Hi!
I'm Indira
indira.sen@gesis.org

RWTH AACHEN UNIVERSITY
gesis Leibniz Institute for the Social Sciences

Measuring **Social Constructs** with NLP

Hi!
I’m Indira
indira.sen@gesis.org

RWTH AACHEN UNIVERSITY
gesis
Leibniz Institute for the Social Sciences

In a nutshell

Areas of Interest: Computational Social Science (CSS), Text Analysis

Core Competencies: Applied Machine Learning and Natural Language Processing (NLP)

In a Previous Life

Bachelors [VIT'14] and Masters [IIITD'17] in Computer Science

Master's Thesis: "Detecting Inorganic Likes on Instagram"

Now

Ph.D Student in CSS
[expected 2023]

"Identifying, Characterizing, and Mitigating Errors in the Measurement of Attitudes and Behaviours with Computational Methods"

SOCIAL SCIENCE

Computational Social Science

A field is emerging that leverages the capacity to collect and analyze data at a scale that may reveal patterns of individual and group behaviors.

David Lazer,[1] Alex Pentland,[2] Lada Adamic,[3] Sinan Aral,[2,4] Albert-László Barabási,[5] Devon Brewer,[6] Nicholas Christakis,[1] Noshir Contractor,[7] James Fowler,[8] Myron Gutmann,[3] Tony Jebara,[9] Gary King,[1] Michael Macy,[10] Deb Roy,[2] Marshall Van Alstyne[2,11]

We live life in the network. We check our e-mails regularly, make mobile phone calls from almost any location, swipe transit cards to use public transportation, and make purchases with credit cards. Our movements in public places may be captured by video cameras, and our medical records stored as digital files. We may post blog entries accessible to anyone, or maintain friendships through online social networks. Each of these transactions leaves digital traces that can be compiled into comprehensive pictures of both individual and group behavior, with the potential to transform our understanding of our lives, organizations, and societies.

The capacity to collect and analyze massive amounts of data has transformed such fields as biology and physics. But the emergence of a data-driven "computational social science" has been much slower. Leading journals in economics, sociology, and political science show little evidence of this field. But computational social science is occurring—in Internet companies such as Google and Yahoo, and in government agencies such as the U.S. National Security Agency. Computational social science could become the exclusive domain of private companies and government agencies. Alternatively, there might emerge a privileged set of academic researchers presiding over private data from which they produce papers that cannot be critiqued or replicated. Neither scenario will serve the long-term public interest of accumulating, verifying, and disseminating knowledge.

What value might a computational social science—based in an open academic environment—offer society, by enhancing understanding of individuals and collectives? What are the

[1]Harvard University, Cambridge, MA, USA. [2]Massachusetts Institute of Technology, Cambridge, MA, USA. [3]University of Michigan, Ann Arbor, MI, USA. [4]New York University, New York, NY, USA. [5]Northeastern University, Boston, MA, USA. [6]Interdisciplinary Scientific Research, Seattle, WA, USA. [7]Northwestern University, Evanston, IL, USA. [8]University of California–San Diego, La Jolla, CA, USA. [9]Columbia University, New York, NY, USA [10]Cornell University, Ithaca, NY, USA. [11]Boston University, Boston, MA, USA. E-mail: david_lazer@harvard.edu. Complete affiliations are listed in the supporting online material.

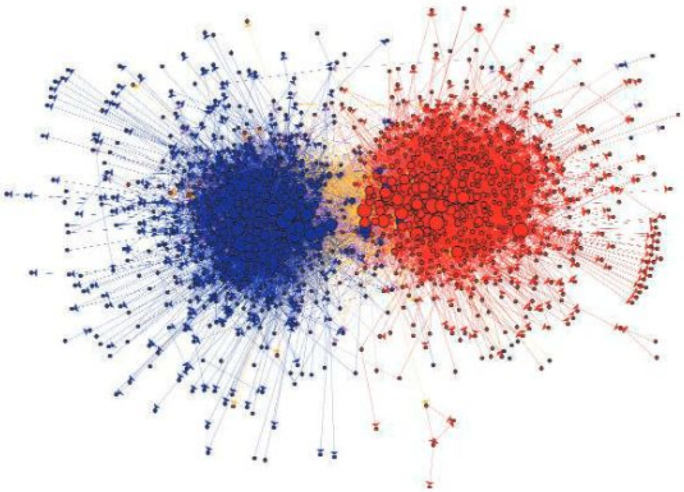

Data from the blogosphere. Shown is a link structure within a community of political blogs (from 2004), where red nodes indicate conservative blogs, and blue liberal. Orange links go from liberal to conservative, and purple ones from conservative to liberal. The size of each blog reflects the number of other blogs that link to it. [Reproduced from (*8*) with permission from the Association for Computing Machinery]

SOCIAL SCIENCE

Computational Social Science

A field is emerging that leverages the capacity to collect and analyze data at a scale that may reveal patterns of individual and group behaviors.

David Lazer,[1] Alex Pentland,[2] Lada Adamic,[3] Sinan Aral,[2,4] Albert-László Barabási,[5] Devon Brewer,[6] Nicholas Christakis,[1] Noshir Contractor,[7] James Fowler,[8] Myron Gutmann,[3] Tony Jebara,[9] Gary King,[1] Michael Macy,[10] Deb Roy,[2] Marshall Van Alstyne[2,11]

We live life in the network. We check our e-mails regularly, make mobile phone calls from almost any location, swipe transit cards to use public transportation, and make purchases with credit cards. Our movements in public places may be captured by video cameras, and our medical records stored as digital files. We may post blog entries accessible to anyone, or maintain friendships through online social networks. Each of these transactions leaves digital traces that can be compiled into comprehensive pictures of both individual and group behavior, with the potential to transform our understanding of our lives, organizations, and societies.

The capacity to collect and analyze massive amounts of data has transformed such fields as biology and physics. But the emergence of a data-driven "computational social science" has been much slower. Leading journals in economics, sociology, and political science show little evidence of this field. But computational social science is occurring—in Internet companies such as Google and Yahoo, and in govern-

[1]Harvard University, Cambridge, MA, USA. [2]Massachusetts Institute of Technology, Cambridge, MA, USA. [3]University of Michigan, Ann Arbor, MI, USA. [4]New York University, New York, NY, USA. [5]Northeastern University, Boston, MA, USA. [6]Interdisciplinary Scientific Research, Seattle, WA, USA. [7]Northwestern University, Evanston, IL, USA. [8]University of California–San Diego, La Jolla, CA, USA. [9]Columbia University, New York, NY, USA [10]Cornell University, Ithaca, NY, USA. [11]Boston University, Boston, MA, USA. E-mail: david_lazer@harvard.edu. Complete affiliations are listed in the supporting online material.

Data from the blogosphere. Shown is a link structure within a community of political blogs (from 2004), where red nodes indicate conservative blogs, and blue liberal. Orange links go from liberal to conservative, and purple ones from conservative to liberal. The size of each blog reflects the number of other blogs that link to it. [Reproduced from (8) with permission from the Association for Computing Machinery]

NATURE|Vol 445|1 February 2007

ESSAY

Putting the pieces together

A twenty-first century science

If handled appropriately, data about Internet-based communication and interactivity could revolutionize our understanding of collective human behaviour.

Duncan J. Watts

Few would deny that many of the major problems currently facing humanity are social and economic in nature. From the apparent wave of religious fundamentalism sweeping the Islamic world (and parts of the Western world), to collective economic security, global warming and the great epidemics of our times, powerful yet mysterious social forces come into play.

But few readers of *Nature* would consider social science to be the science of the twenty-first century. Although economics, sociology, political science and anthropology have produced a plethora of findings regarding human social behaviour, they self-reports from participants, which suffer from cognitive biases, errors of perception and framing ambiguities.

The striking proliferation over the past decade of Internet-based communication and interactivity, however, is beginning to lift these constraints. For the first time, we can begin to observe the real-time interactions of millions of people at a resolution that is sensitive to effects at the level of the individual. Meanwhile, ever-faster computers permit us to simulate large networks of social interactions. The result has been tremendous interest in social networks: thousands of papers and a growing number of books have been published in less than a decade, leading some to herald framework of collective social dynamics. People do not just interact: their interactions have consequences for the choices they, and others, make.

Studies that combine all these features are currently beyond the state of the art, but two of my group's recent projects indicate tentative progress. The first used the anonymized e-mail logs of a university community of around 40,000 people to track daily network evolution over a year as a function of existing network structure, shared activities (such as classes) and individual attributes. Dynamic data of this type may shed light on the relative roles of structural constraints and individual preferences in determining, for example, observed

SOCIAL SCIENCE

Computational Social Science

David Lazer,[1] Alex Pentland,[2] Lada Adamic,[3] Sinan Aral,[2,4] Albert-László Barabási,[5] Devon Brewer,[6] Nicholas Christakis,[1] Noshir Contractor,[7] James Fowler,[8] Myron Gutmann,[3] Tony Jebara,[9] Gary King,[1] Michael Macy,[10] Deb Roy,[2] Marshall Van Alstyne[2,11]

A field is emerging that leverages the capacity to collect and analyze data at a scale that may reveal patterns of individual and group behaviors.

We live life in the network. We check our e-mails regularly, make mobile phone calls from almost any location, swipe transit cards to use public transportation, and make purchases with credit cards. Our movements in public places may be captured by video cameras, and our medical records stored as digital files. We may post blog entries accessible to anyone, or maintain friendships through online social networks. Each of these transactions leaves digital traces that can be compiled into comprehensive pictures of both individual and group behavior, with the potential to transform our understanding of our lives, organizations, and societies.

The capacity to collect and analyze massive amounts of data has transformed such fields as biology and physics. But the emergence of a data-driven "computational social science" has been much slower. Leading journals in economics, sociology, and political science show little evidence of this field. But computational social science is occurring—in Internet companies such as Google and Yahoo, and in govern-

[1]Harvard University, Cambridge, MA, USA. [2]Massachusetts Institute of Technology, Cambridge, MA, USA. [3]University of Michigan, Ann Arbor, MI, USA. [4]New York University, New York, NY, USA. [5]Northeastern University, Boston, MA, USA. [6]Interdisciplinary Scientific Research, Seattle, WA, USA. [7]Northwestern University, Evanston, IL, USA. [8]University of California–San Diego, La Jolla, CA, USA. [9]Columbia University, New York, NY, USA [10]Cornell University, Ithaca, NY, USA. [11]Boston University, Boston, MA, USA. E-mail: david_lazer@harvard.edu. Complete affiliations are listed in the supporting online material.

Data from the blogosphere. Shown is a link structure within a community of political blogs (from 2004), where red nodes indicate conservative blogs, and blue liberal. Orange links go from liberal to conservative, and purple ones from conservative to liberal. The size of each blog reflects the number of other blogs that link to it. [Reproduced from (8) with permission from the Association for Computing Machinery]

NATURE|Vol 445|1 February 2007

ESSAY

Putting the pieces together

A twenty-first century science

If handled appropriately, data about Internet-based communication and interactivity could revolutionize our understanding of collective human behaviour.

Duncan J. Watts

great ep... mics of our times, powerful yet mysterious social forces come into play.

But few readers of *Nature* would consider social science to be the science of the twenty-first century. Although economics, sociology, political science and anthropology have produced a plethora of findings regarding human social behaviour, they

self-reports from participants, which suffer from cognitive biases, errors of perception and framing ambiguities. ... of the individual. Meanwhile, ever-faster computers permit us to simulate large networks of social interactions. The result has been tremendous interest in social networks: thousands of papers and a growing number of books have been published in less than a decade, leading some to herald

framework of collective social dynamics. People do not just interact: their interactions have consequences for the choices they, and others, make.

Studies that combine all these features are currently beyond the state of the art, but two of my group's recent projects indicate tentative progress. The first used the anonymized e-mail logs of a university community of around 40,000 people to track daily network evolution over a year as a function of existing network structure, shared activities (such as classes) and individual attributes. Dynamic data of this type may shed light on the relative roles of structural constraints and individual preferences in determining, for example, observed

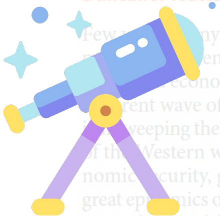

Digital Traces + computational methods, incl. NLP, are the ***Telescope*** for the Social Sciences

Inferences from Digital Traces: Challenges

- [Tufekci'14] on the representative and methodological issues of using Twitter
- Social Data Biases and Pitfalls [Olteanu'20]
- Call for a Measurement Theory for Digital Trace Data [Jungherr'18]

Inferences from Digital Traces: Challenges

- ❖ Issues with the use of 'AI' or computational methods
  - ➢ black-box, intransparent
  - ➢ perpetuate social biases

Shah et al., '[Predictive Biases in Natural Language Processing Models: A Conceptual Framework and Overview]'

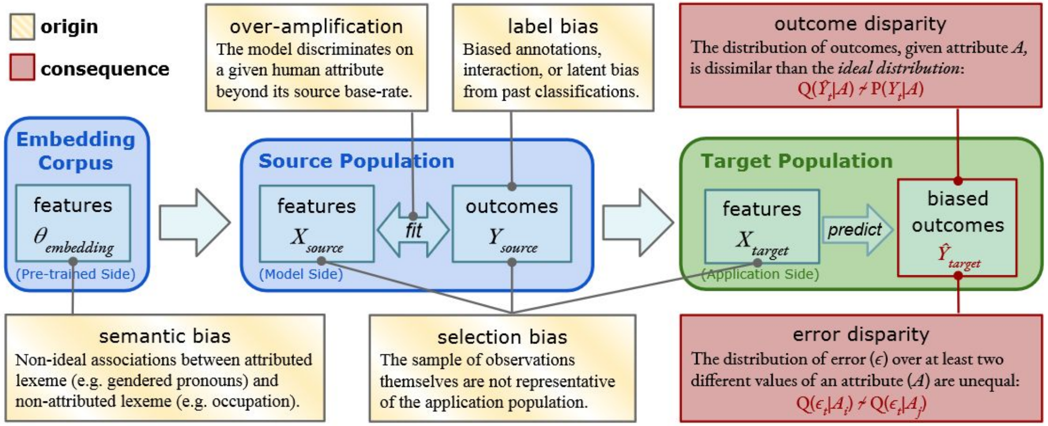

Figure 1: *The Predictive Bias Framework for NLP*: Depiction of where bias may originate within a standard supervised NLP pipeline. Evidence of bias is seen in $\hat{y}$ via *outcome disparity* and *error disparity*.

