CS 6474/CS 4803 Social Computing: Health and Well-Being

Munmun De Choudhury

munmund@gatech.edu

Week 10 | March 15, 2023

Assignment II available (Due: Apr 12)


Le-shaker.fr

Social Media Derived Behavioral and Affective Markers Predict Postpartum Changes

(De Choudhury, Counts, Horvitz, CSCW 2013; CHI 2013)

Social Media Derived Behavioral and Affective Markers Predict Postpartum Changes

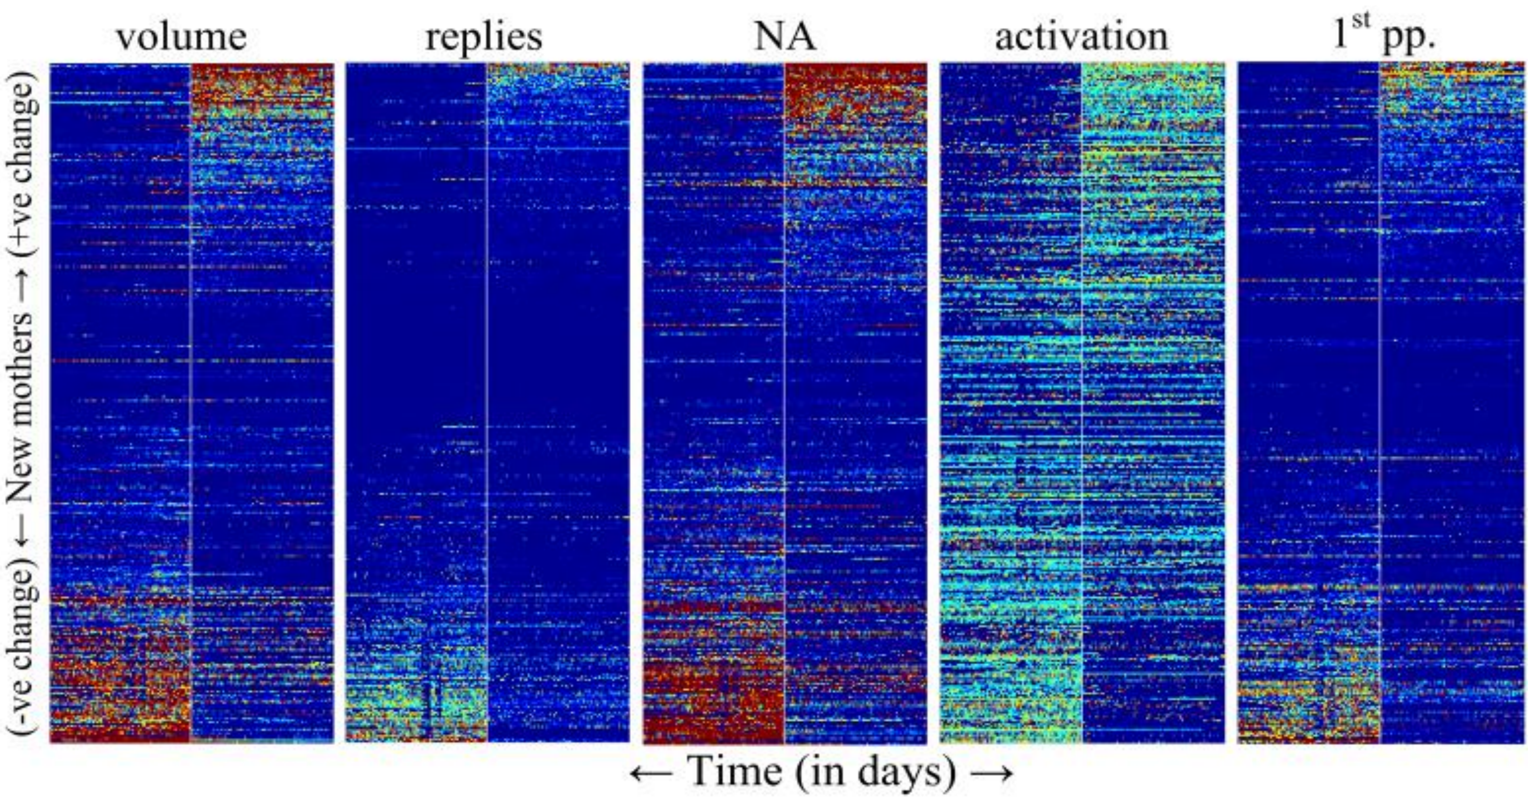

376 users (new mothers); 40,426 posts between March 2011 and July 2012

(Saha and De Choudhury, PACM/CSCW 2018)

