Research Statement: Johannes C. Eichstaedt

I am a computational psychologist studying the mechanisms that give rise to mental and physical illness and well-being. In my work, I seek to elucidate these mechanisms by analyzing digital text using Natural Language Processing (NLP), machine learning, and artificial intelligence (AI). Digital text arises in many communication (e.g., SMS, chat, email) and medical contexts (such as in clinical notes). Foremost among these data sources in terms of size, with billions of users, social media has become the largest cross-sectional and longitudinal dataset on human behavior in history.

Digital natural language is a complex phenotype that encodes cognitions, emotions, and behaviors in ecological contexts. At the individual level, I leverage this novel data stream to unobtrusively measure mental health and well-being (e.g., depression, stress, personality, and life satisfaction) and use it as a data-driven lens to characterize the behavioral and social mechanisms that impact health outcomes. At the community level, geo-aggregated social media language allows for understanding the psychological context of communities and social determinants of health, such as environmental stress and social capital. Through enabling the unobtrusive measurement of community well-being (including life satisfaction and positive and negative affect), my work points to its importance as a high-level marker of social health and functioning. For example, well-being components are associated with heart disease mortality [1], impacted by traumatic shocks caused by police brutality and racism [2], and low well-being can drive electoral choice towards populists [3].

To leverage the new generation of population-scale social media data sets, I co-developed a methodological toolkit that combines theory from the social sciences with NLP and AI. In 2011, I co-founded the World Well-Being Project at the University of Pennsylvania. This lab, which is now an inter-institutional consortium of labs, pioneered NLP methods (such as *differential language analysis*) [4, 5, 6] to measure and characterize psychological and social processes through digital text. As of 2023, I have produced over 60 papers and procured $4.9m+ in external funding with the World Well-Being Project. The methods we created in this consortium have substantively contributed to psychology, medicine, public health, and public policy, beyond the methodological advances in NLP and machine learning.

My work spans the individual (§1) and community level (§2), which structure the summary of my work in this statement for organizational purposes. These are followed by sections about ongoing grants with a population health focus (§3) and future directions (§4).

§1 Using Social Media to Measure Individual-level Mental Health and Psychology

The analysis of natural language allows for the tracking of affective states over time (§1.1) and can serve as an early warning system for the onset of mood disorders, such as depression (§1.2). Beyond prediction, the fine-grained language profile associated with depression can pinpoint behavioral risk factors and enable symptom dashboards (§1.3). Beyond states, psychological traits (such as Big Five Personality) can be assessed unobtrusively from subjects' language (§1.4, §1.5). Beyond social media, many sources of natural language are suitable for psychological assessment through NLP (§1.6).

§1.1 Tracking states: Machine learning classifiers can track the weekly emotions of Facebook users across many months. NLP and machine-learning methods allow for tracking within-person changes in mood over time in valence and arousal (the affective circumplex) based on social media posts. Specifically, we have built a valence and arousal classifier from 2,895 annotated Facebook statuses (with reliabilities of $r = .77$ and $r = .83$, respectively) [8] and applied it to 303,575 Facebook statuses of 640 Facebook posters who had also self-reported their personality. This allowed us to infer users' moods across 17,937 weeks, generating the largest dataset of its kind to study affect dynamics [9] (links to animations: 21-y.o. woman, 24-y.o. man). Studying these week-to-week transitions across all users, we found, for example, that men have a narrower "set point" for affect than women, whose natural equilibrium encompasses a broader range of valence and arousal–suggesting higher emotional flexibility. This work points to how unobtrusive NLP methods can track affect dynamics across relatively short time scales (e.g., in response to current events). Beyond general affect, the NLP and ML/AI-based methods can also be used to track the onset of psychopathology, such as mood disorders.

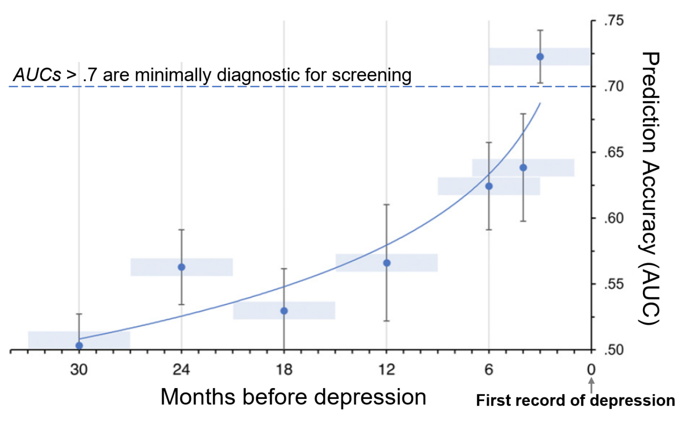

Figure 1 (§1.2): Out-of-sample prediction accuracies (*AUCs*) of future depression status based on different 6-month windows of Facebook statuses. Clinically relevant predictions are possible 3 months prior to the first formal record of depression [7].

