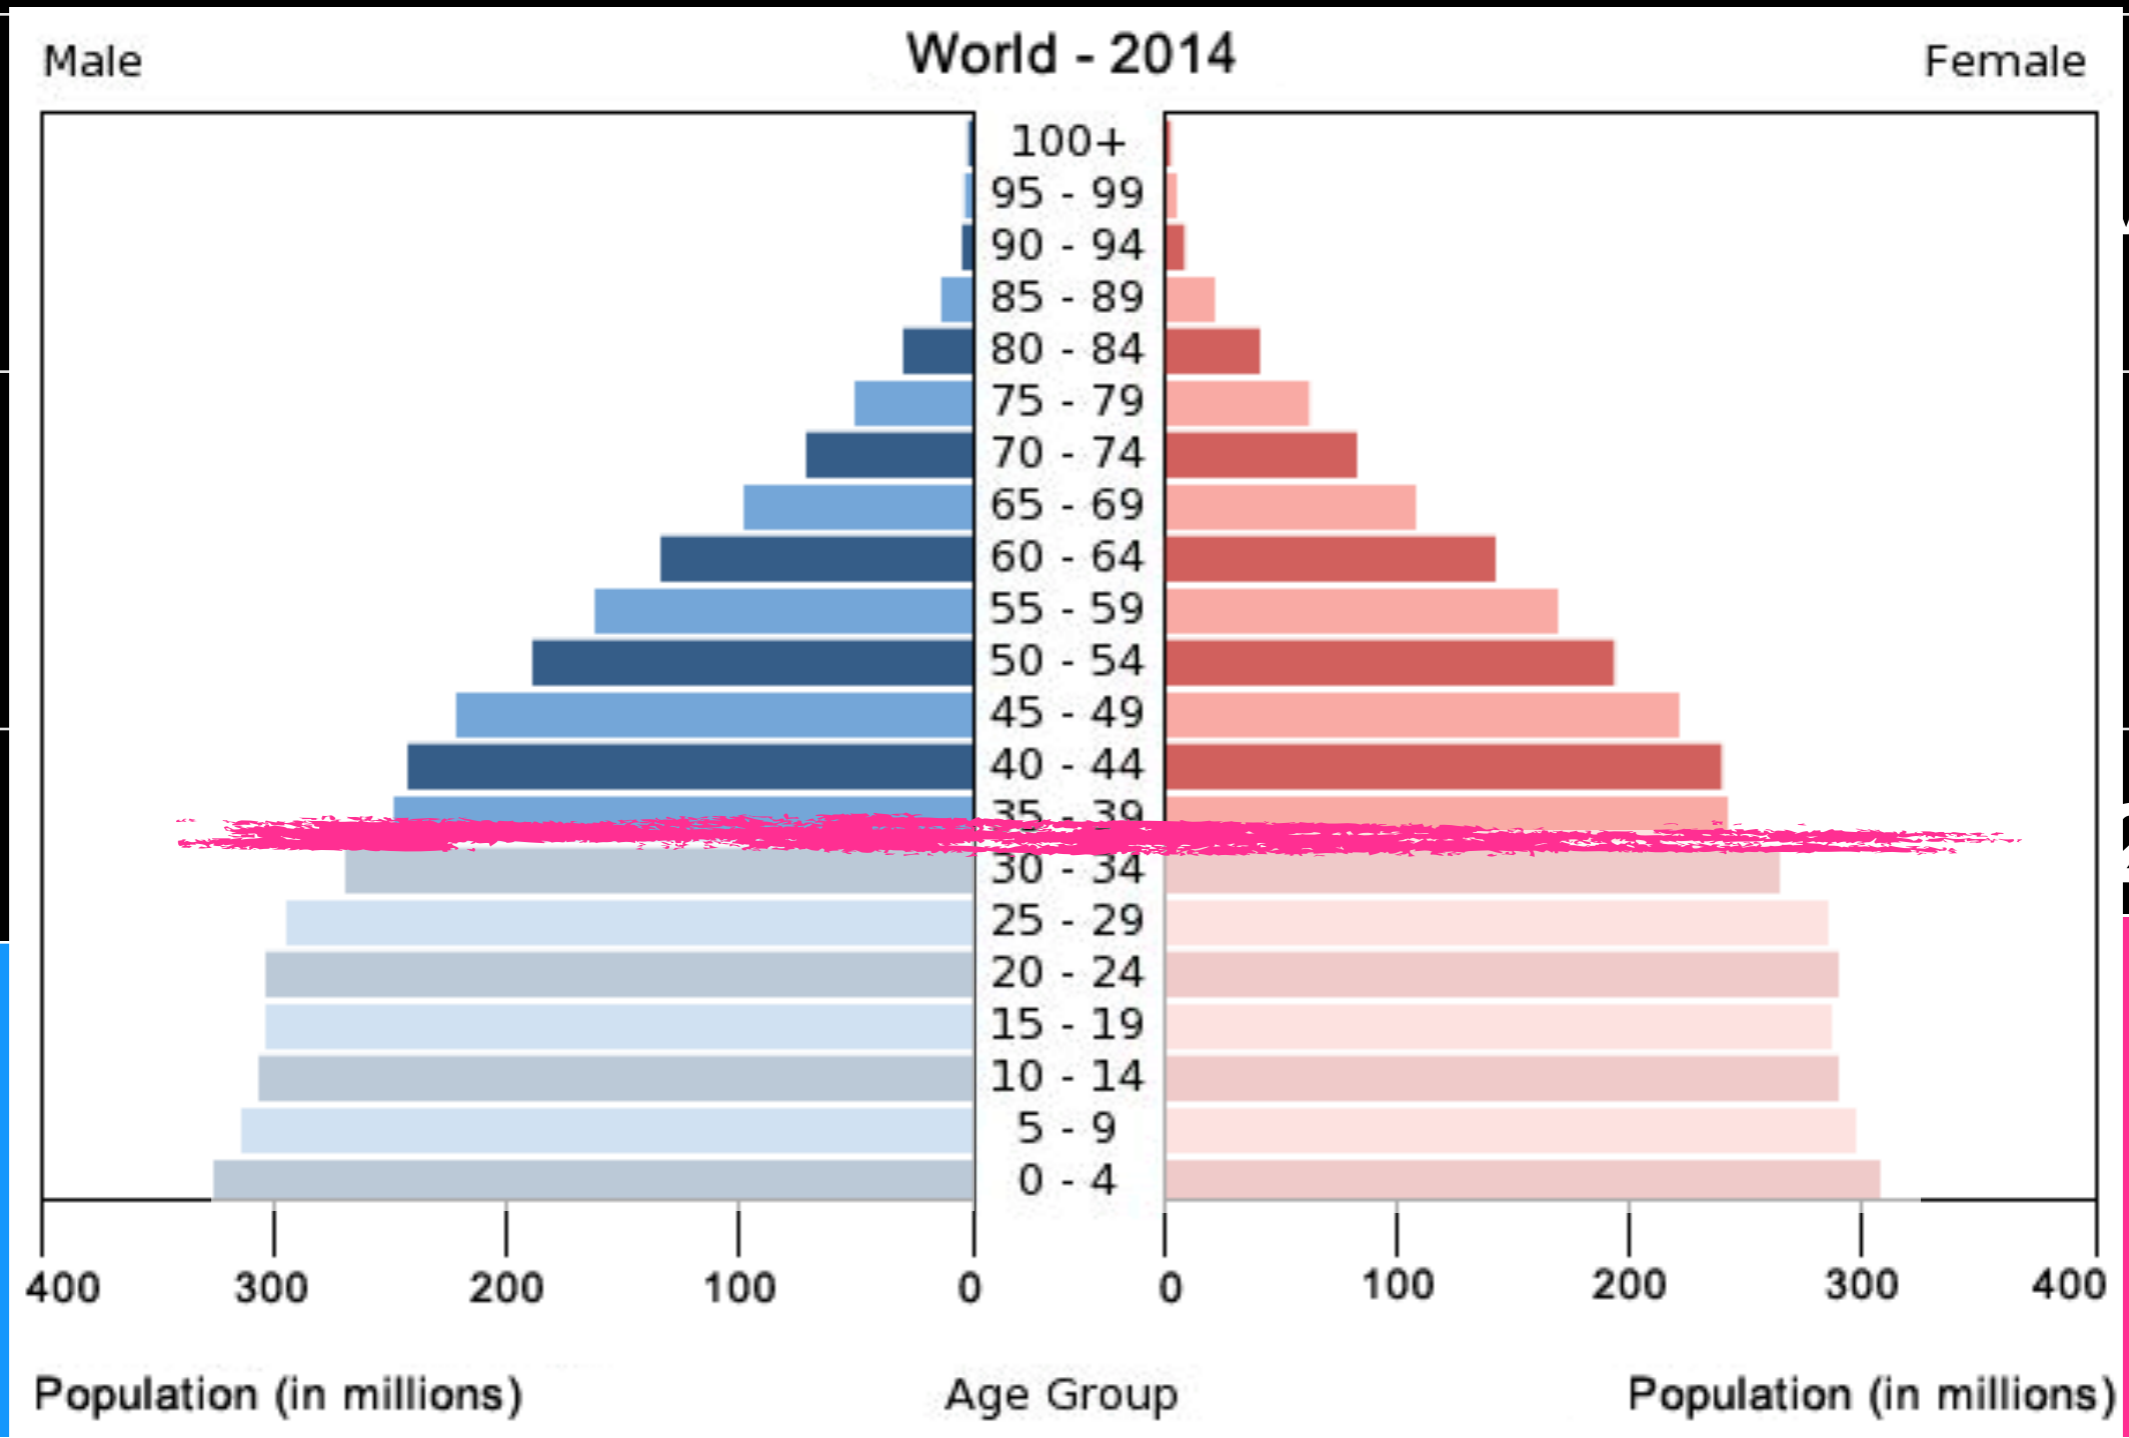
100

95

90

85

80



ing
views
ned

2

O45

U35

O45

U35

Overgeneralization

FALSE POSITIVES

Aug 6 2016

Dear Ms Hovy,

Congratulations on reaching
retirement age!

Also, you're on a no-fly list
because of your political
views and religious beliefs.

The Signal and the Noise



(a) Jewish % Population



(b) Asian % Population



(c) Average House Price



(d) IMD

violent	0.44
gang	0.43
drug	0.42
rob	0.4
danger	0.39
knife	0.39
integration	0.38
black	0.38
boy	0.38
evenly	0.38
dangerous	0.37
stab	0.37

Hate Speech Detection

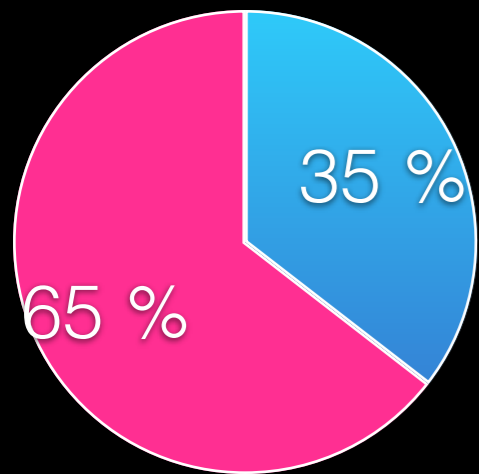
Abusive Language Workshop
ACL Vancouver, Aug 4 2017