Measuring Levels of Acute Stress in College Campuses with Social Media

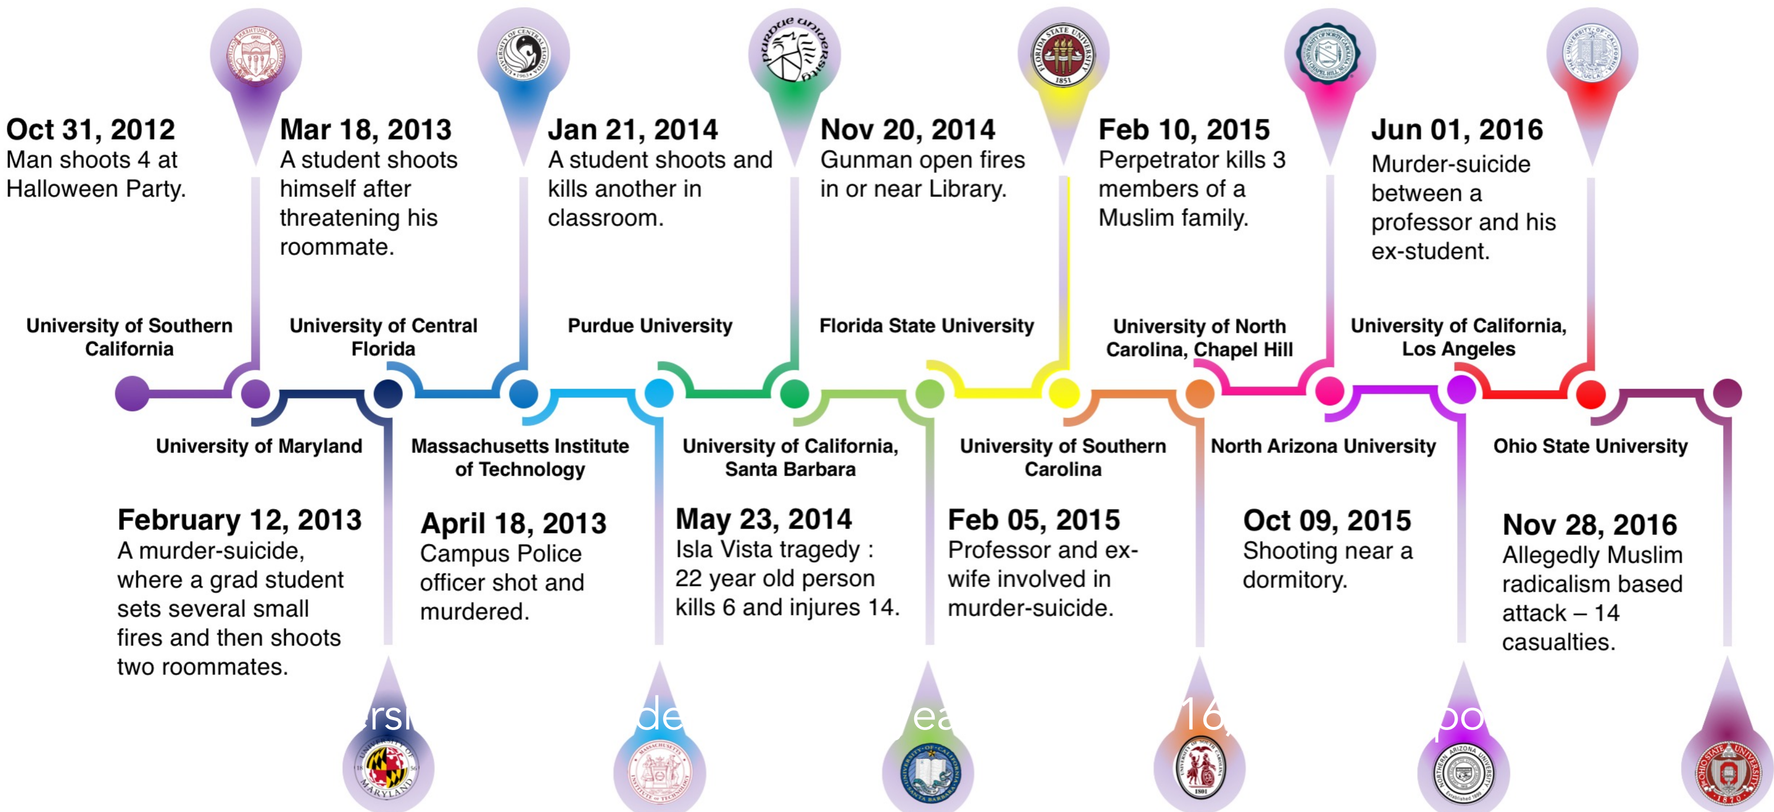

Temporal and Linguistic Patterns of Stress

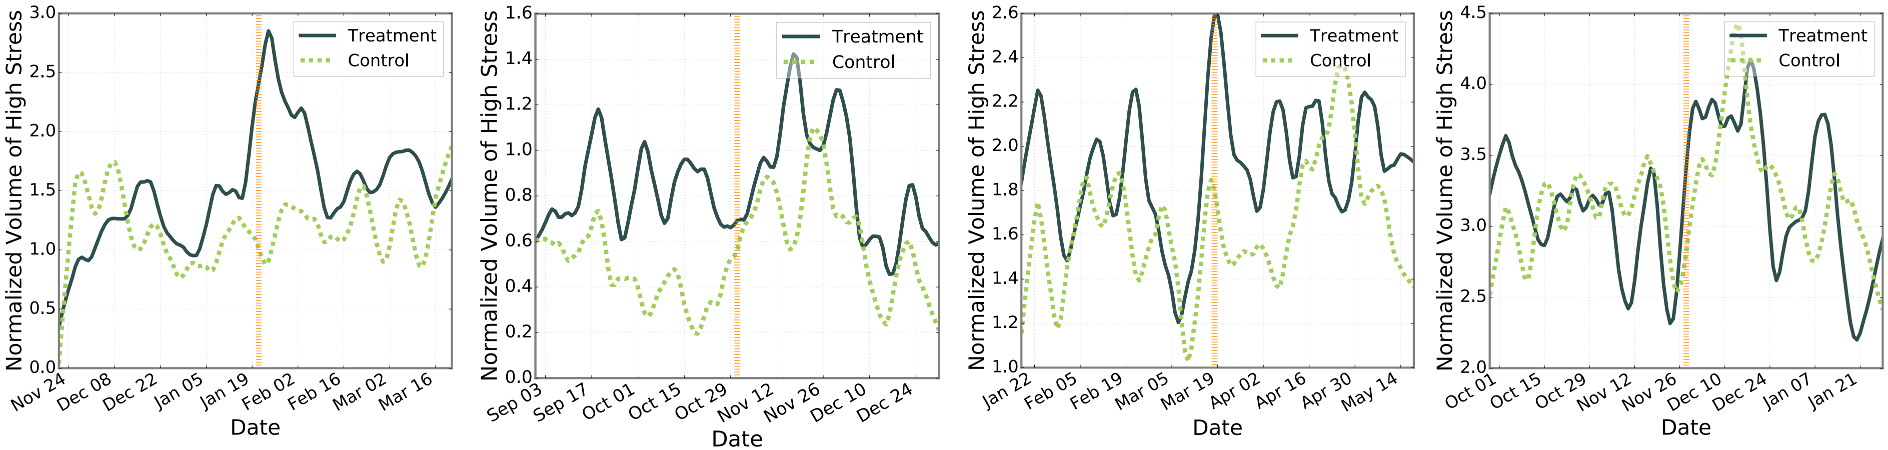

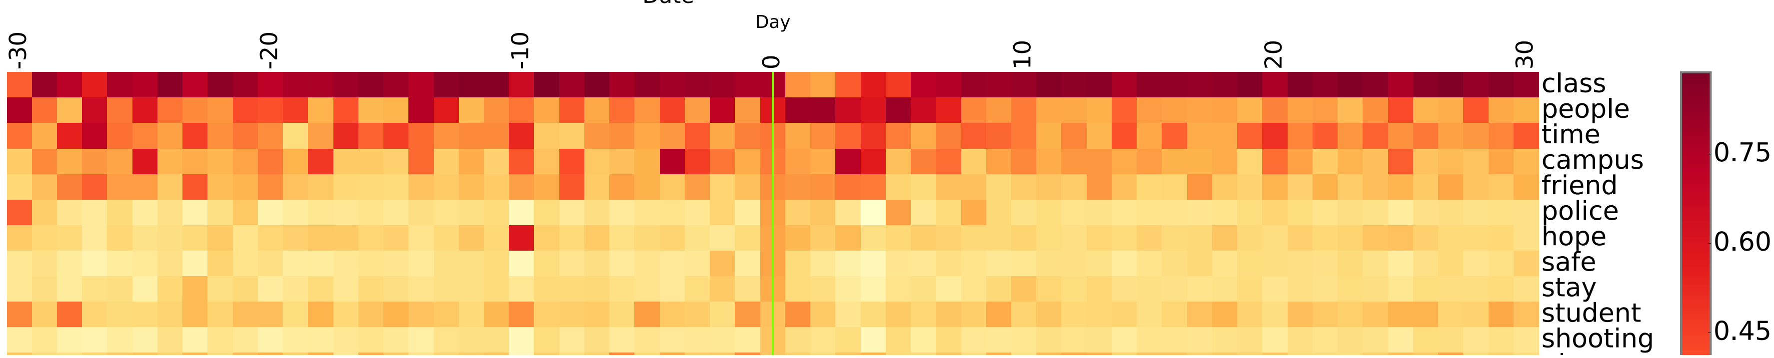

(De Choudhury, Counts, Horvitz, ICWSM 2013; WebSci 2013)

Social media depression index

actual (BRFSS data)

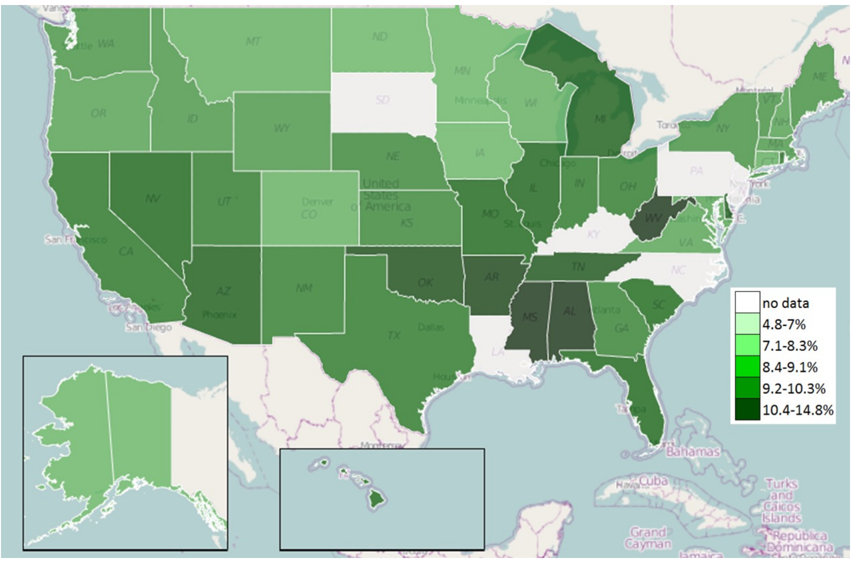

predicted (SMDI)