How can we make research with digital traces + NLP more valid and reliable?

- Inspired by the quantitative social sciences and measurement theory, we develop a conceptual ‘**error framework**’

- Error Frameworks link potential errors to each step of a study’s lifecycle

- Allows researchers to systematically identify and characterize errors

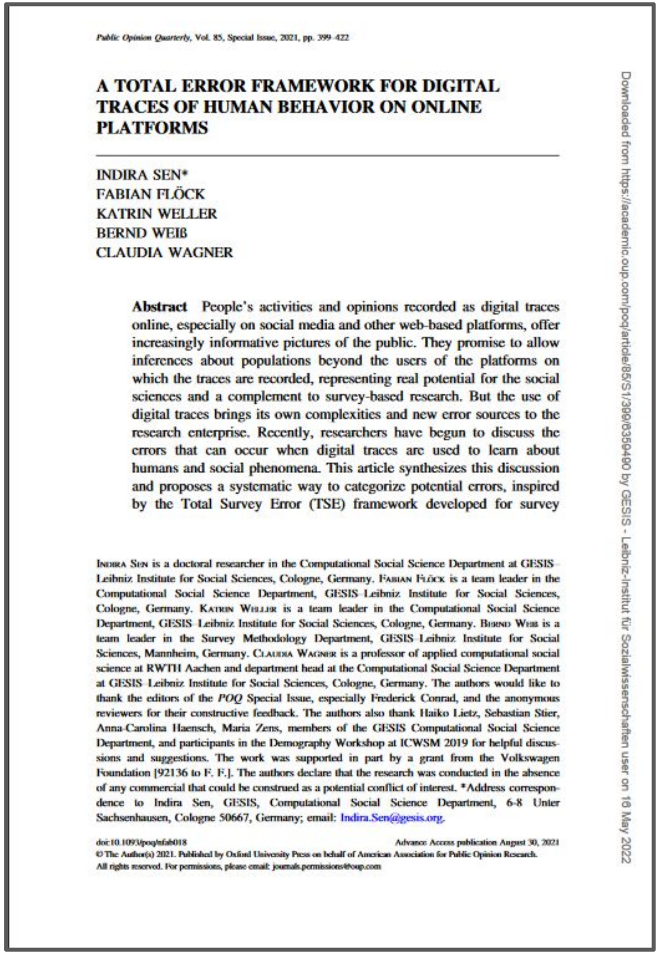

Public Opinion Quarterly, Vol. 85, Special Issue, 2021, pp. 399–422

A TOTAL ERROR FRAMEWORK FOR DIGITAL TRACES OF HUMAN BEHAVIOR ON ONLINE PLATFORMS

INDIRA SEN*
FABIAN FLÖCK
KATRIN WELLER
BERND WEIß
CLAUDIA WAGNER

Abstract People's activities and opinions recorded as digital traces online, especially on social media and other web-based platforms, offer increasingly informative pictures of the public. They promise to allow inferences about populations beyond the users of the platforms on which the traces are recorded, representing real potential for the social sciences and a complement to survey-based research. But the use of digital traces brings its own complexities and new error sources to the research enterprise. Recently, researchers have begun to discuss the errors that can occur when digital traces are used to learn about humans and social phenomena. This article synthesizes this discussion and proposes a systematic way to categorize potential errors, inspired by the Total Survey Error (TSE) framework developed for survey

Indira Sen is a doctoral researcher in the Computational Social Science Department at GESIS–Leibniz Institute for Social Sciences, Cologne, Germany. Fabian Flöck is a team leader in the Computational Social Science Department, GESIS–Leibniz Institute for Social Sciences, Cologne, Germany. Katrin Weller is a team leader in the Computational Social Science Department, GESIS–Leibniz Institute for Social Sciences, Cologne, Germany. Bernd Weiß is a team leader in the Survey Methodology Department, GESIS–Leibniz Institute for Social Sciences, Mannheim, Germany. Claudia Wagner is a professor of applied computational social science at RWTH Aachen and department head at the Computational Social Science Department at GESIS–Leibniz Institute for Social Sciences, Cologne, Germany. The authors would like to thank the editors of the *POQ* Special Issue, especially Frederick Conrad, and the anonymous reviewers for their constructive feedback. The authors also thank Haiko Lietz, Sebastian Stier, Anna-Carolina Haensch, Maria Zens, members of the GESIS Computational Social Science Department, and participants in the Demography Workshop at ICWSM 2019 for helpful discussions and suggestions. The work was supported in part by a grant from the Volkswagen Foundation [92136 to F. F.]. The authors declare that the research was conducted in the absence of any commercial that could be construed as a potential conflict of interest. *Address correspondence to Indira Sen, GESIS, Computational Social Science Department, 6-8 Unter Sachsenhausen, Cologne 50667, Germany; email: Indira.Sen@gesis.org.

doi:10.1093/poq/nfab018 Advance Access publication August 30, 2021

TED-On: A Total Error Framework for Digital Traces of Human Behavior on Online Platforms

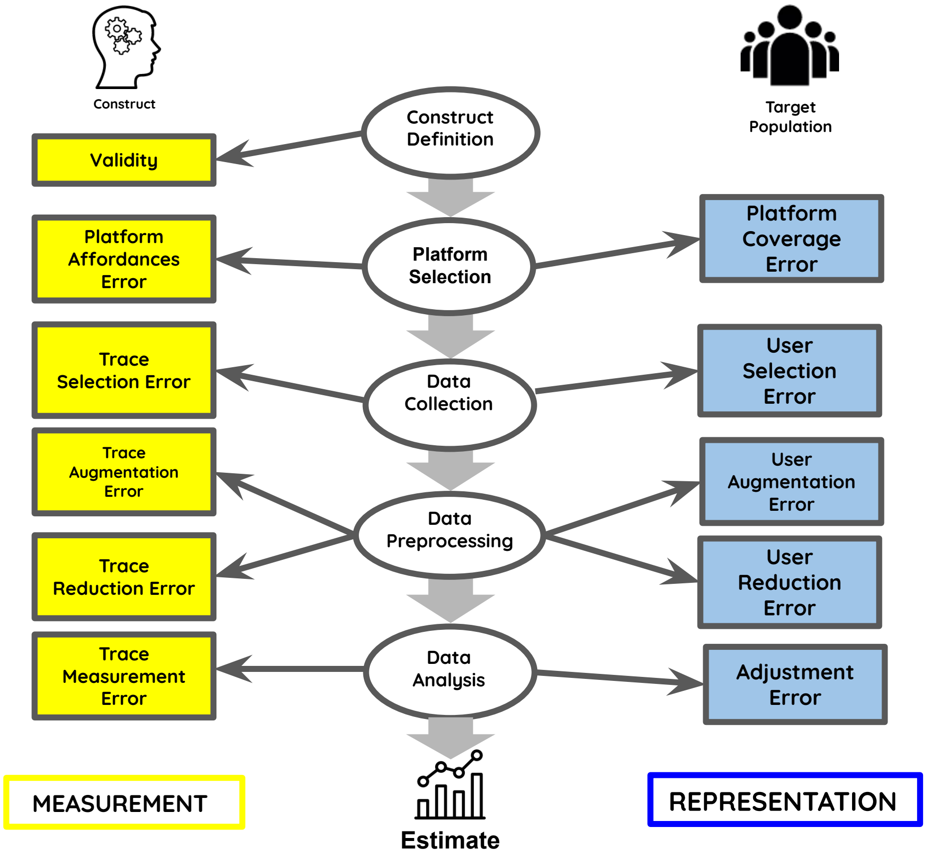

TED-On: A Total Error Framework for Digital Traces of Human Behavior on Online Platforms

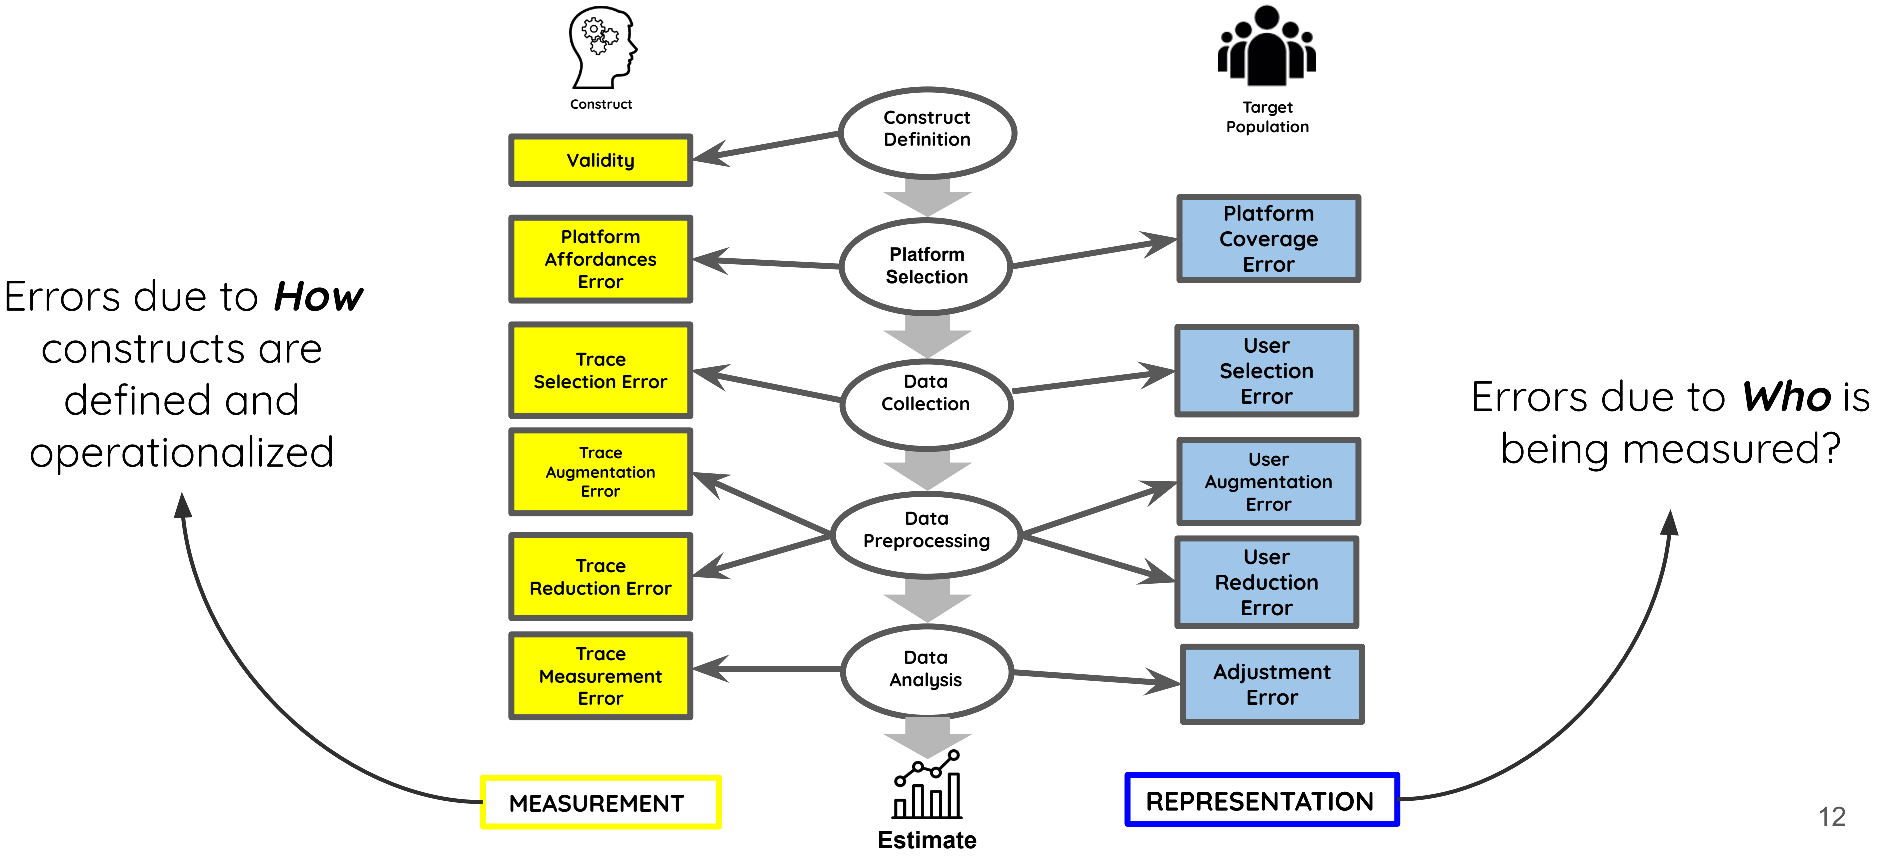

TED-On: A Total Error Framework for Digital Traces of Human Behavior on Online Platforms

How do our modeling choices affect our NLP models?