§1.2 Depression can be predicted from social media before it appears in medical records. The public health and medical communities need scalable approaches to identify mental illness in ways that preserve privacy (see [10] for our review). I lead the first and only study of its kind to date, in which we approached 11,224 patients in an Emergency Department of a large urban hospital, allowing us to build a sample of N = 683 patients for which we linked the language used in their Facebook status histories with diagnoses of depression in their medical records. We demonstrated that machine-learning techniques, when applied to natural language from Facebook posts, can provide a rough screening system for future depression up to three months before the first recorded diagnosis (predictive accuracy: AUC =.72; see **Fig. 1**) [7]. Language patterns associated with depression include depressed affect (*tears, sad, feelings*), somatic complaints (*pain, hurt*), loneliness (*miss*) (see **Fig. 2.**), in addition to cognitive processes (*preoccupation with the self, rumination*) and hostility (see [7]).

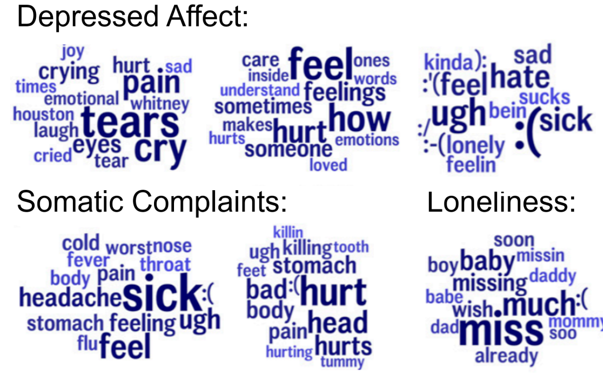

Figure 2 (§1.2): Six of the language topics most predictive of future depression status. Latent Dirichlet Allocation topics are semantically coherent clusters of words based on co-occurrence. Words are sized by prevalence within topic [7].

Our study established for the first time that social media-based screening is feasible in real-life patient populations and, in principle, can yield a "symptom dashboard" for clinicians. Patients can consent to share access to their social media feeds or phone keyboards, and algorithms can automatically derive real-time indicators of sadness, loneliness, sleep disturbance, and others, to which physicians can respond with specific treatment suggestions. In a study extending this work, we showed that 18 out of 21 major health conditions in medical records were predictable from patients' Facebook language over and above demographic baselines [11].

§1.3 Language analyses reveal fine-grained nuances about the cognitions and processes behind depression. Beyond using language features to measure a construct through machine learning, open-vocabulary exploratory language analyses, such as *differential language analysis*, can be used to characterize psychological constructs [4]. Extending the language patterns observed in the clinical study reported in §1.2 and **Fig. 2**, in a large dataset of N = 15,034 Facebook users, I differentiated the low-mood and low self-worth facets of survey-measured depression by correlating them individually with 19,389 words and phrase frequencies. **Fig. 3** shows the most associated words and phrases, sized by association coefficients (controlling for age, gender, hostility, and multiple comparisons).

Figure 3 (§1.3): Correlational word clouds. Words and phrases in Facebook statuses most associated with **top:** *low mood* (4 survey items) and **bot.:** *low self worth* (3 items) across a sample of 15,034 consenting Facebook users who shared their statuses and took the survey. Word sizes indicate association coefficient (β) with survey construct (controlled for age and gender), color indexes frequency (from grey: rare to red: frequent). All correlations significant at $p < .01$, adjusted for multiple comparisons.

While none of the questions asked about loneliness, feeling *alone* emerges as the strongest language feature predicting low mood (but not mentions of sadness). Similarly, asking *why* is the strongest language marker for low self-worth, which dovetails with accounts of depression as a loss of meaning. In addition to markers of typical depression symptoms (*sleep, death, suicide, can't remember, pain*), words and phrases suggesting self-negation (*I don't, I'm not*) and hedging (*apparently, probably, maybe*) suggest the loss of self-efficacy. This example showcases the potential of exploratory language analyses to identify the cognitions and behaviors that covary with a psychological construct (depression) in an ecological setting (Facebook) assessed through a minimal survey (7 questions about depression) to inform and build theory [12].