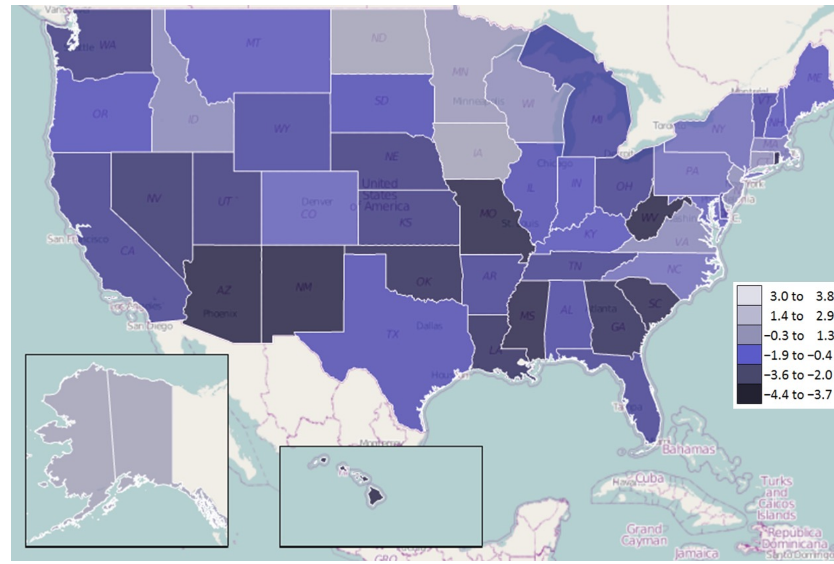

Socio-demographic, spatio-temporal patterns of prevalence of depression

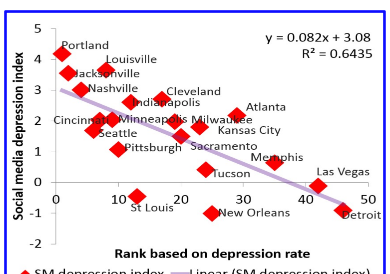

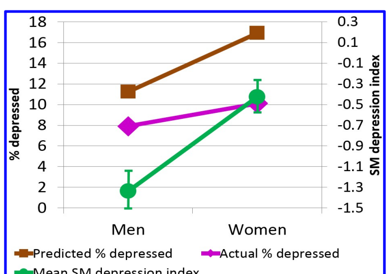

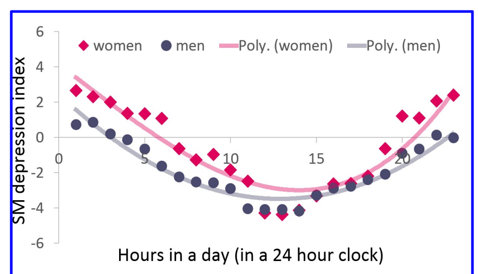

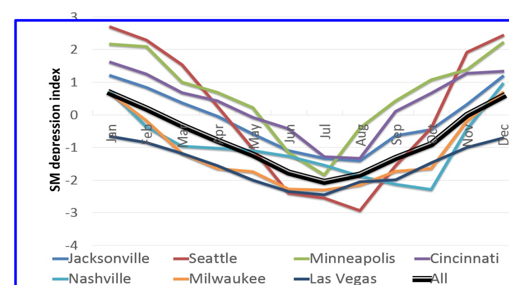

Multi-Task Learning for Mental Health using Social Media Text
Adrian Benton Johns Hopkins University adrian@cs.jhu.edu
Margaret Mitchell Microsoft Research* mmitchellai@google.com
Abstract
We introduce initial groundwork for estimating suicide risk and mental health in a deep learning framework. By modeling multiple conditions, the system learns to make predictions about suicide risk and mental health at a low false positive rate. Conditions are modeled as tasks in a multi-task learning (MTL) framework, with gender prediction as an additional auxiliary task. We demonstrate the effectiveness of multi-task learning by comparison to a well-tuned single-task baseline with the same number of parameters. Our best MTL model predicts potential suicide attempt, as well as the pre
1 Introduction
Discovering Shifts to Suicidal Ideation from Mental Health Content
Munmun De Choudhury Georgia Tech Atlanta GA 30332 munmund@gatech.edu
Emre Kiciman Microsoft Research Redmond WA
Glen Coppersmith Qntfy.io Crownsville MD, 21032 glen@qntfy.io
ABSTRACT
History of mental illness is a major factor behind suicide risk and ideation. However research efforts toward characterizing and forecasting this risk is limited due to the paucity of information regarding suicide ideation, exacerbated by the stigma of mental illness. This paper fills gaps in the literature by developing a statistical methodology to infer which individuals could undergo transitions from mental health discourse to suicidal ideation.
Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Current Opinion in Behavioral Sciences
Detecting depression and mental illness on social media: an integrative review
Sharath Chandra Guntuku[1], David B Yaden[1], Margaret L Kern[2], Lyle H Ungar[1] and Johannes C Eichstaedt[1]
Although rates of diagnosing mental illness have improved over the past few decades, many cases remain undetected. Symptoms associated with mental illness are observable on Twitter, Facebook, and web forums, and automated methods are increasingly able to detect depression and other mental illnesses. In this paper, recent studies that aimed to predict mental illness using social media are reviewed. Mentally ill users have been identified using screening surveys, their public sharing of a diagnosis on Twitter, or by their membership in an online forum, and they were distinguishable from control users by patterns in their language and online activity. Automated detection methods may help to identify depressed or otherwise at-risk individuals through the large-scale passive monitoring of social media, and in the future may complement existing screening procedures.
Automated analysis of social media potentially provides methods for early detection. If an automated process could detect elevated depression scores in a user, that individual could be targeted for a more thorough assessment, and provided with further resources, support, and treatment. Studies to date have either examined how the use of social media sites correlates with mental illness in users [3] or attempted to detect mental illness through analysis of the content created by users. This review focuses on the latter: studies aimed at predicting mental illness using social media. We first consider methods used to predict depression, and then consider four approaches that have been used in the literature. We compare the different approaches, provide direction for future studies, and consider ethical issues.
Prediction methods
Identifying Social Media Markers of ... Machine Learning and Clinical ...
Birnbaum et al
Facebook language predicts depression in medical records
Johannes C. Eichstaedt, Robert J. Smith, Raina M. Merchant, Lyle H. Ungar, Patrick Crutchley, Daniel Preoţiuc-Pietro, David A. Asch, and H. Andrew Schwartz
PNAS
Significance
Results
Natural Language Processing of Social Media as Screening for Suicide Risk
Glen Coppersmith, Ryan Leary, Patrick Crutchley and Alex Fine
SAGE
Introduction

Predicting Depression via Social Media

Summary

- Can social media activities and connectedness predict risk to major depressive disorder?
- Recruitment of a sample of Twitter users through a survey methodology over Amazon's Mechanical Turk
  - ~40% provided access to Twitter data