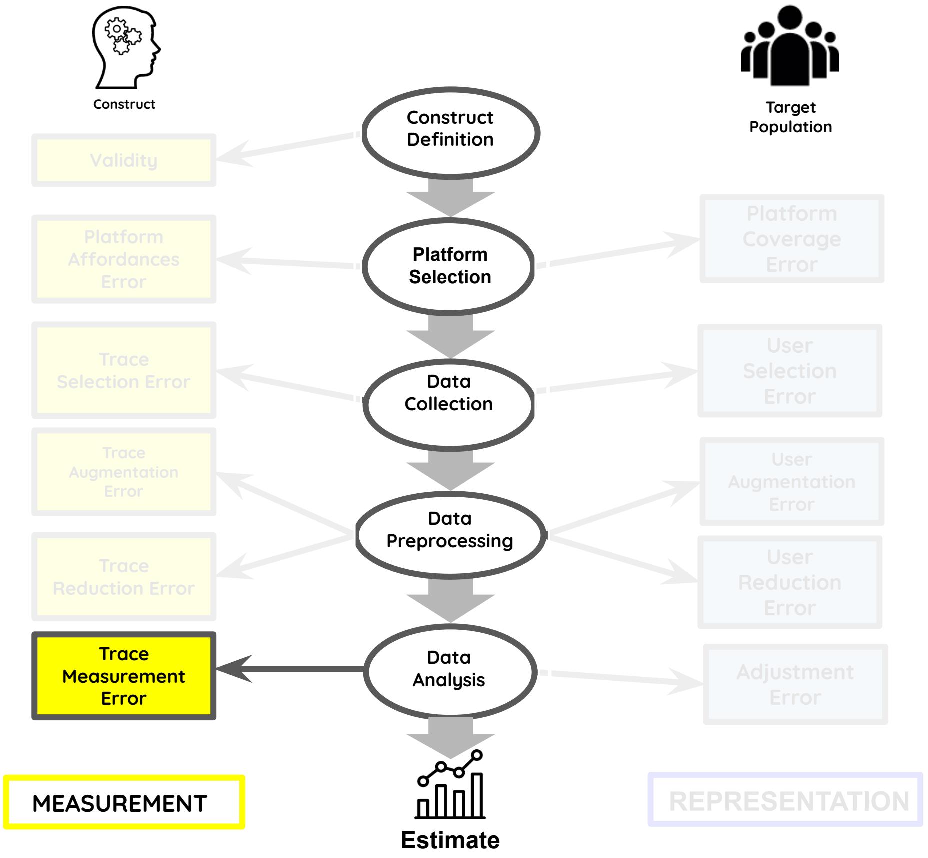

How does Counterfactually Augmented Data Impact Models for Social Computing Constructs?

Indira Sen, Mattia Samory, Fabian Flöck,
Claudia Wagner, Isabelle Augenstein

gesis Leibniz Institute for the Social Sciences
RWTH AACHEN UNIVERSITY
UNIVERSITY OF COPENHAGEN
COPENLU

EMNLP 2021

Jigsaw's AI-powered toxic language detector is now processing 500 million requests daily

Kyle Wiggers
@Kyle_L_Wiggers

February 8, 2021 4:00 AM

Jigsaw

Inside Facebook, Twitter and Google's AI battle over your social lives

From stamping out trolls to removing fake bot accounts, here's how social networks are waging war using AI weapons.

Alfred Ng July 17, 2018 5:00 a.m. PT

Social Computing NLP Models

- Language technology and NLP tools underpin AI methods for platform governance, employing **social computing models**

Social Computing NLP Models

- Language technology and NLP tools underpin AI methods for platform governance, employing **social computing models**

- Supervised NLP models often learn **dataset artefacts** [Le Bras'20, Linzen'20, Samory'21] rather than the construct they're supposed to model

Social Computing NLP Models

- Language technology and NLP tools underpin AI methods for platform governance, employing **social computing models**

- Supervised NLP models often learn **dataset artefacts** [Le Bras'20, Linzen'20, Samory'21] rather than the construct they're supposed to model

> a sentiment model trained on movie reviews considers words like '**plot**', '**script**', and '**romantic**' to be indicative of sentiment

Social Computing NLP Models

- Language technology and NLP tools underpin AI methods for platform governance, employing **social computing models**

- Supervised NLP models often learn **dataset artefacts** [Le Bras'20, Linzen'20, Samory'21] rather than the construct they're supposed to model

> a sentiment model trained on movie reviews considers words like '**plot**', '**script**', and '**romantic**' to be indicative of sentiment

- Hinders one type of model robustness => **out-of-domain generalizability**

Potential Solution: Counterfactually Augmented Data (CAD)

What is Counterfactually Augmented Data?

What is Counterfactually Augmented Data?

Obtained by making **minimal changes** to existing data to flip the label while still being meaningful

What is Counterfactually Augmented Data?

Obtained by making **minimal changes** to existing data to flip the label while still being meaningful

original: I don’t want a female commentator on soccer matches

Sexist

What is Counterfactually Augmented Data?

Obtained by making **minimal changes** to existing data to flip the label while still being meaningful

original: I don’t want a female commentator on soccer matches

Sexist

CAD: I don’t want an AI commentator on soccer matches

non-sexist

What is Counterfactually Augmented Data?

Obtained by making **minimal changes** to existing data to flip the label while still being meaningful

original: I don't want a female commentator on soccer matches

Sexist

CAD: I don't want an AI commentator on soccer matches

non-sexist

Assumption. Obtain a pair of data points with opposite labels with everything in common **but the construct** => better capture meaningful features

This Work: Does CAD work and how?

- Does training on CAD improve model performance?

- Are some types of CAD better than others?

- Why is CAD Better?

Constructs, Data, and Methods

constructs

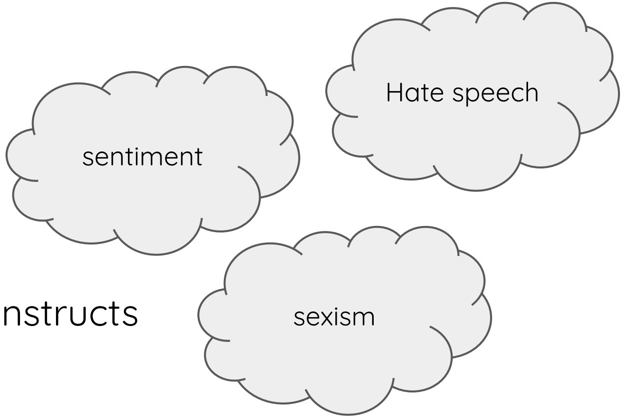

Constructs, Data, and Methods

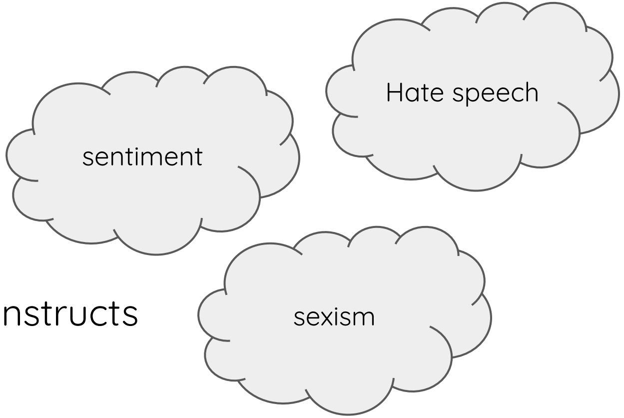

constructs

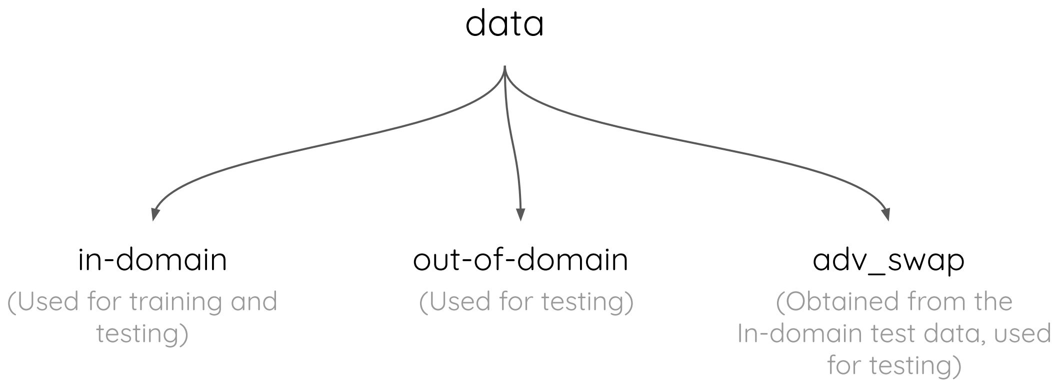

Constructs, Data, and Methods

constructs

- sentiment
- Hate speech
- sexism

data

- in-domain (Used for training and testing)
- out-of-domain (Used for testing)
- adv_swap (Obtained from the In-domain test data, used for testing)

methods

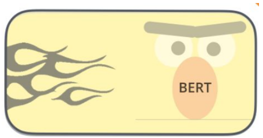

Fine-tuned BERT

BoW LogReg

Constructs, Data, and Methods

Hate speech

sentiment

constructs

sexism

data

in-domain

(Used for training and testing)

out-of-domain

(Used for testing)

adv_swap

(Obtained from the In-domain test data, used for testing)

methods

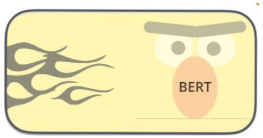

Fine-tuned BERT

BoW LogReg

Constructs, Data, and Methods

constructs

sentiment

Hate speech

sexism

data

in-domain
(Used for training and testing)

out-of-domain
(Used for testing)

adv_swap
(Obtained from the In-domain test data, used for testing)

methods

Fine-tuned BERT

BoW LogReg

Experimental Setup

- **Counterfactual models** trained on 50% original data, **50% CAD** (25% for sexism as there's only positive CAD)

- **Non-Counterfactual models**: models trained on 100% original data

- Results measured with **Macro F1**

Part I: Does training on CAD improve model performance?

Does training on CAD improve model performance?

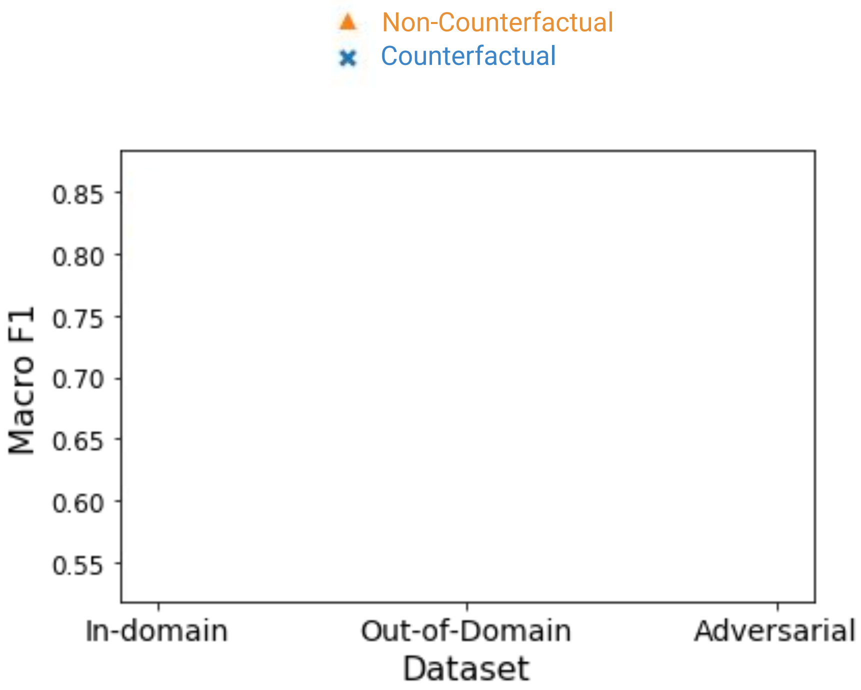

Does training on CAD improve model performance?

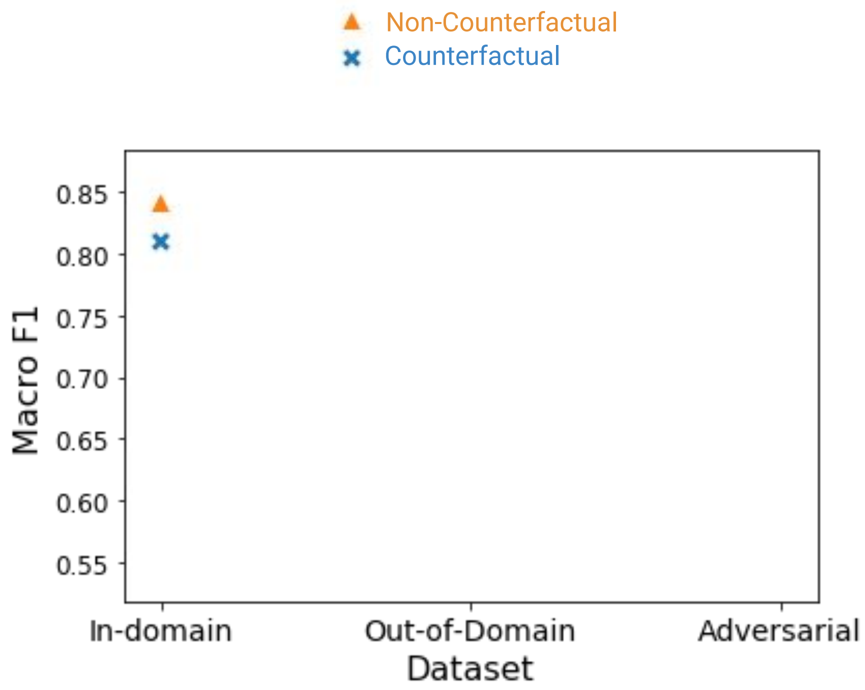

Does training on CAD improve model performance?

❖ in-domain performance drops

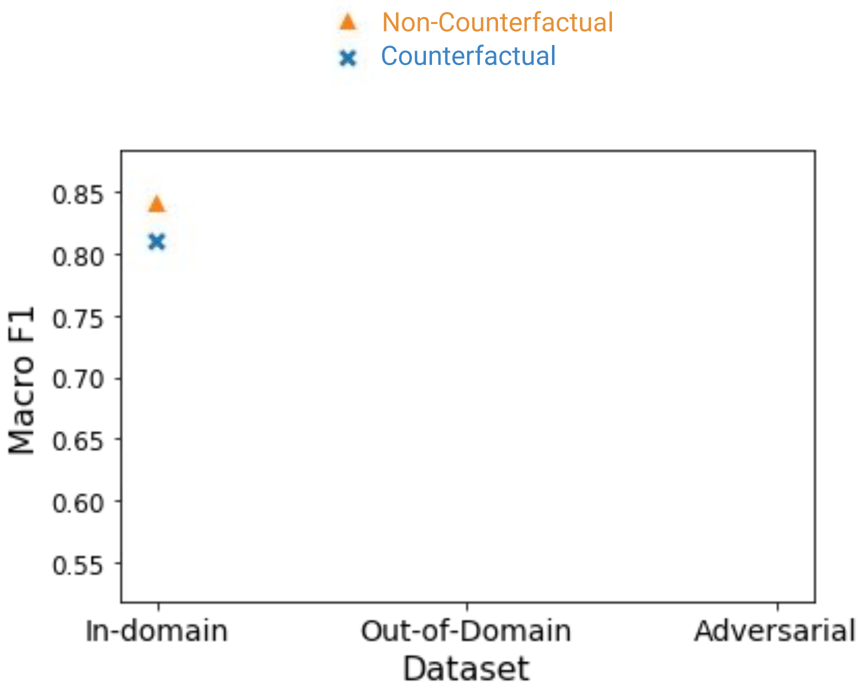

Does training on CAD improve model performance?