§1.4 Many psychological traits can be reliably estimated from language. Across 20+ studies, we have observed that most traits can be measured with decent accuracy from autobiographical language using machine learning; such assessments require about 500 to 1,000 words of writing from subjects [4, 5]. Using language-based models, we have successfully estimated age [13, 14] and gender [14, 15], measured personality [16, 17, 18], and more subtle constructs, such as temporal orientation–the extent to which people orient towards the future vs. the past, which is associated with better education, health, and income [19, 20]. Beyond mental health, we can measure a wide variety of psychological processes that, in turn, are risk and protective factors. For example, we have measured stress [21], and how much subjects embrace an internal (vs. external) locus of control [22].

§1.5 Language models can estimate Big Five personality as well as friends can. As one case study, 67,477 Facebook users took Big Five personality surveys and shared access to their Facebook statuses [16]. The words, phrases, and topics in their Facebook statuses were extracted, and a machine learning model used them to predict survey personality out-of-sample in a cross-validated framework, which ensures that the model generalizes well to new data (e.g., [23, 24]). We compared these prediction accuracies to those obtained by informants who reported their impression of a friend's personality via a personality survey [16]. As can be seen in **Fig. 4**, personality can be predicted from Facebook posts about as well or better than a friend can estimate it for the same person.

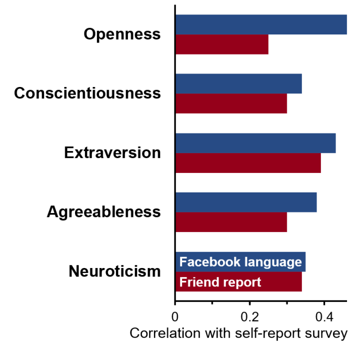

Figure 4 (§1.5): Out-of-sample accuracies for prediction models based on Facebook statuses (blue) compared to friend report (red), benchmarked against self report (ground truth); training N = 66,732, test N = 745 [16].

§1.6 Many sources of language can be used for measurement (e.g., SMS, keyboard logs, and clinical notes). While social media data collection is comparatively straightforward, natural language arises in many other life contexts, including automatically-transcribed spoken language, clinician notes, and communication systems such as email and Slack. We have found a wide variety of language sources suitable for psychological assessment [25, 26], including the text entered on phone keyboards across all apps [27]. One of the advantages of language is its versatility: language models tend to work across many types and genres of language. For example, prediction models for depression can be trained on Facebook statuses and then applied to transcripts of oral histories, and predict symptoms severity with no almost no performance degradation [27].

§2 Community-level Assessment of Health and Well-Being

Measuring psychosocial factors for communities and populations (e.g., counties or countries) has been limited in traditional survey designs due to cost. Population surveys (such as those carried out by Gallup) are generally limited to a small number of constructs, and annual estimates are not available for small geographies like counties. Social-media-based estimates can fill this gap (§2.1), allow for the characterization of county population psychological risk factors for physical illness (like heart disease) (§2.2), and provide measures of population well-being to evaluate policy (§2.3). These advances suggest the potential for such unobtrusive measurement to measure the psychological impact of shocks and current events on the population, such as those observed after police killings (§2.4).

§2.1 Twitter allows for county-level estimation of mental health, well-being, and psychological traits. We have developed a geolocation pipeline [28] to determine the US counties of origin for 1.5 billion tweets shared between 2009 and 2015 (see Github), with data processing ongoing for county-level data sets for 2019, 2020, and 2021 covering 5+ million Americans. As discussed in §3.1, while signal from random Twitter samples is noisy, we have developed a pipeline to derive stable county-level language estimates [29] with user-level post-stratification weights to adjust sample demographics for representativeness [30]. Once supervised language-based assessment models have been built at the individual level, as described in §1.1–§1.2 and §1.4–§1.5, these models can be applied to the stabilized county-level language profiles. This allows for the estimation of the expression/prevalence of (in principle, unlimited) psychological constructs cost-effectively for counties without the need for additional self-report data, opening the door for a wide variety of applications in public health and public policy.