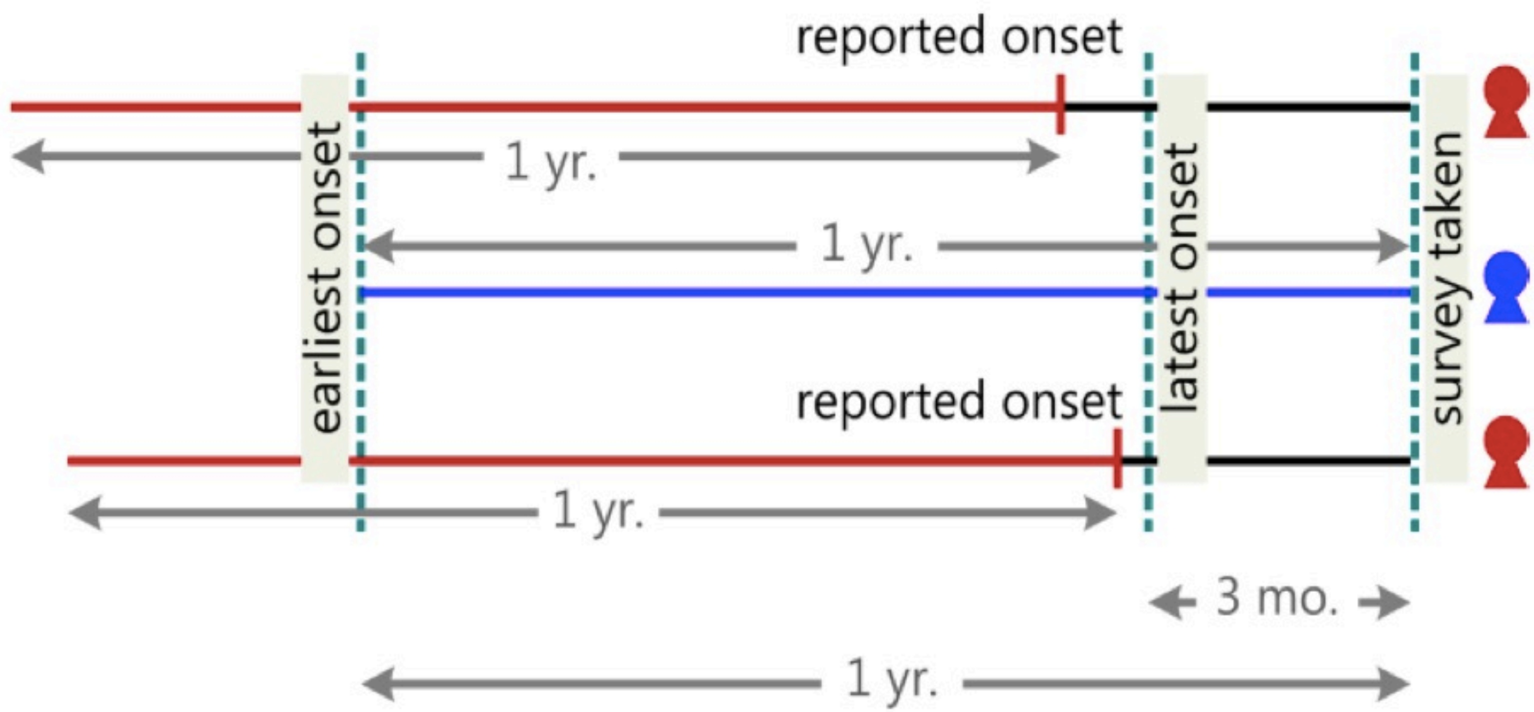

Summary

- Social engagement
- "Insomnia index" – mean *z*-score of an individual's volume of Twitter activity per hour
- Ego-centric social graph – nodal properties (*inlinks, outlinks*); dyadic properties (*reciprocity, interpersonal exchange*); neighborhood properties (*density, clustering coefficient, two-hop neighborhood, embeddedness, number of ego components*)
- Language
    - Depression lexicon – top uni- and bigrams compiled from Yahoo! Answers category on mental health
    - Linguistic style

Summary

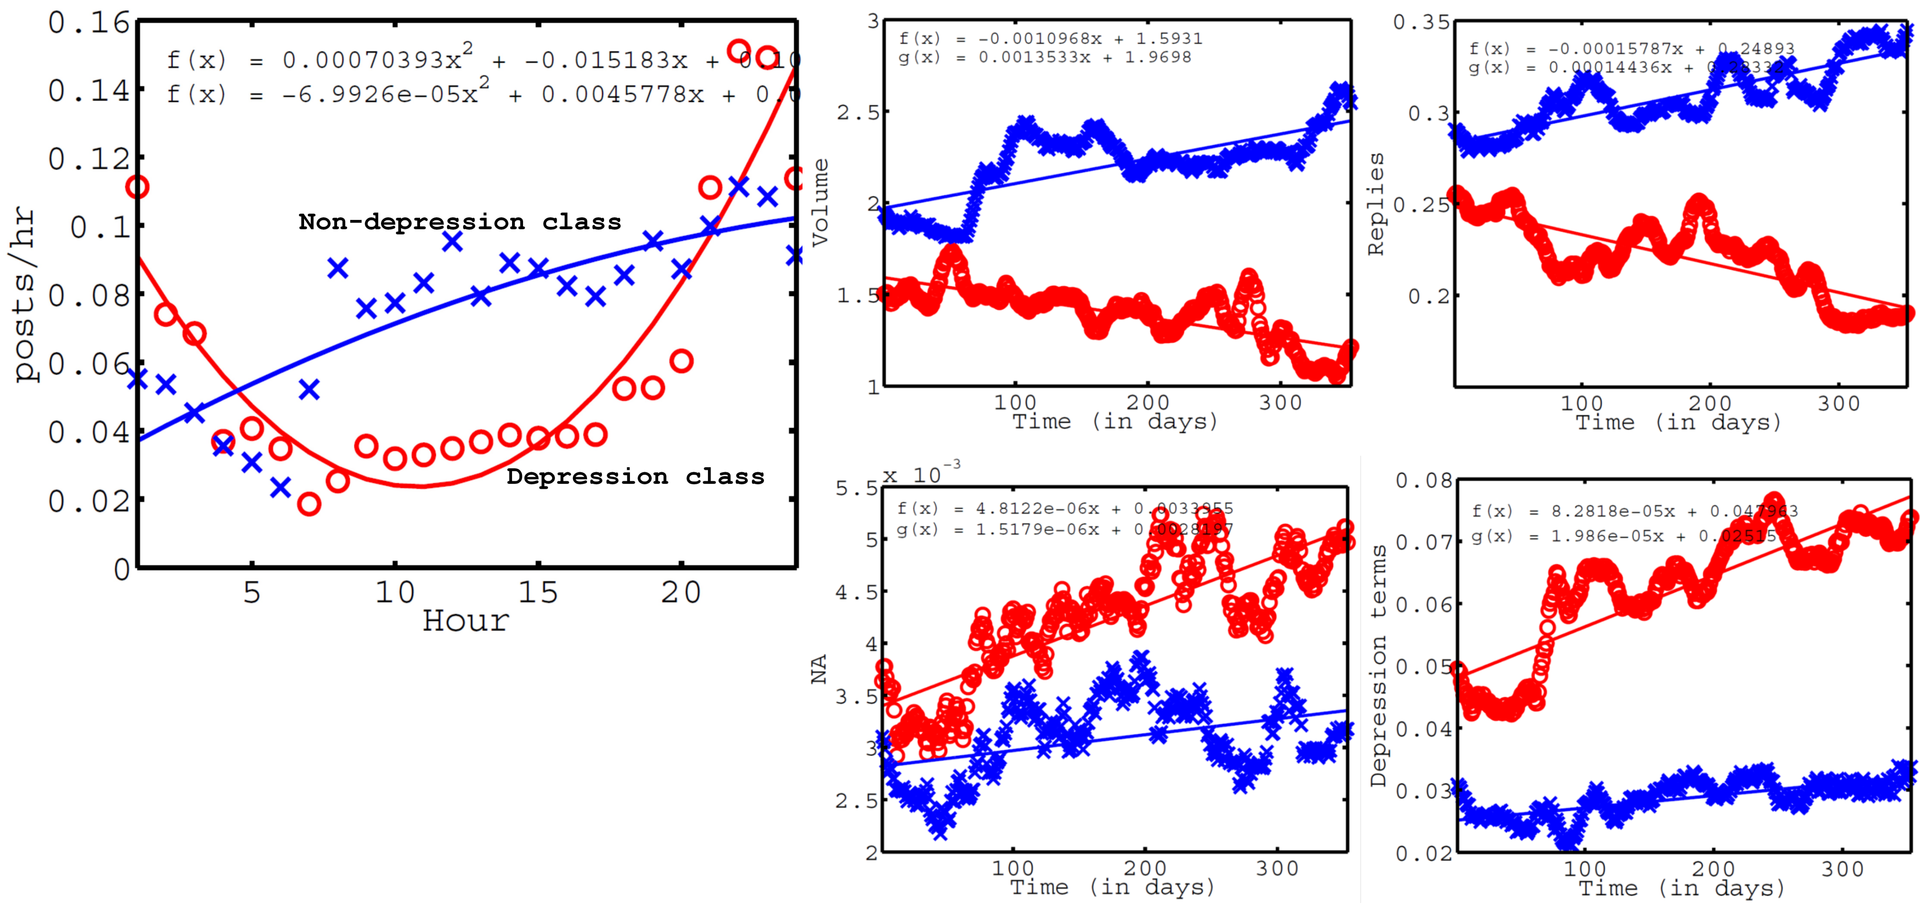

Non-depression class
Depression class
posts/hr
Hour
Volume
Replies
NA
Depression terms
Time (in days)
f(x) = -0.0010968x + 1.5931
g(x) = 0.0013533x + 1.9698
f(x) = 4.8122e-06x + 0.0033955
g(x) = 1.5179e-06x + 0.0028197
f(x) = 8.2818e-05x + 0.047963
g(x) = 1.986e-05x + 0.02515
x 10^-3