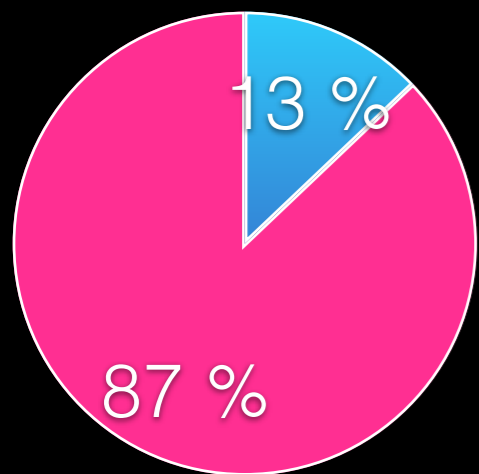
Exposure

Underexposure

- available
- not available



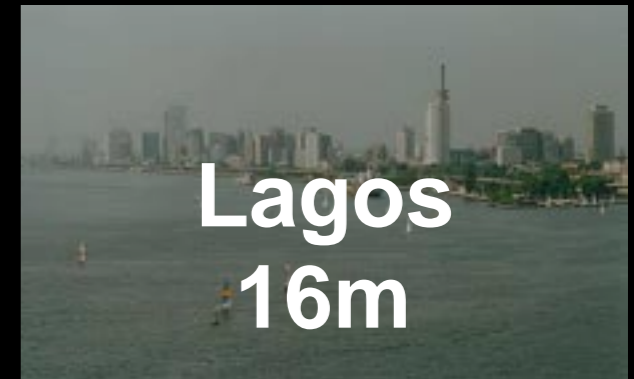
treebanks



semantic
resources

evaluation

Overexposure



**sentiment
analysis**

discourse
parsing

bias



Dual Use

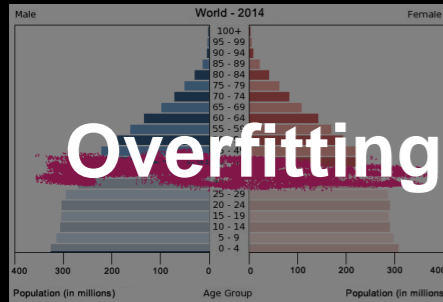
	Pro	Con
authorship attribution	historical documents	dissenter anonymity
text classification	sentiment analysis	censorship
personalization	better user experience	tailored ads

What can we do?

Problem

Source

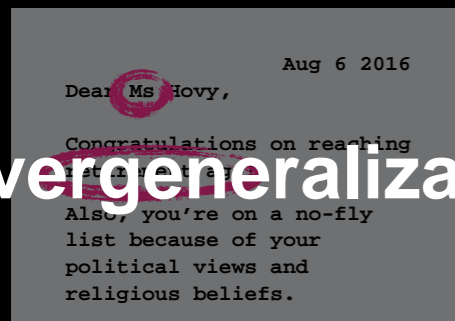
Countermeasures



data
selection

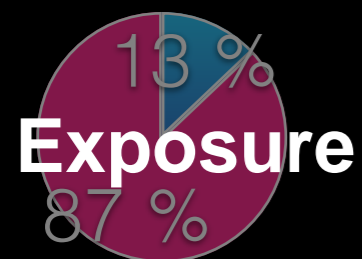
regularization, priors, sampling

Overgeneralization



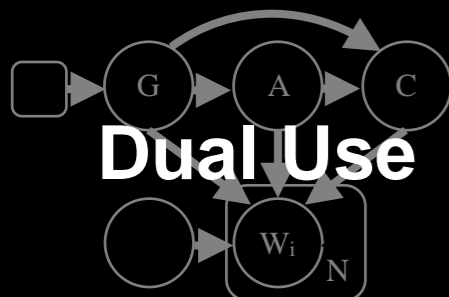
models

dummy labels, error weighting,
confidence thresholds



research
design

consider possible impact



community
goals

educate users, keep discussion
going





HOME

ACCEPTED
PAPERS

INVITED
SPEAKERS

DISCUSSANTS
& PANELISTS

SPONSORS

ORGANIZERS

RELATED
CONFERENCES

ARCHIVED

Slow internet, fast access version is here: <http://www.ethicsinnlp.org/workshop/program.html>

Twitter Hashtag:
#EthNLP
Ethics
in Natural Language Processing
A Workshop at EACL 2017

4.April.2017
Valencia, Spain

Previous
Updates

<http://www.ethicsinnlp.org/>



Wrapping up

Conclusion

- social media + NLP = new methods for CSS
- but: NLP is more than engineering
 - language \geq information, reflects demographics
 - accounting for demographics affects performance
- Work on social media has social impact and responsibility

Open questions

- how do we reflect ever-changing language?
- how do we account for demographic variation?
- how do we ensure fairness?

Thanks!

Questions?

www.dirkhovy.com

 @dirk_hovy

Workshop on NLP and Computational Social Sciences

@ACL

Aug 3, 2017

Vancouver, Canada

<https://sites.google.com/site/nlpandcss/>



@nlpandcss