§2.2 More people die from heart disease in hostile and stressed communities, as measured via Twitter. Heart disease is known to be impacted by psychological risk factors which have proven hard to measure with geographic precision. Using county-level Twitter data, we showed that machine learning models could predict heart disease mortality rates better than models combining the ten leading risk factors (including income, education, smoking, diabetes, and hypertension), suggesting that the heart-disease-related variance in Twitter language contains additional psychosocial information not contained in the traditional predictors (see **Fig. 5**) [1], replicated in [31].

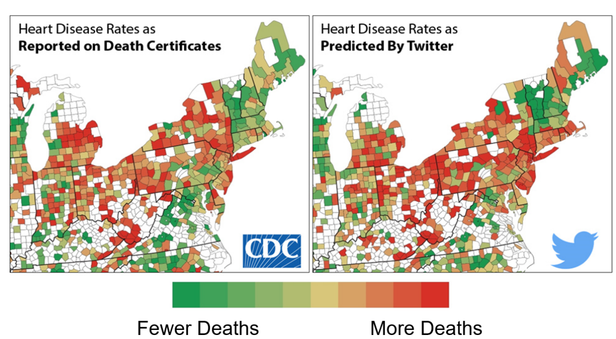

Figure 5 (§2.2): Map of counties showing age-adjusted mortality from atherosclerotic heart disease **left:** as reported by the CDC and **right:** as estimated out-of-sample through Twitter in a cross-validation framework [1].

Specifically, we have shown that the affective profile of US counties on Twitter (especially engagement and excitement vs. disengagement and boredom) is associated with heart disease mortality (reduced vs. increased, respectively), even when controlling for income and education. Similarly, expressions of hostility and interpersonal tension are psychological risk factors, while positive social experiences and optimism are protective factors. These findings are ecological and observational but dovetail with the heart disease literature.

The people posting on Twitter are not the ones dying from heart disease. Instead, these findings suggest that beyond socioeconomic indicators, communities have psychological characteristics (norms, social capital and connectedness, and environmental stress [21]) that substantively impact physical health risk. Mining social media can capture critical population psychological variables that are otherwise prohibitively costly to measure, suggesting that social media is a novel data stream for social epidemiology.

§2.3 Twitter can estimate community life satisfaction with higher external validity than phone surveys. Governments and NGOs (such as the OECD) are advocating for the measurement of population subjective well-being.

We have applied well-being prediction models trained on individual-level survey data to 1.5 billion geotagged tweets to accurately estimate *life satisfaction* and positive and negative affect when compared against estimates derived from 2.1 million Gallup surveys (r = .51 to r = 62) [32]. We have shown that the Twitter-based estimates of life satisfaction correlate more highly with external economic and health criterion variables (incl. mentally unhealthy days, income, and all-cause mortality) than Gallup survey estimates. This pattern of results suggests that the machine learning de-noises and stabilizes the Twitter estimates, and that the totality of language behaviors from a community contains more information about the community's underlying well-being than survey responses collected from respondents at a single time point. While the costly Gallup well-being tracking survey effort was effectively terminated in 2016, the cost-effective Twitter-based measurement is ongoing and can track ongoing trends in population mental health (see §3.1).

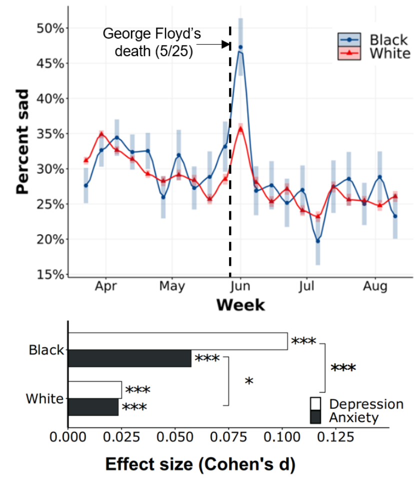

Figure 6 (§2.4) **Top:** Weekly percentage of Black and White Americans reporting having experienced sadness 'a lot of the day yesterday' in the Gallup Panel COVID-19 data set. **Bot.:** Increases in anxiety and depression symptom endorsements (on PHQ-2 and GAD-2) as standardized effect sizes comparing the five weeks before Floyd's murder to the week following his murder using data from the US Census Household Pulse survey. *p < 0.05; ***p < 0.001; adapted from [2].