Summary

Egonetwork measures	Depres. class	Non-depres. class
#followers/inlinks	26.9 (σ=78.3)	45.32 (σ=90.74)
#followees/outlinks	19.2 (σ=52.4)	40.06 (σ=63.25)
Reciprocity	0.77 (σ=0.09)	1.364 (σ=0.186)
Prestige ratio	0.98 (σ=0.13)	0.613 (σ=0.277)
Graph density	0.01 (σ=0.03)	0.019 (σ=0.051)
Clustering coefficient	0.02 (σ=0.05)	0.011 (σ=0.072)
2-hop neighborhood	104 (σ=82.42)	198.4 (σ=110.3)
Embeddedness	0.38 (σ=0.14)	0.226 (σ=0.192)
#ego components	15.3 (σ=3.25)	7.851 (σ=6.294)

Discussion Point I

In this paper, the ground truth was obtained from Amazon mechanical turk workers. Anything unique about this population that may have affected the findings? What would be alternative ways of recruiting people or gathering high quality ground truth?

Discussion Point II

Depression is not an online condition, but one that spans both the online and the offline life. The paper does not take offline attributes into their models.

Is there a way to that into account? What would be the most significant offline attributes to consider?

PLOS | COMPUTATIONAL BIOLOGY

RESEARCH ARTICLE

Combining Search, Social Media, and Traditional Data Sources to Improve Influenza Surveillance

Mauricio Santillana[1,2,3]*, **André T. Nguyen**[1], **Mark Dredze**[4], **Michael J. Paul**[5], **Elaine O. Nsoesie**[6,7], **John S. Brownstein**[2,3]

1 Harvard School of Engineering and Applied Sciences, Cambridge, Massachusetts, United States of America, **2** Boston Children's Hospital Informatics Program, Boston, Massachusetts, United States of America, **3** Harvard Medical School, Boston, Massachusetts, United States of America, **4** Department of Computer Science, Johns Hopkins University, Baltimore, Maryland, United States of America, **5** Department of Information Science, University of Colorado, Boulder, Colorado, United States of America, **6** Department of Global Health, University of Washington, Seattle, Washington, United States of America, **7** Institute for Health Metrics and Evaluation, Seattle, Washington, United States of America

* msantill@fas.harvard.edu

Citation: Santillana M, Nguyen AT, Dredze M, Paul MJ, Nsoesie EO, Brownstein JS (2015) Combining Search, Social Media, and Traditional Data Sources to Improve Influenza Surveillance. PLoS Comput Biol 11(10): e1004513. doi:10.1371/journal.pcbi.1004513

Abstract

We present a machine learning-based methodology capable of providing real-time ("now-cast") and forecast estimates of influenza activity in the US by leveraging data from multiple data sources including: Google searches, Twitter microblogs, nearly real-time hospital visit records, and data from a participatory surveillance system. Our main contribution consists of combining multiple influenza-like illnesses (ILI) activity estimates, generated indepen-

Original Investigation | Health Informatics

December 23, 2020

Development of a Machine Learning Model Using Multiple, Heterogeneous Data Sources to Estimate Weekly US Suicide Fatalities

Daejin Choi, PhD[1]; Steven A. Sumner, MD[2]; Kristin M. Holland, PhD[3]; John Draper, PhD[4]; Sean Murphy, PhD[4]; Daniel A. Bowen, MPH[3]; Marissa Zwald, PhD[3]; Jing Wang, MD[3]; Royal Law, PhD[5]; Jordan Taylor, BS[6]; Chaitanya Konjeti, BS[6]; Munmun De Choudhury, PhD[6]

» Author Affiliations | Article Information

JAMA Netw Open. 2020;3(12):e2030932. doi:10.1001/jamanetworkopen.2020.30932

Key Points

Question Can real-time streams of secondary information related to suicide be used to accurately estimate suicide fatalities in the US in real time?

Discussion Point III

But are models trained on aggregated group-level differences useful at the individual level?