❖ in-domain performance drops

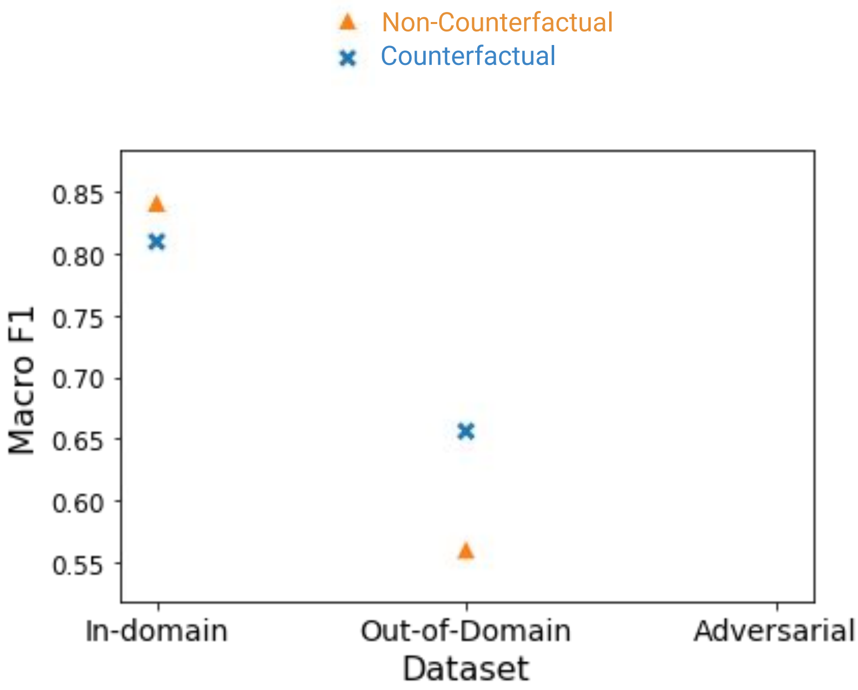

Does training on CAD improve model performance?

- in-domain performance drops
- out-of-domain performance improves

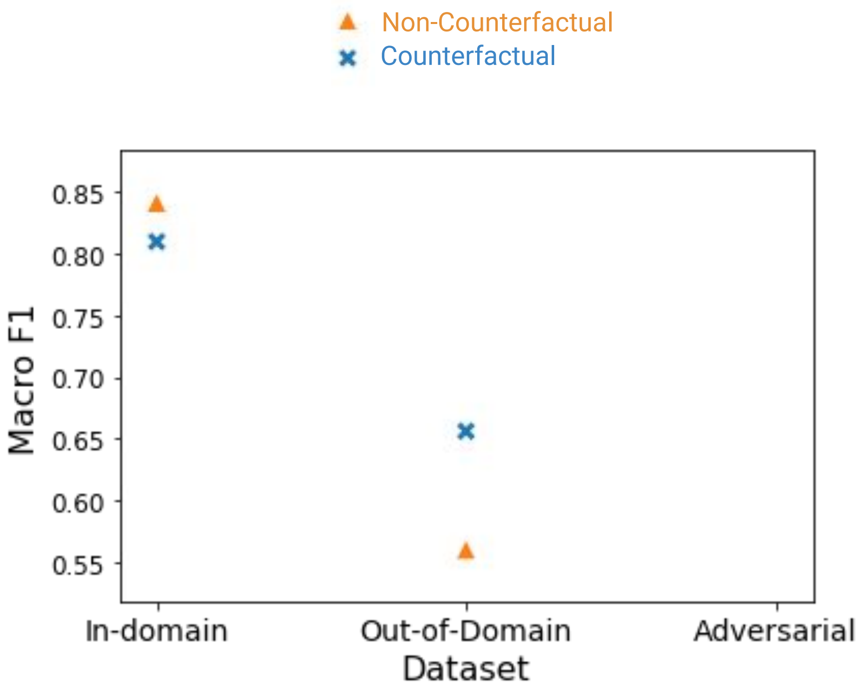

Does training on CAD improve model performance?

- in-domain performance drops
- out-of-domain performance improves

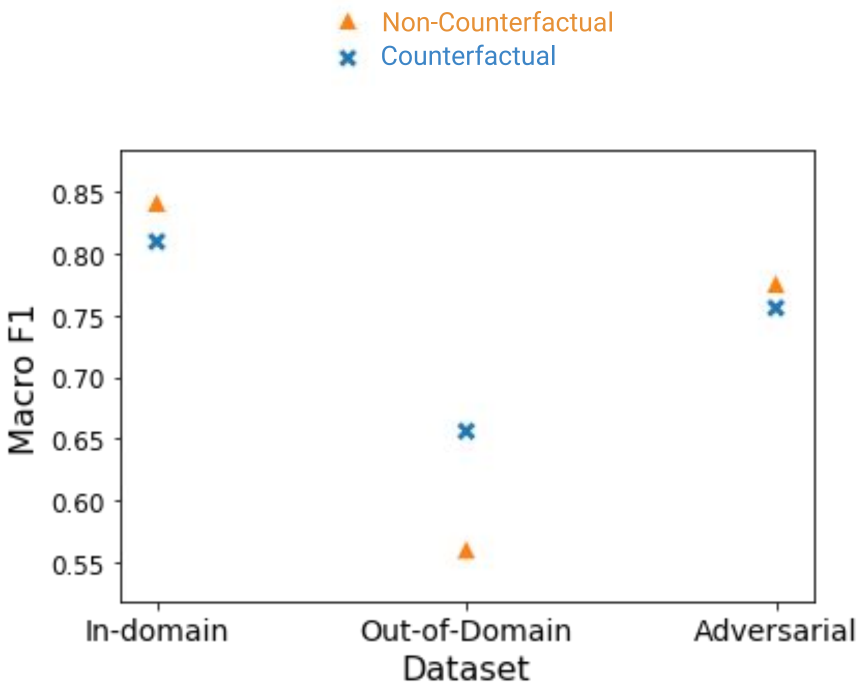

Does training on CAD improve model performance?

- ❖ in-domain performance drops
- ❖ out-of-domain performance improves
- ❖ performance on adversarial data drops -- but the gap is smaller

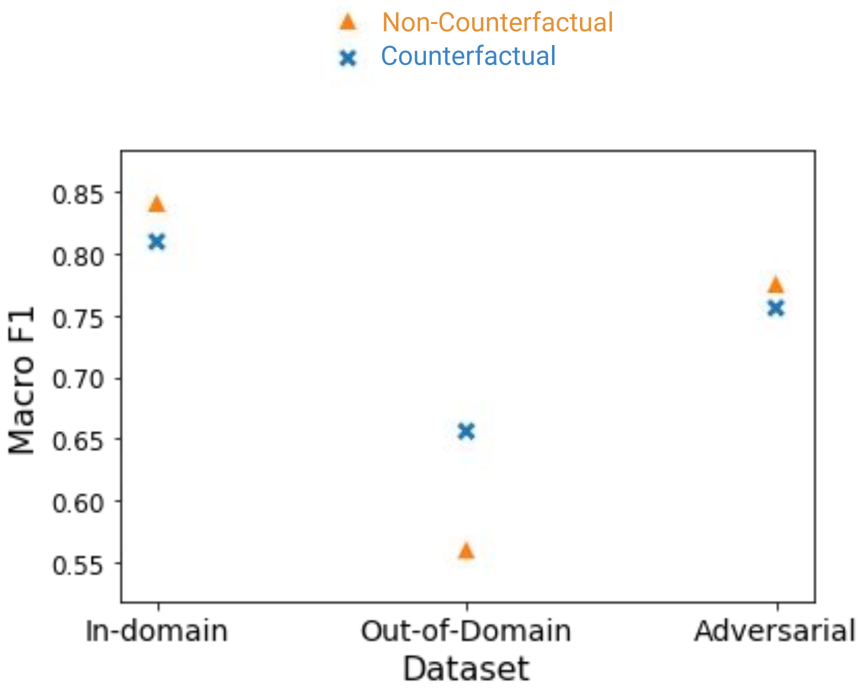

This Work: Does CAD work and how?

- **Does training on CAD improve model performance?**
  Yes, for out-of-domain performance

This Work: Does CAD work and how?

- **Does training on CAD improve model performance?**
  Yes, for out-of-domain performance

- **Are some types of CAD better than others?**

- **Why is CAD Better?**

TED-On: A Total Error Framework for Digital Traces of Human Behavior on Online Platforms

CAD helps us tackle some modeling errors

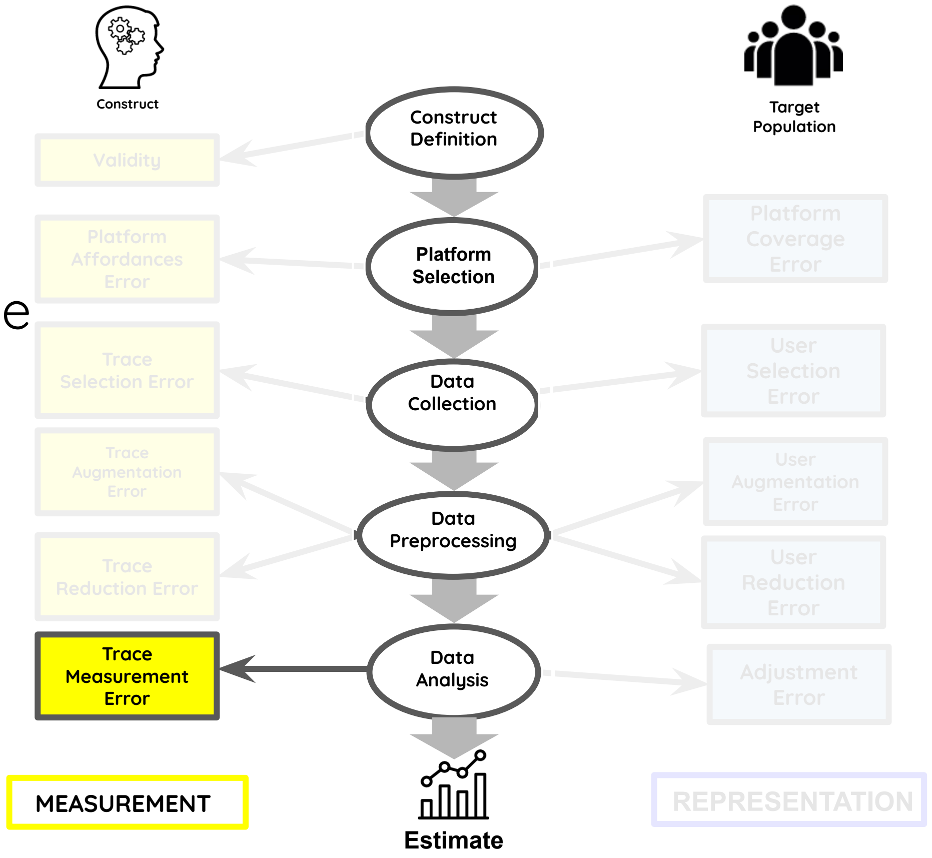

TED-On: A Total Error Framework for Digital Traces of Human Behavior on Online Platforms

CAD helps us tackle some modeling errors

But what if errors occur **upstream?**

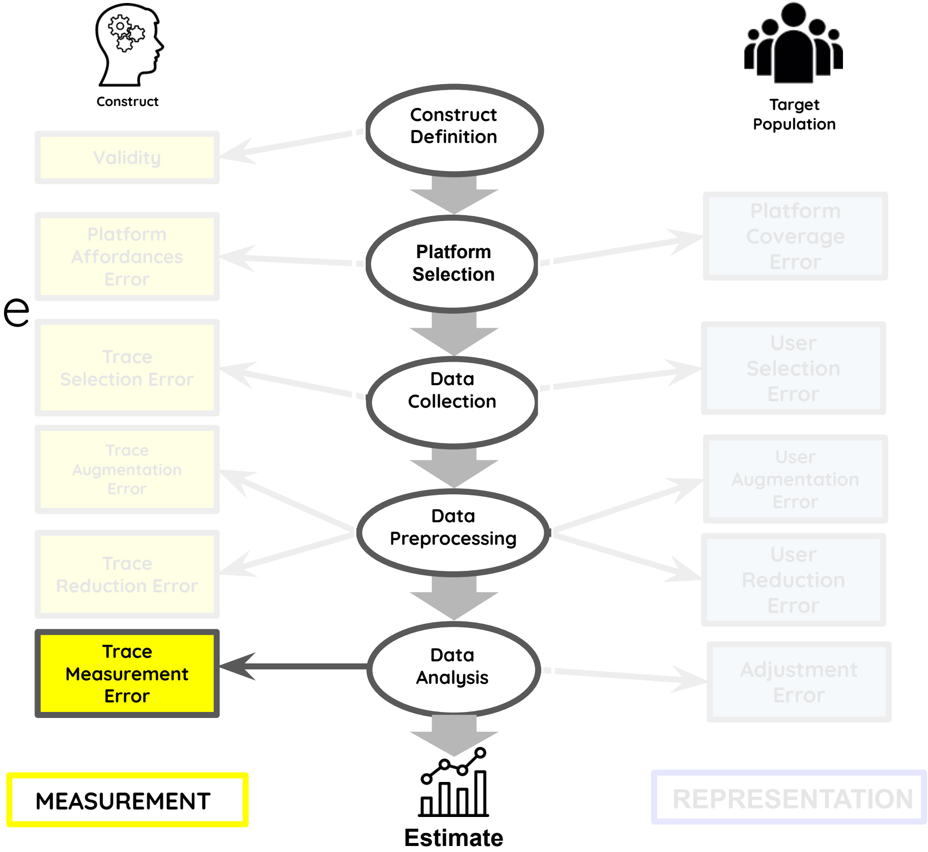

TED-On: A Total Error Framework for Digital Traces of Human Behavior on Online Platforms

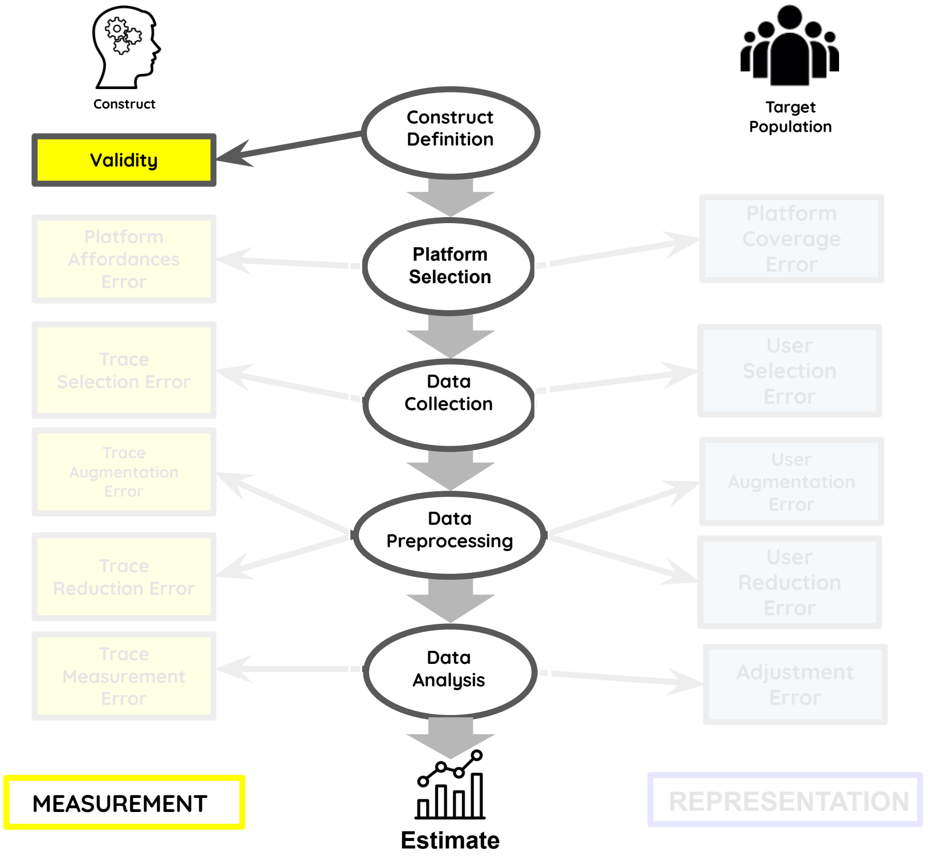

TED-On: A Total Error Framework for Digital Traces of Human Behavior on Online Platforms

How do we define and operationalize the construct in a theory-driven way?

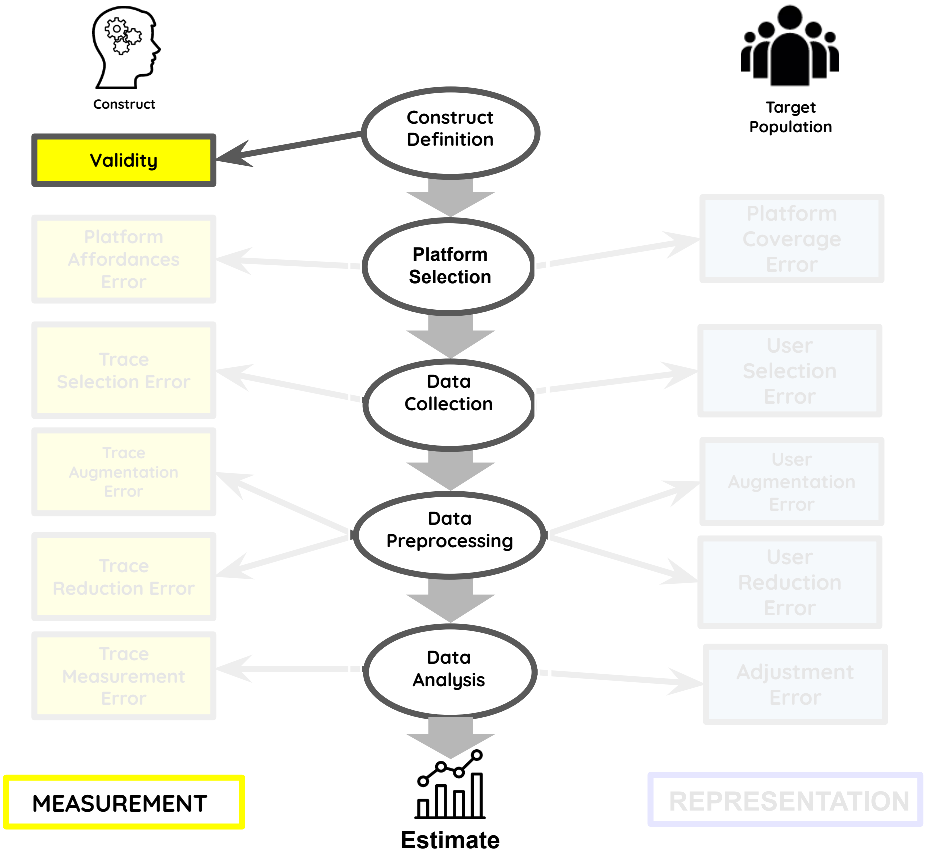

Depression at Work: Exploring Depression in Major US Companies from Online Reviews

Indira Sen, Daniele Quercia, Marios Constantinides, Matteo Montecchi, Licia Capra, Sanja Šćepanović, Renzo Bianchi

gesis Leibniz Institute for the Social Sciences
NOKIA Bell Labs
UCL
KING'S College LONDON
NTNU

CSCW 2022

Depression in the workplace

- Damages employee wellbeing

- Typically measured via surveys; But:
  - Reporting biases
  - Small, self-selected samples

- Measurement tools are either tailored towards burnout or general depression, **not specific work-related depression**

Company Reviews as an alternative to surveys

Crowdsourcing Job Satisfaction Data: Examining the Construct Validity of Glassdoor.com Ratings

Richard N. Landers[1], Robert C. Brusso[2], and Elena M. Auer[1]

1. University of Minnesota
2. Capital One

Employee Satisfaction and Corporate Performance: Mining Employee Reviews on Glassdoor.com

Completed Research Paper

Ning Luo
Gabelli School of Business
Fordham University, New York, USA
nluo1@fordham.edu

Yilu Zhou
Gabelli School of Business
Fordham University, New York, USA
yzhou62@fordham.edu

John J. Shon
Gabelli School of Business
Fordham University, New York, USA
jshon@fordham.edu

CHI 2020 Paper

CHI 2020, April 25–30, 2020, Honolulu, HI, USA

Modeling Organizational Culture with Workplace Experiences Shared on Glassdoor

Vedant Das Swain*, Koustuv Saha*, Manikanta D. Reddy, Hemang Rajvanshy, Gregory D. Abowd, Munmun De Choudhury
Georgia Institute of Technology
Atlanta, Georgia, USA
{vedantswain, koustuv.saha, mani, hemangr, abowd, munmund}@gatech.edu

Data: Company Reviews

5.0 ★★★★★

Current Employee, more than 10 years

Good place to work

Sep 20, 2022 - Senior Software Engineer in Irving, TX

✓ Recommend CEO Approval ✓ Business Outlook

Pros

They have a flexible schedule and it is a good work life balance.

Cons

Not much to contribute on the negatives. They are very european.

Be the first to find this review helpful

Helpful Share Report

Data: Company Reviews

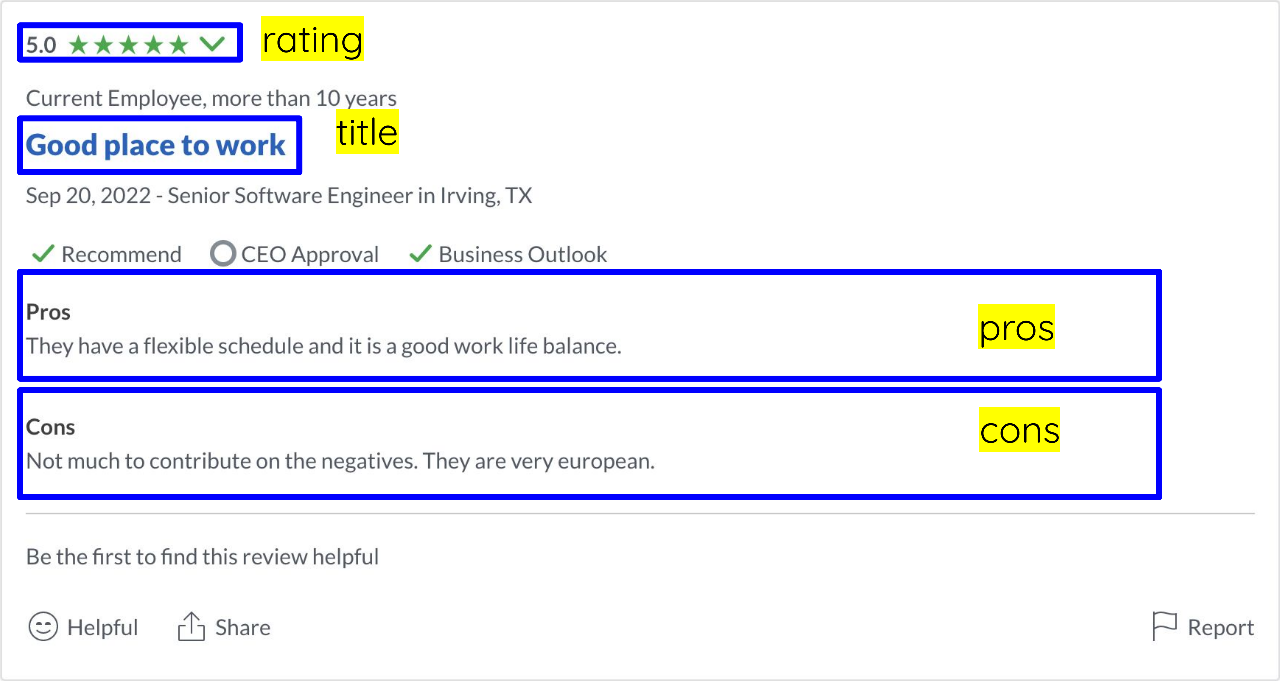

Data: Company Reviews

- We use 380K+ reviews between 2008-2020 about 104 S&P 500 companies from **a popular company reviewing platform**

- Each company has at least 1000 reviews

Theory: The Occupational Depression Inventory (ODI)

- Specifically created for measuring workplace depression

- 9 item survey covering many dimensions — anhedonia, fatigue, suicidal ideation, and more

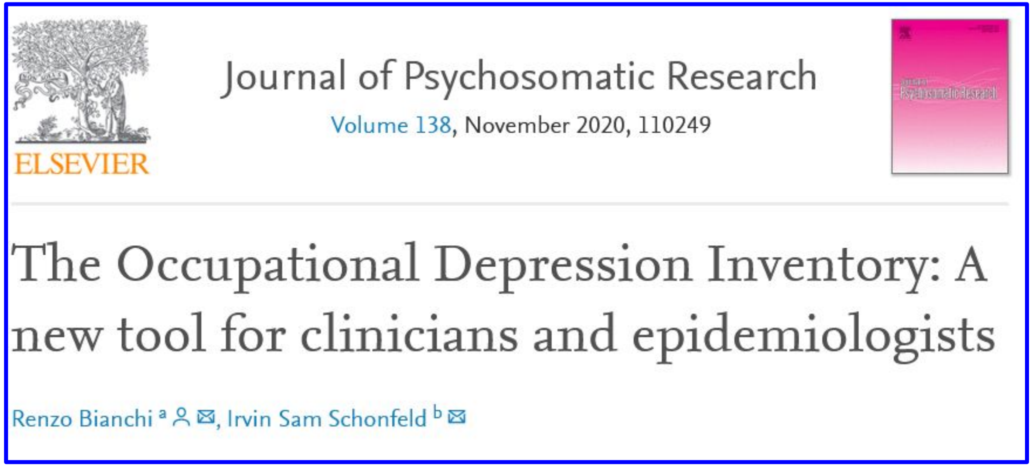
ELSEVIER

Journal of Psychosomatic Research

Volume 138, November 2020, 110249

The Occupational Depression Inventory: A new tool for clinicians and epidemiologists

Renzo Bianchi[a], Irvin Sam Schonfeld[b]

Theory: The Occupational Depression Inventory (ODI)

- ❖ Specifically created for measuring workplace depression

- ❖ 9 item survey covering many dimensions — anhedonia, fatigue, suicidal ideation, and more

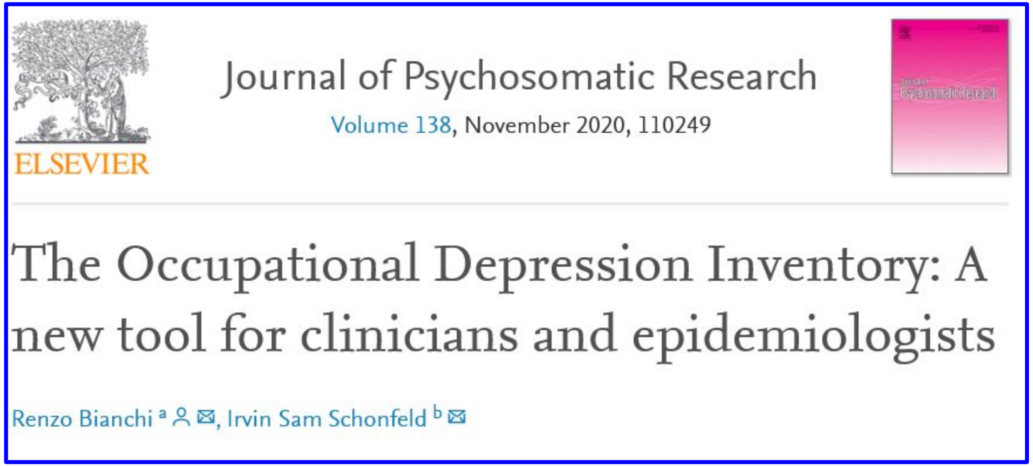
ELSEVIER

Journal of Psychosomatic Research

Volume 138, November 2020, 110249

The Occupational Depression Inventory: A new tool for clinicians and epidemiologists