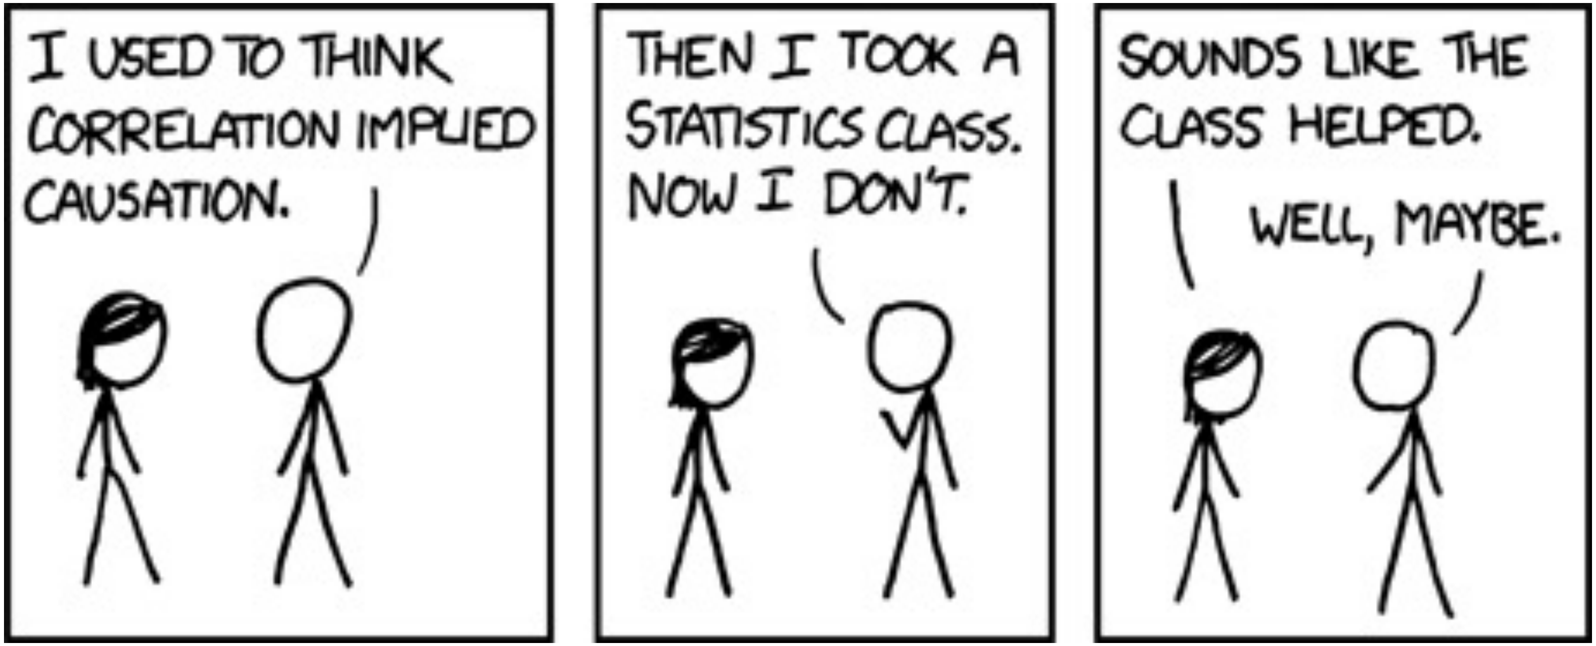

Correlation and causation

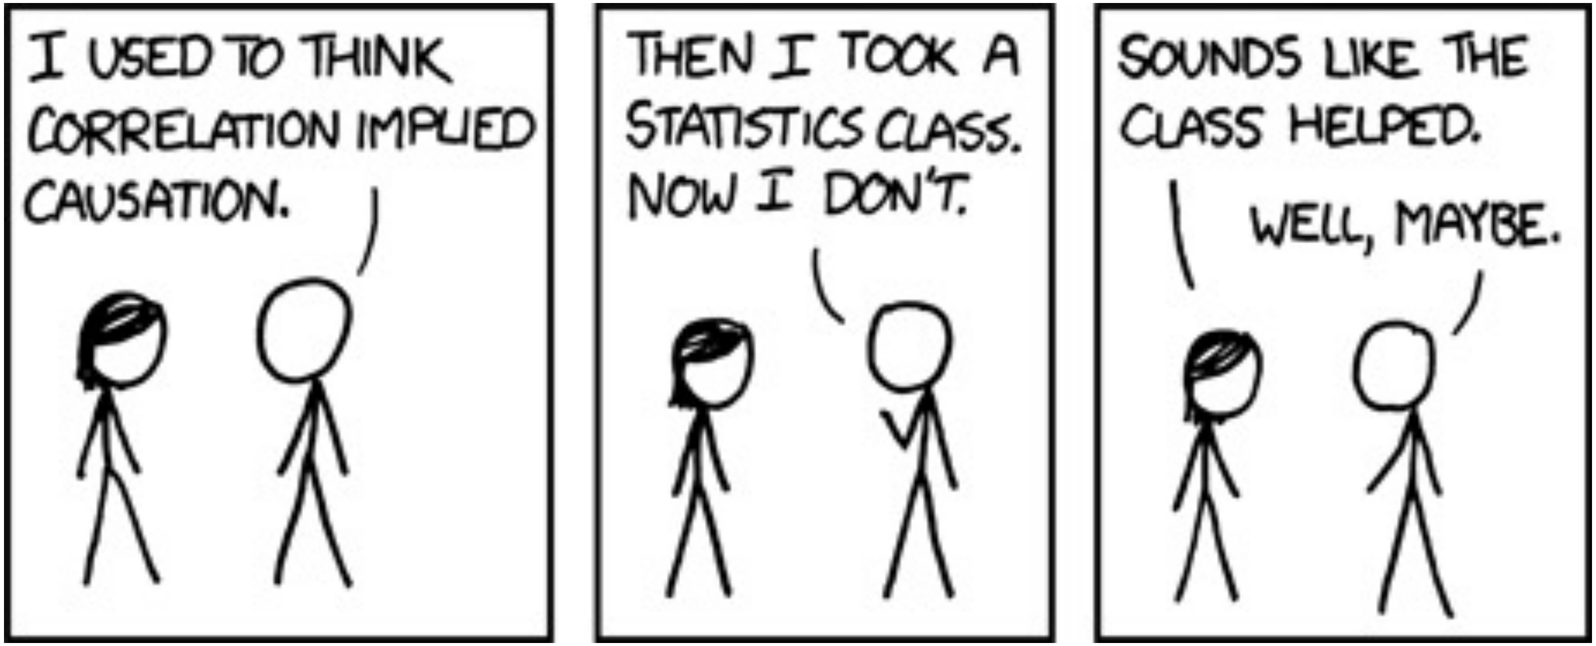
I USED TO THINK CORRELATION IMPLIED CAUSATION.
THEN I TOOK A STATISTICS CLASS. NOW I DON'T.
SOUNDS LIKE THE CLASS HELPED.
WELL, MAYBE.

What comes next?

What comes next?

Social Media + Machine Learning for clinical interventions

Efficacy

Validity

SOCIAL MEDIA + MACHINE LEARNING

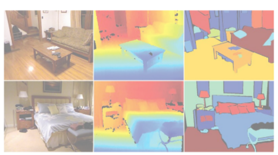

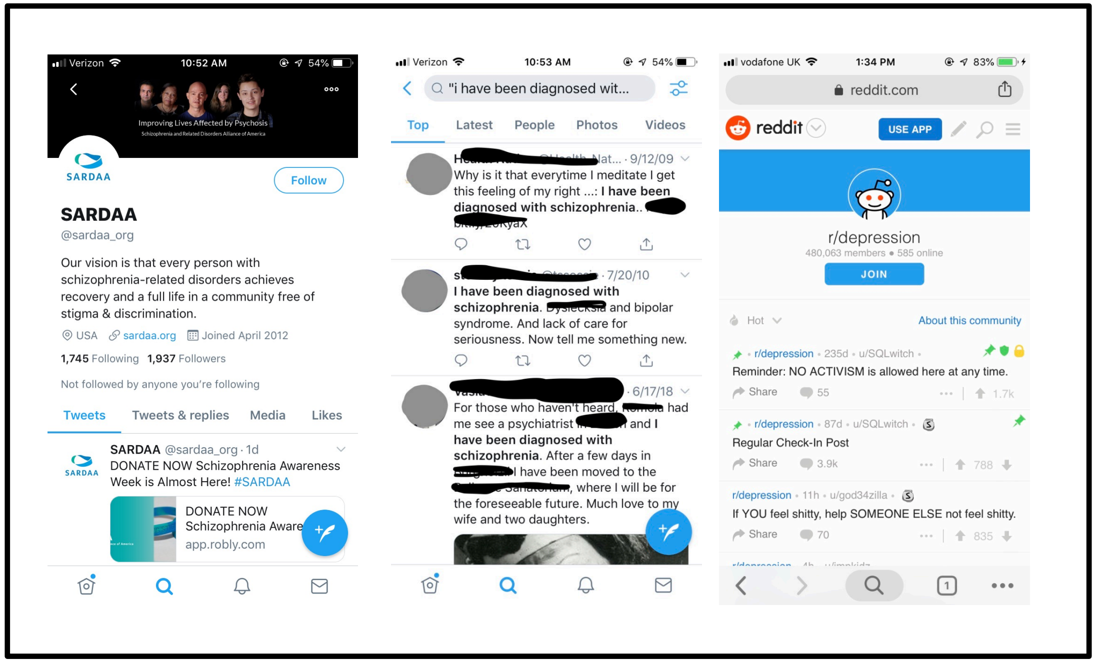

Ground truth label:
Readily available

Proxy ground truth labels

Ground truth label:
clinical assessment

Construct Validity: Do the proxy diagnostic signals **objectively and accurately measure what they claim to measure** (clinical mental illness diagnosis)

Theoretical/Clinical grounding: Is what is being measured by the proxy diagnostic signals

valid in itself?

Proxy data sets: diagnostic signals for schizophrenia on Twitter

Affiliation Data

N = 861

Self-reports Data

N = 412

Appraised Data

N = 153

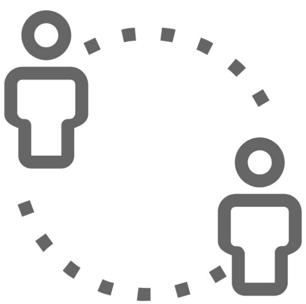

Matched Control Data

N = 640

Patient's social media data

Schizophrenia Patient Data

N = 88

Healthy Control Data

N = 55

Methodology: Triangulation

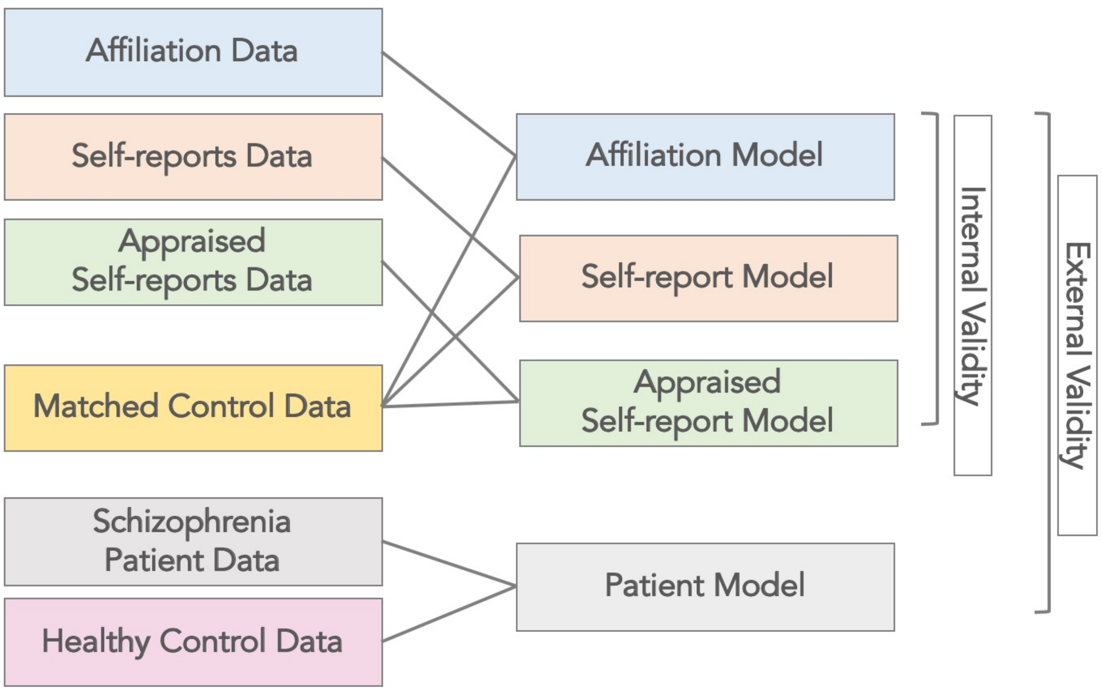

Binary classification task:

Distinguishing those with schizophrenia from control populations

Efficacy

High internal validity
Very low external validity

	Cross Validation	Testing on patient data
Affiliation Model	0.89	0.21
Self-report Model	0.72	0.48
Appraised Model	0.80	0.55
Patient Model	0.72	0.76

Issues with Construct Validity

Affiliation	*β*	*Appraised*	*β*	*Patient*	*β*
i'm	-0.825	*NegAffect*	0.063	*cog mech*	-0.003
stigma	0.665	negation	0.074	*present*	-0.002
mhchat	0.696	*present*	0.40	*body*	-0.002
body	0.729	help	0.401	*verbs*	-0.002
bipolar	0.774	thought	0.41	*social*	-0.002
work	0.919	i'm	0.44	*aux verbs*	-0.002
self	0.961	die	0.45	help	0.0002
social	1.109	alone	0.45	feeling	0.001
care	1.111	hard	0.457	i'm	0.002
depression	1.116	cry	0.50	gonna	0.002
suicide	1.133	*body*	0.52	angel	0.002
thanks	1.445	feeling	0.523	burning	0.002
illness	1.447	*verbs*	0.58	pray	0.003
help	1.632	sorry	0.662	lifetime	0.005
mental health	1.866	gonna	0.63	attack	0.006

Main Takeaway

If the broader research agenda is to use social media data to inform clinical decision-making, such as early diagnosis, treatment or patient-provider interventions, **(social media) data collection and machine learning model development should happen in context**

Class Exercise

Describe a design idea where we can use social media based depression (or other mental health condition like schizophrenia) predictors to help people. How would it negotiate privacy and ethical issues?

Improving “Blanket” Interventions

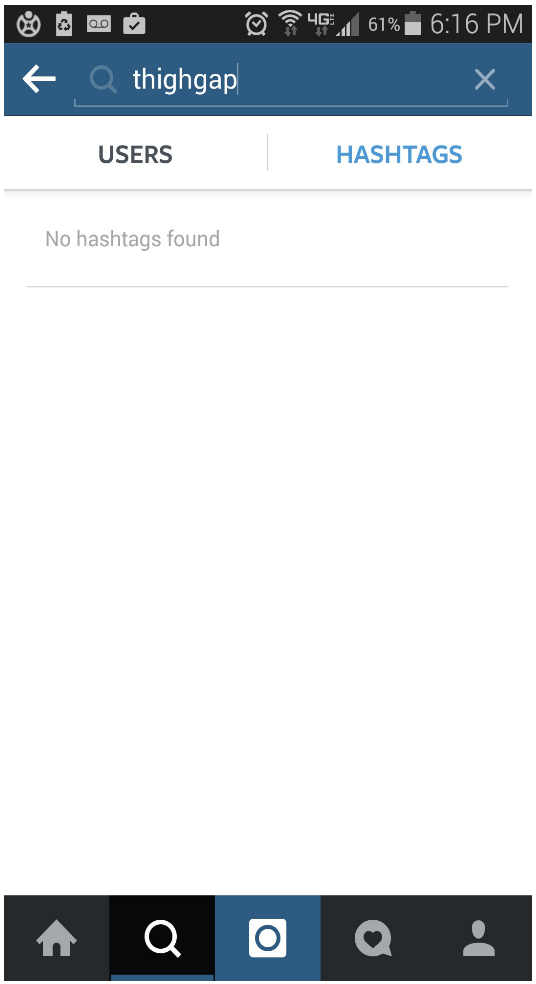

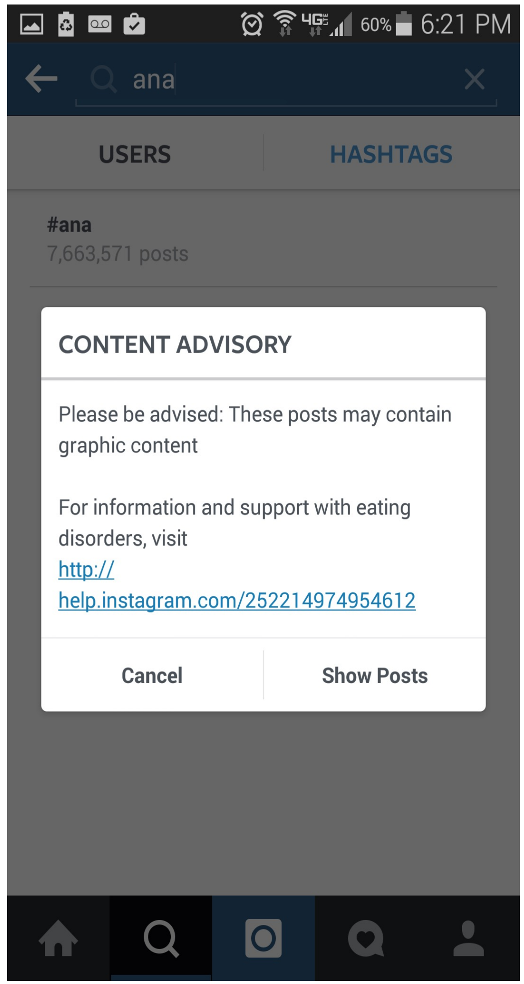

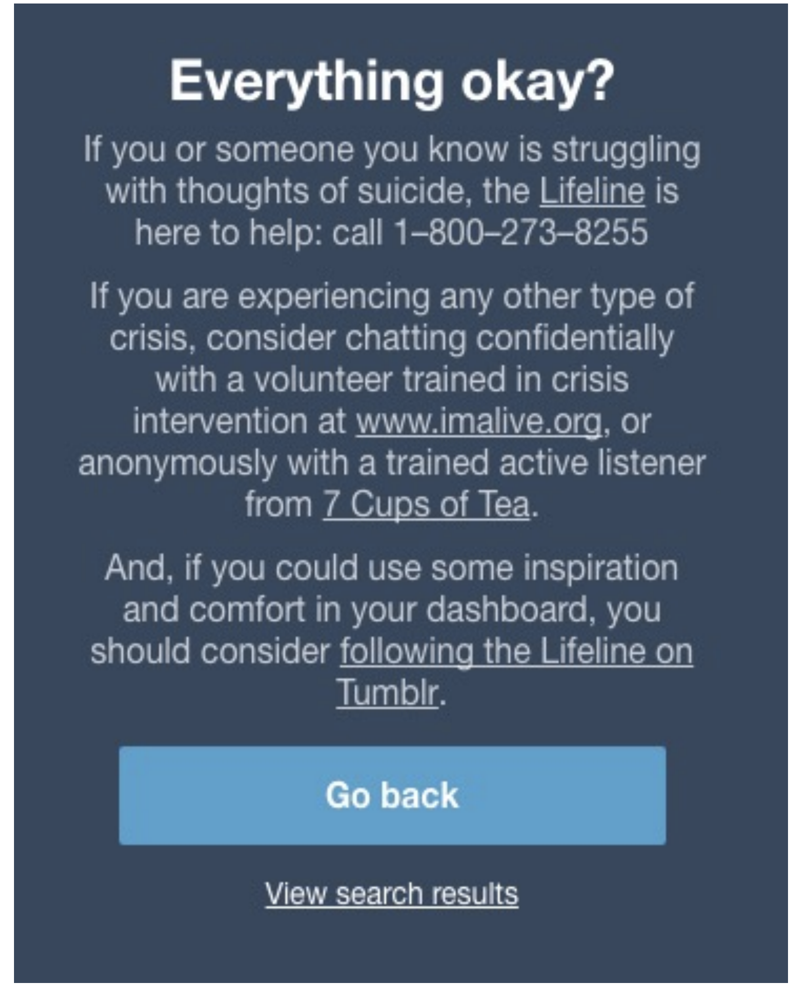


Google
suicide
Web
News
Images
Videos
Books
More
About 214,000,000 results (0.44 seconds)
Need help? United States:
1 (800) 273-8255
National Suicide Prevention Lifeline
Hours: 24 hours, 7 days a week
Languages: English, Spanish
Website: www.suicidepreventionlifeline.org