Renzo Bianchi[a], Irvin Sam Schonfeld[b]

But: How do we use this in passive reviews?

autoODI: Discovering reviews with markers of ODI

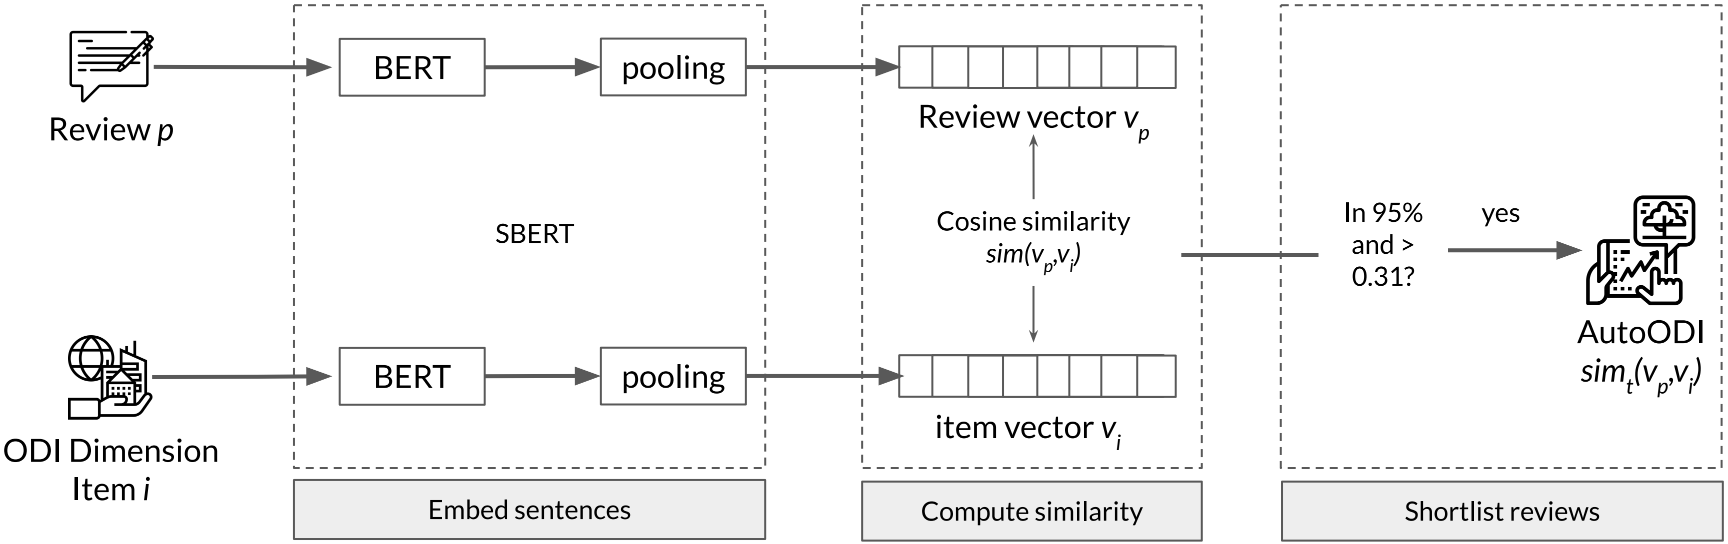

autoODI: Discovering reviews with markers of ODI

Unlabeled data *d*

Scale item *s*

input

autoODI: Discovering reviews with markers of ODI

Unlabeled data *d*

Scale item *s*

input

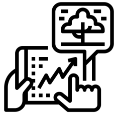

Relevant traces

output

autoODI: Discovering reviews with markers of ODI

Unlabeled data *d*

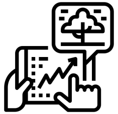

Relevant traces

Scale item *s*

autoODI: Discovering reviews with markers of ODI

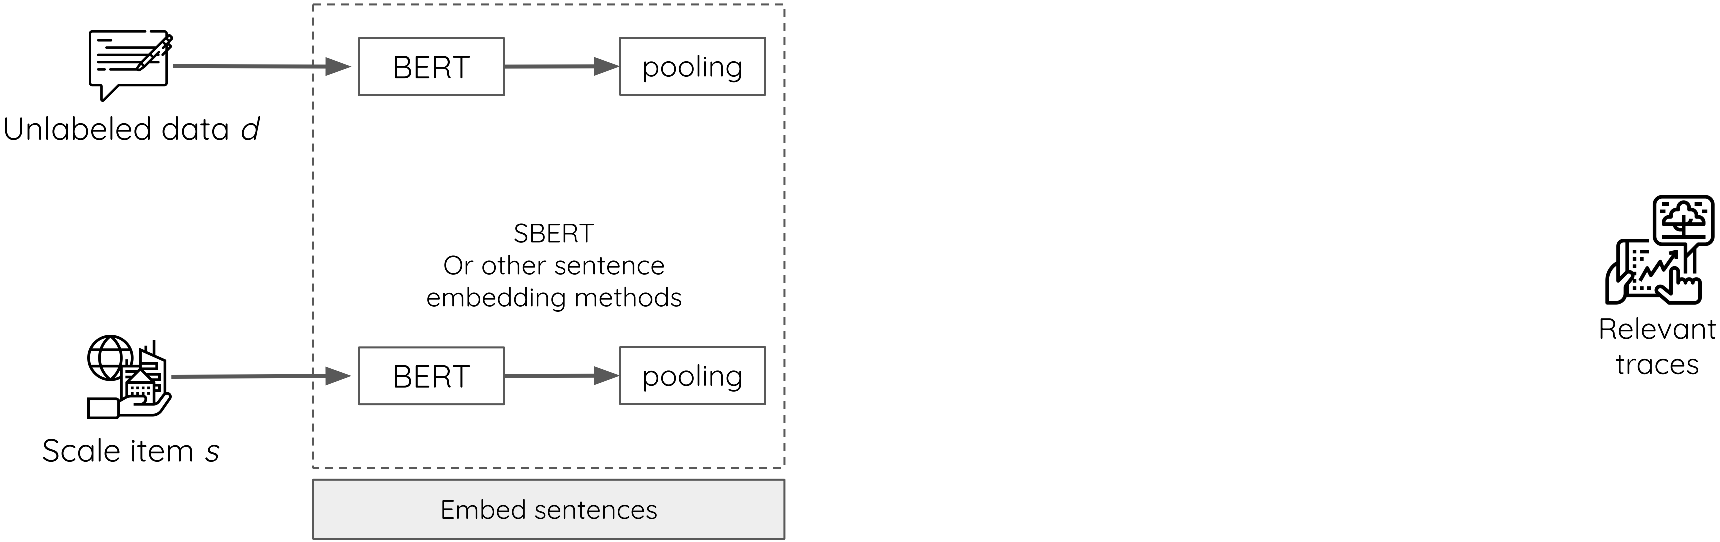

autoODI: Discovering reviews with markers of ODI

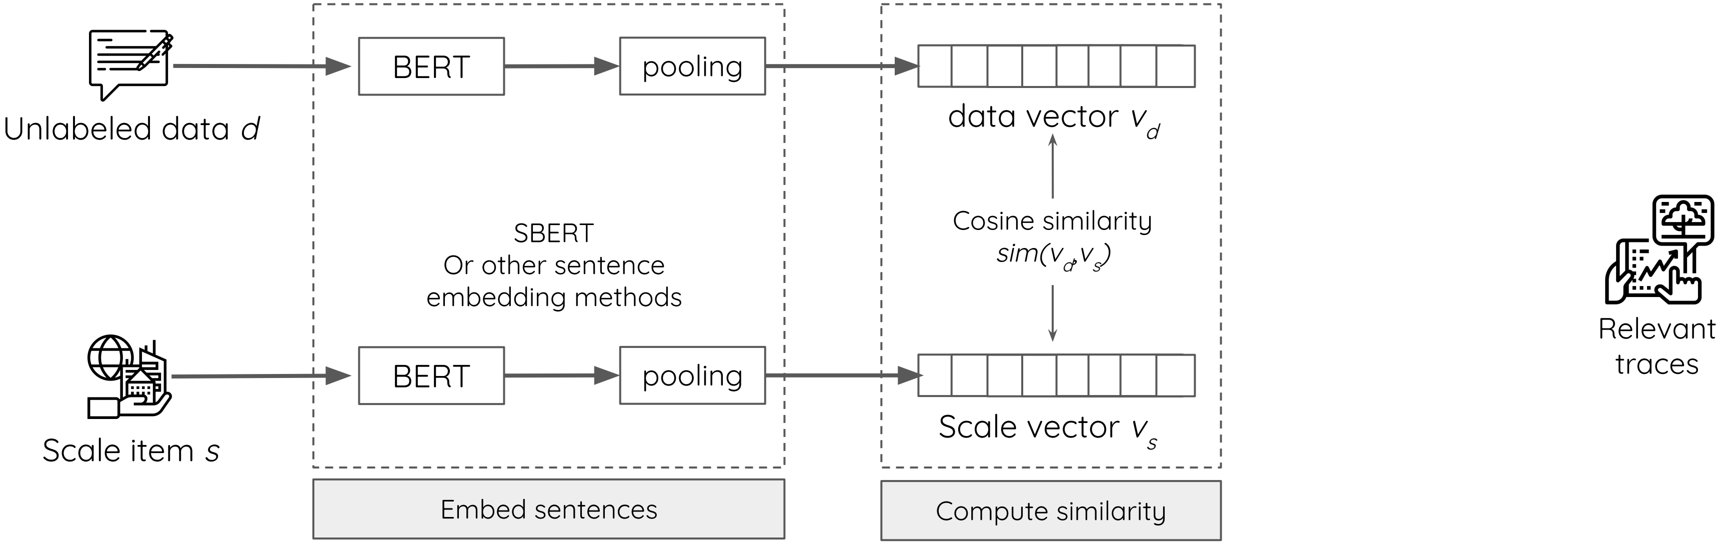

autoODI: Discovering reviews with markers of ODI

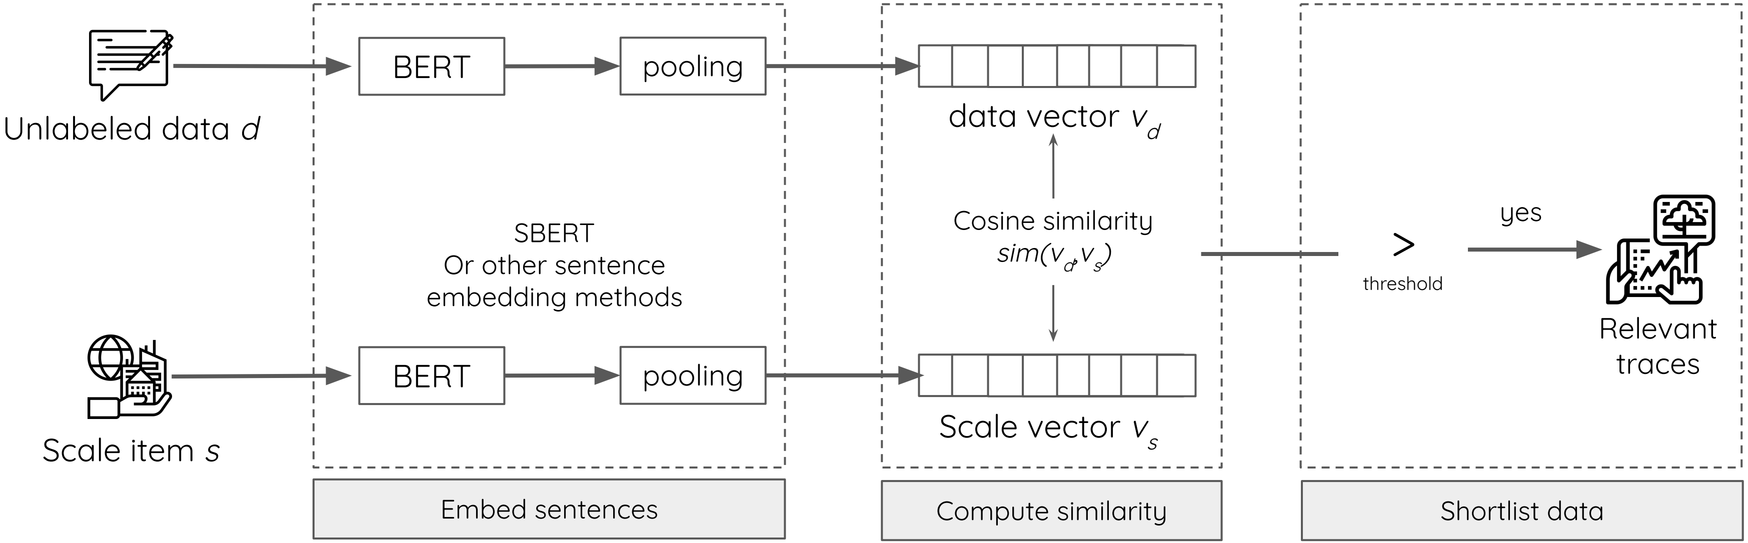

autoODI: Example

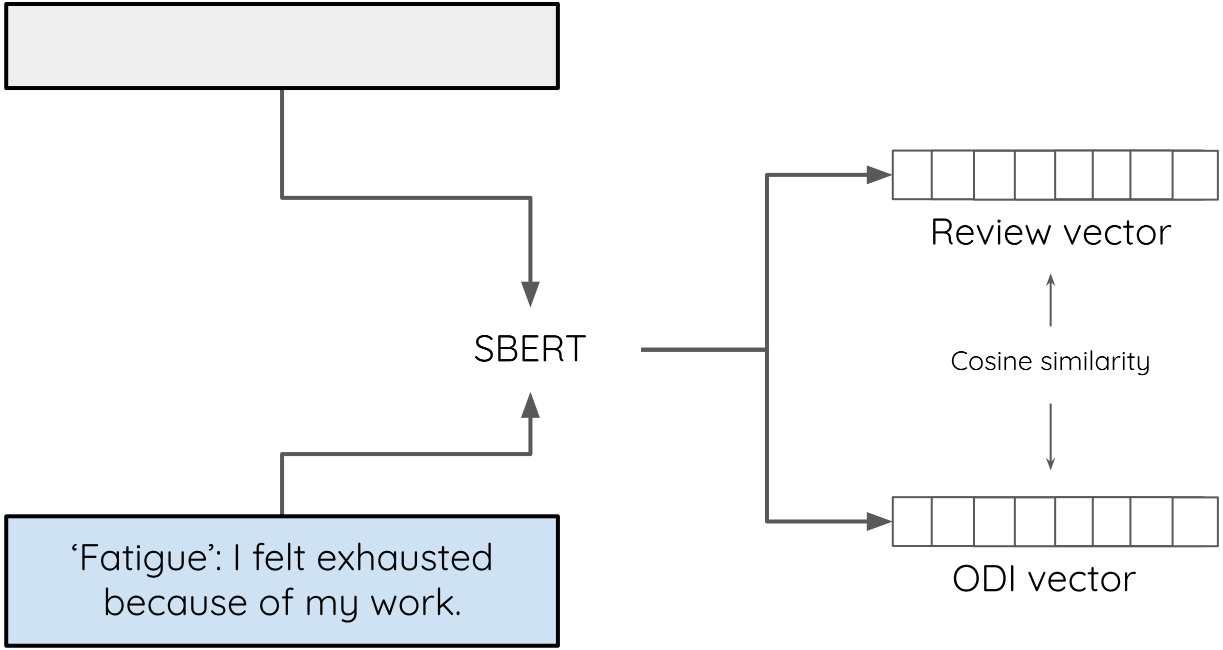

autoODI: Example

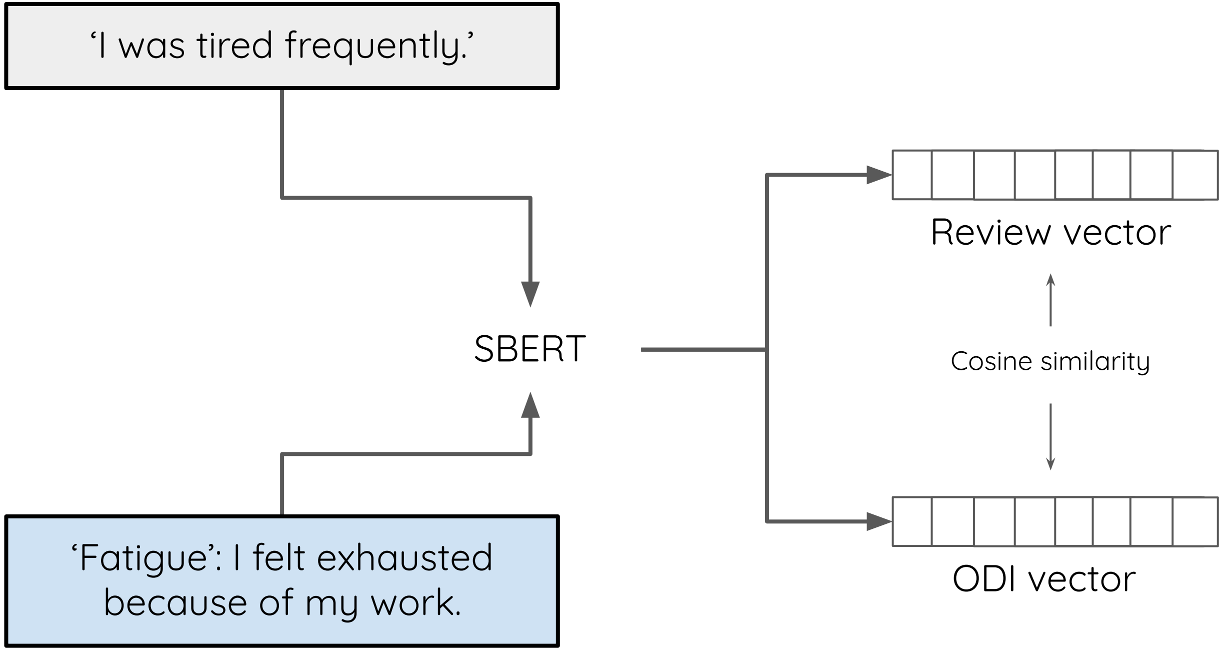

autoODI: Example

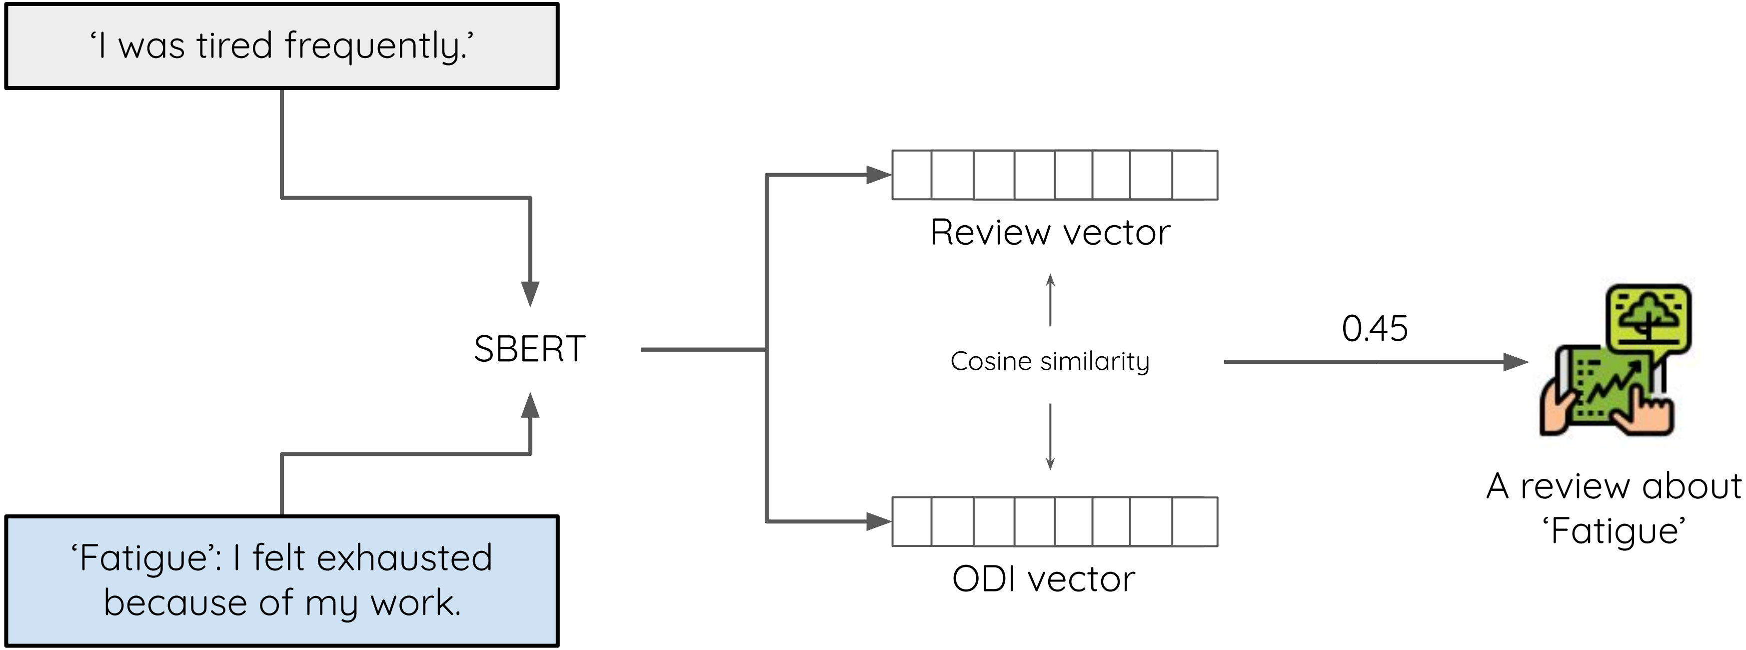

autoODI: Example

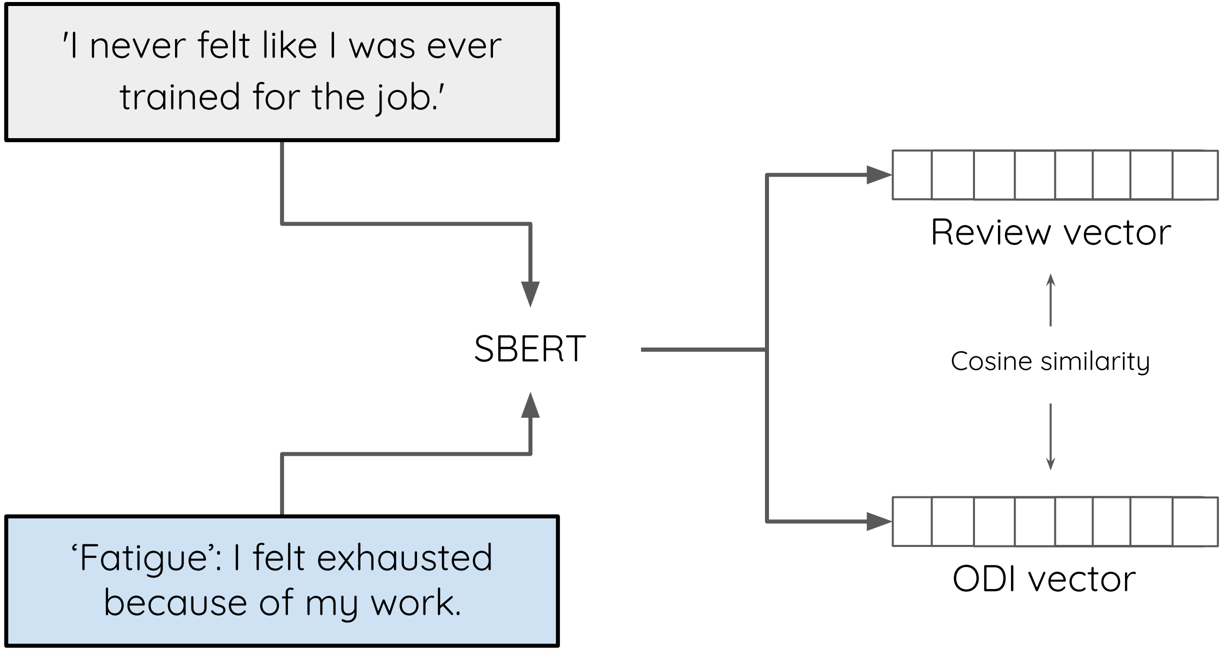

autoODI: Example

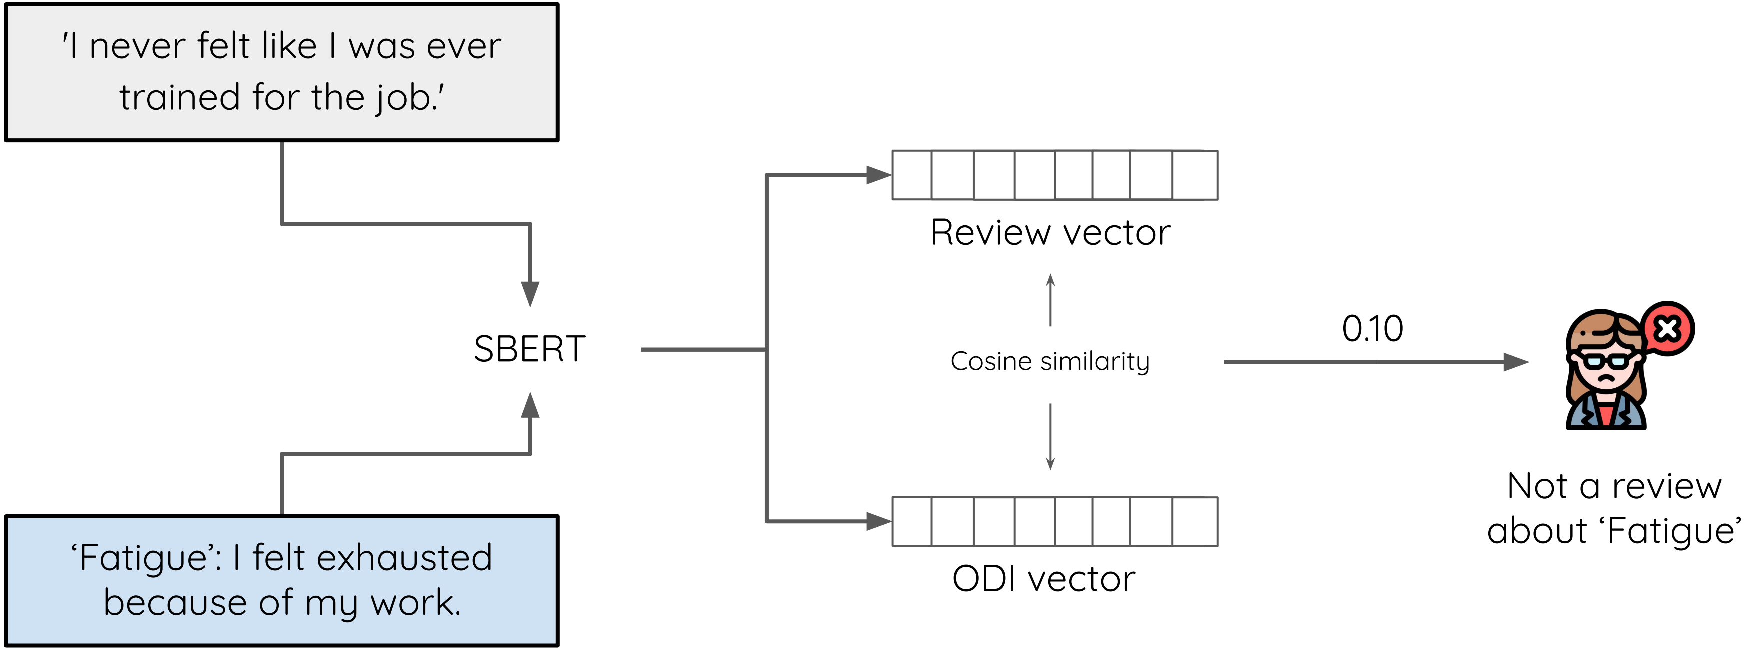

The Four Dimensions of autoODI

ODI Category	Survey Item	Example of a Company Review
Depressed mood	I felt depressed because of my job.	This was the most depressing job I've ever had.
Sleep Alterations	My sleep was disturbed as a result of job stress (I had difficulties falling asleep or staying asleep, or I slept much more than usual).	I fell asleep countless times.
Fatigue / Loss of Energy	I felt exhausted because of my work.	I came into work already tired many days.
Feelings of Worthlessness	My experience at work made me feel like a failure.	All in all, a very negative work experience for me.