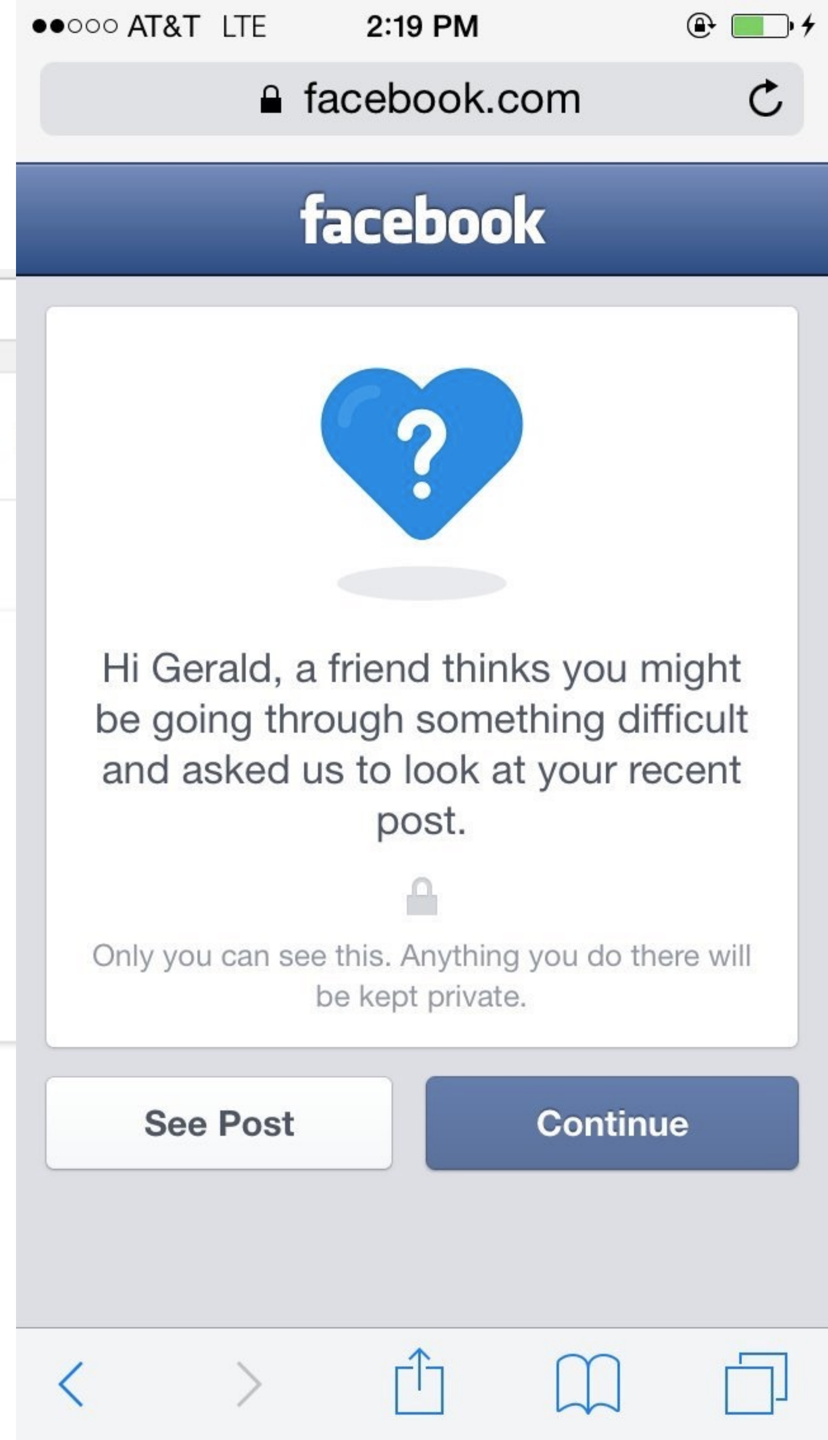
AT&T LTE
2:19 PM
facebook.com
facebook
?
Hi Gerald, a friend thinks you might be going through something difficult and asked us to look at your recent post.
Only you can see this. Anything you do there will be kept private.
See Post
Continue

A Taxonomy of Ethical Tensions in Inferring Mental Health States from Social Media

Stevie Chancellor
Georgia Tech
Atlanta, GA, US
schancellor3@gatech.edu

Michael L Birnbaum
Northwell Health
Glen Oaks, NY, US
mbirnbaum@northwell.edu

Eric D. Caine
University of Rochester
Rochester, NY, US
Eric_Caine@urmc.rochester.edu

Vincent M. B. Silenzio
University of Rochester
Rochester, NY, US
vincent.silenzio@rochester.edu

Munmun De Choudhury
Georgia Tech
Atlanta, GA, US
munmund@gatech.edu

ABSTRACT

Powered by machine learning techniques, social media provides an unobtrusive lens into individual behaviors, emotions, and psychological states. Recent research has successfully employed social media data to predict mental health states of individuals, ranging from the presence and severity of mental disorders like depression to the risk of suicide. These algorithmic inferences hold great potential in supporting early detection and treatment of mental disorders and in the design of interventions. At the same time, the outcomes of this research can pose great risks to individuals, such as issues of incorrect, opaque algorithmic predictions, involvement of bad or unaccountable actors, and potential biases from intentional or inadvertent misuse of insights. Amplifying these tensions, there are also divergent and sometimes inconsistent methodological gaps and under-explored ethics and privacy dimensions. This paper presents a taxonomy of these concerns and ethical challenges, drawing from existing literature, and poses questions to be resolved as this research gains traction. We identify three areas of tension: ethics committees and the gap of social media research; questions of validity, data, and machine learning; and implications of this research for key stakeholders. We conclude with calls to action to begin resolving these interdisciplinary dilemmas.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing**; **Social media**; • **Applied computing** → *Psychology*;

Conference on Fairness, Accountability, and Transparency (FAT '19)*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3287560.3287587

1 INTRODUCTION

Last year, Facebook unveiled automated tools to identify individuals contemplating suicide or self-injury [75, 62]. The company claims that they "use pattern recognition technology to help identify posts and live streams as likely to be expressing thoughts of suicide," which then can deploy resources to assist the person in crisis [75]. Reactions to Facebook's suicide prevention artificial intelligence (AI) are mixed, with some concerned about the use of AI to detect suicidal ideation as well as potential privacy violations [86]. Other suicide prevention AIs, however, have been met with stronger public backlash. Samaritan's Radar, an app that scanned a person's friends for concerning Twitter posts, was pulled from production, citing concerns for data collection without user permission [54], as well as enabling harassers to intervene when someone was vulnerable [4].

Since 2013, a new area of research has incorporated techniques from machine learning, natural language processing, and clinical psychology to categorize individuals' moods and expressed well-being from social media data. These algorithms are powerful enough to infer with high accuracy whether an individual might be suffering from disorders such as major depression [28, 19, 84, 73, 78], postpartum depression [26, 27], post-traumatic stress [21], schizophrenia [60, 6], and suicidality [15, 22]. These algorithms can also reveal symptomatology linked to psychiatric challenges,