Does the **autoODI** technique actually work?

- We validate the method in the following ways:
  - Linguistic validation: obtain and rank most highly occurring keywords (1,2,3-grams) in shortlisted reviews by TF-IDF

Does the **autoODI** technique actually work?

- We validate the method in the following ways:
  - Linguistic validation: obtain and rank most highly occurring keywords (1,2,3-grams) in shortlisted reviews by TF-IDF

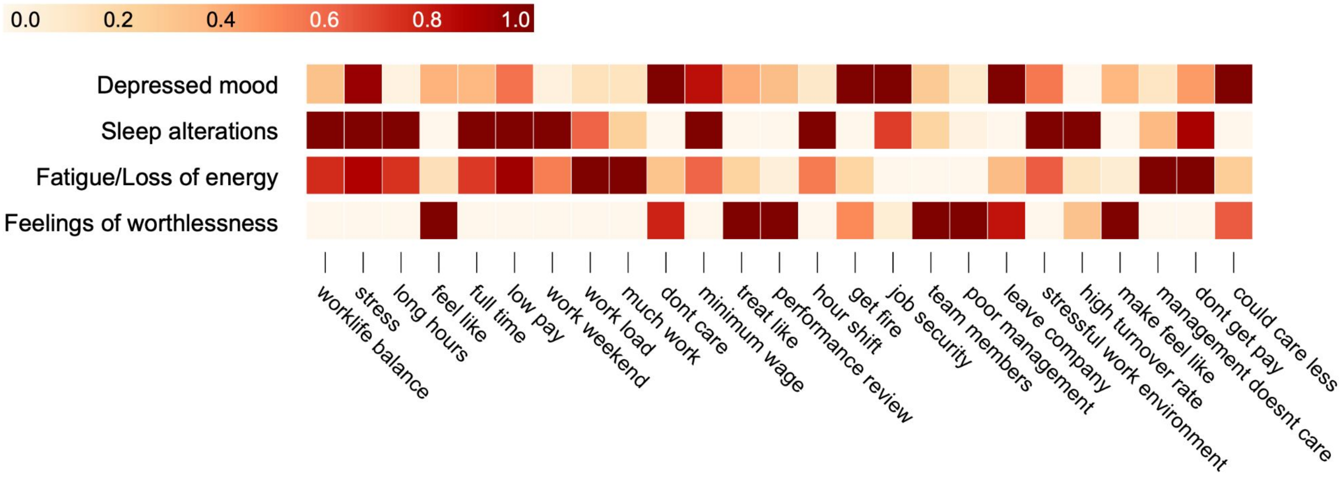

Does the **autoODI** technique actually work?

- We validate the method in the following ways:
  - Linguistic validation: obtain and rank most highly occurring keywords (1,2,3-grams) in shortlisted reviews by TF-IDF

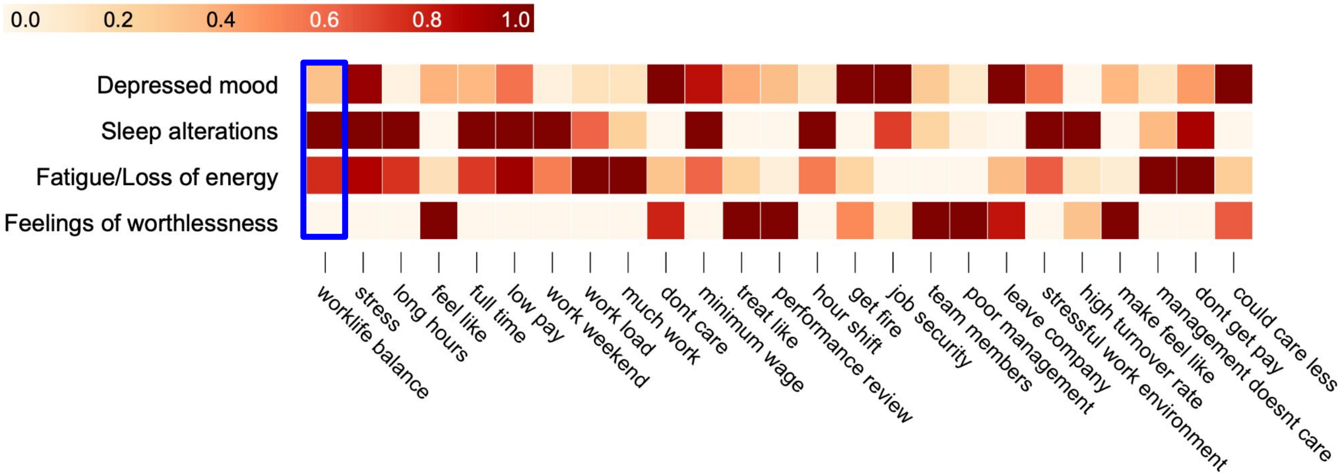

Applying the method: On a company and state level

Are high ODI scores associated with **negative state outcomes?**

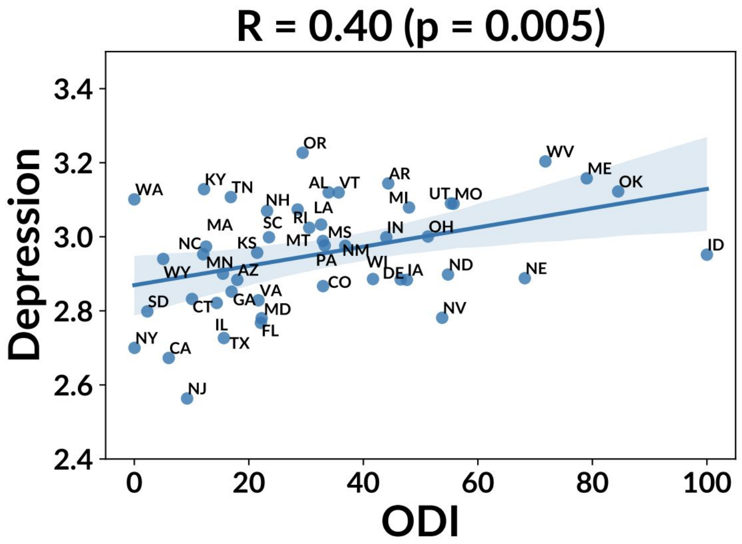

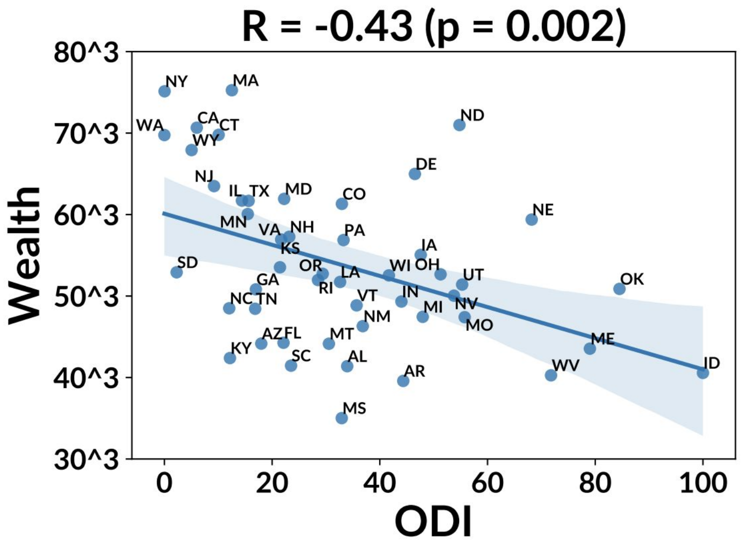

Takeaways from autoODI

Takeaways from autoODI

- We devise an unobtrusive measure of **occupational depression that is grounded in theory** — using sentence embeddings and survey items

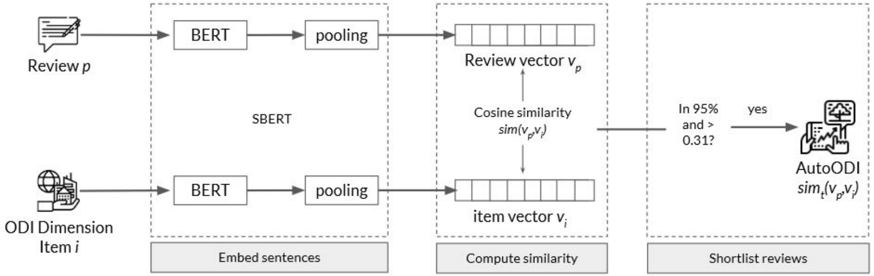

Takeaways from autoODI

- We devise an unobtrusive measure of **occupational depression that is grounded in theory** — using sentence embeddings and survey items

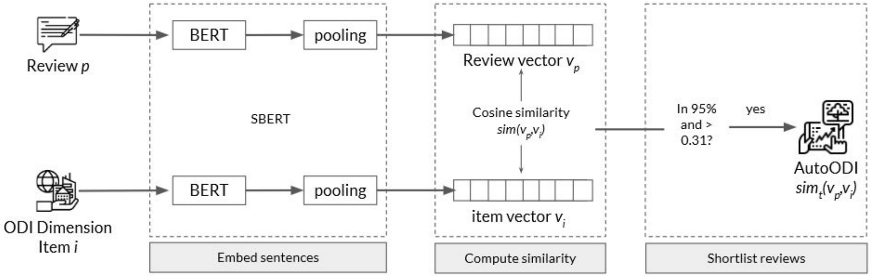

- We find that companies and states with high ODI scores have less financial success

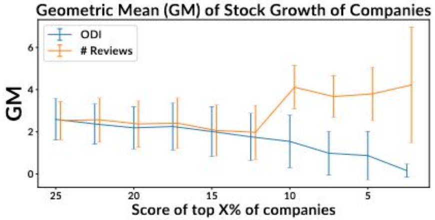

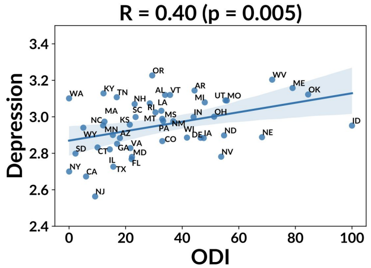

Strengths and weaknesses of the method

- \+ semi-supervised with input from domain experts => no labeled training data needed, applicable to other constructs
- \+ Captures related concepts => health and work-life balance

Strengths and weaknesses of the method

+ semi-supervised with input from domain experts => no labeled training data needed, applicable to other constructs
+ Captures related concepts => health and work-life balance

- SBERT is good at identifying topical similarity, not valence
  - ➢ Review considered relevant for ‘At my work, there is not so much to do’: ‘All I do is work.’
- Word-sense disambiguation issues

Concluding thoughts on CSS + NLP Measurements

- Digital traces have great potentials for the social sciences
- But also pitfalls!

Concluding thoughts on CSS + NLP Measurements

- Digital traces have great potentials for the social sciences
- But also pitfalls!
- Error frameworks help us identify, characterise, and document errors

Concluding thoughts on CSS + NLP Measurements

- Digital traces have great potentials for the social sciences
- But also pitfalls!
- Error frameworks help us identify, characterise, and document errors
- We need **theory-driven NLP + CSS**
  - Starting from, integrating, and updating social science theories
  - High-quality data for modeling and model development
  - Innovative methods, quantitative and qualitative

Concluding thoughts on CSS + NLP Measurements

- Digital traces have great potentials for the social sciences
- But also pitfalls!
- Error frameworks help us identify, characterise, and document errors
- We need **theory-driven NLP + CSS**
    - Starting from, integrating, and updating social science theories
    - High-quality data for modeling and model development
    - Innovative methods, quantitative and qualitative

to calibrate the ‘telescope’

Thanks to my amazing collaborators!


Katrin Weller


Claudia Wagner


Fabian Flöck


Bernd Weiss


Mattia Samory


Isabelle Augenstein


Licia Capra


Matteo Montecchi


Sanja Šćepanović


Daniele Quercia


Marios Constantinides


Renzo Bianchi

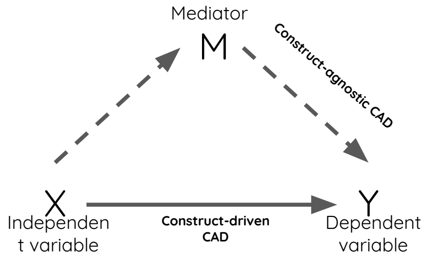

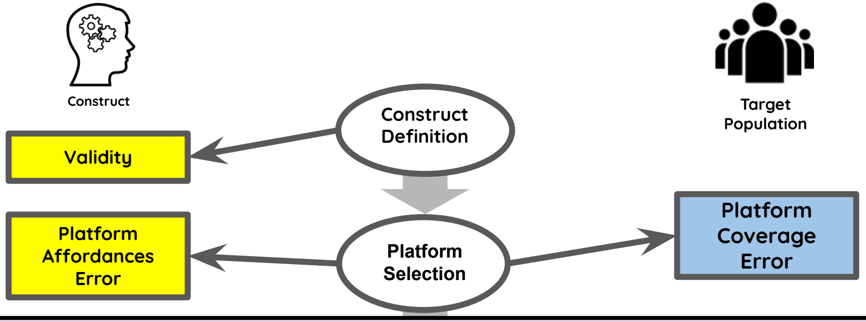

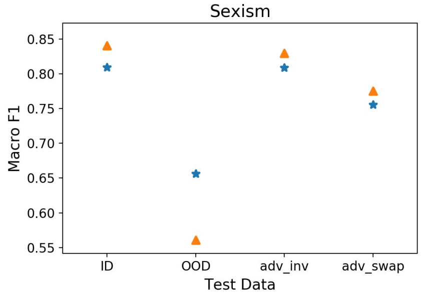

And thank **you** for listening! Questions?

indira.sen@gesis.org