§2.4 Police killings cause traumatic shocks to Black mental health. The killing of George Floyd on May 25, 2020, was captured in COVID-related data collection efforts by both Gallup and the US Census. Using Gallup data (n = 47,355), we showed that anger and sadness in the population rose to levels not seen in the last 15 years of Gallup data collection. This was particularly true for Black Americans, roughly half of whom reported these emotions on the days following Floyd's murder (**Fig. 6, top**), [2]. Using the Census data (n = 319,471) (**Fig. 6, bot.**), we estimated that an additional 900,000 Black Americans would have screened positive for depression following Floyd's murder, associated with a burden of roughly 2.7 million to 6.3 million mentally unhealthy days. This work demonstrates that the vicarious trauma experienced through police brutality subjects Black Americans to an additional mental health burden, which both reflects and reinforces the reality of inequality and structural racism in the US. Beyond static conceptions of mental health, our findings suggest that racial police killings represent a dynamic shock to the well-being of the population and to that of Black Americans in particular. This also stresses the importance of measurement technologies for population mental health with a sufficient temporal resolution to capture dynamics shocks–such as unobtrusive measurement through social media.

§3 Ongoing Grants with Population Health Applications

While social media is a promising source for the unobtrusive and cost-effective monitoring of population health, particularly early work yielded noisy time series estimates, often distorted by bots, with implausible spikes and discontinuities. In currently funded work, I co-lead projects to stabilize longitudinal estimates from Twitter to make them reliable predictors and outcome measures for population health applications (§3.1) and demonstrate their practical usefulness as an early warning system for the opioid epidemic (§3.2).

§3.1 R01: Enabling quasi-experimental methods at the county-level with weighted sampling and digital cohort designs for mental health surveillance. In a 2021 R01 Smart and Connected Health grant from NIMH, we are tracking 5+ million Twitter users over time. This effectively yields a digital cohort study comprising 1.5+% of the US population, allowing us to observe population trends by generalizing from the changes we see within this cohort over time **(see Fig. 7)**. In the current approach, in the first step, we effectively control for the distorting impact of single "super-posting" Twitter accounts (such as bots) by deriving estimates by aggregating over users, not Twitter posts (**Fig. 7, left**). In the second step, we estimate the age, gender, income, and education of the Twitter users with high accuracy from their language [14]. We then compare the demographics of our cohort against government-reported county demographics to derive

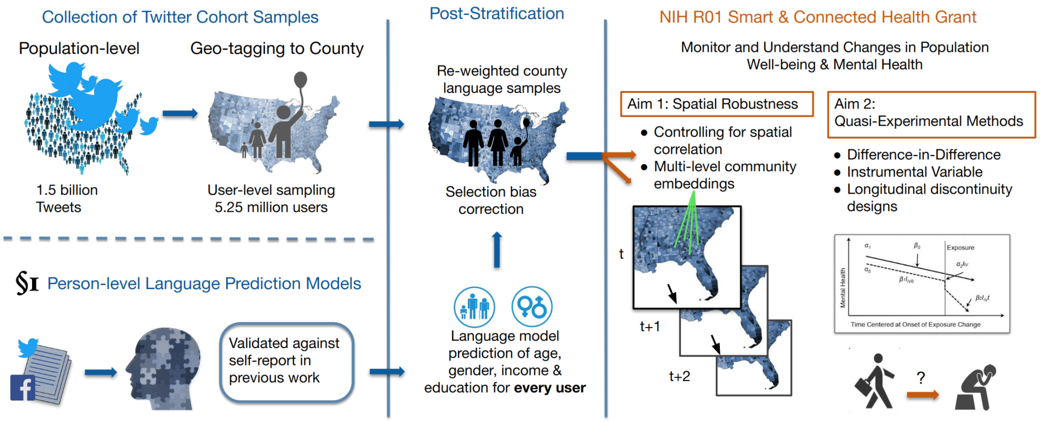

Figure 7 (§3.1): Psychological population surveillance through tracking cohorts of Twitter users and demographic post-stratification.

person-level post-stratification weights [30]. As a result, we obtain a county-level Twitter language sample that is corrected for selection biases using methods similar to the post-stratification in population phone surveys (**Fig. 7, middle**). Based on these more stabilized language samples that track the same users over time, we are evolving quasi-experimental methods (difference-in-difference, longitudinal discontinuity, and instrumental variable designs) to detect changes and test hypotheses about how events and policy decisions impact psychosocial community processes (**Fig. 7, right**). For example, in a first study, we are examining the impact of COVID-19 lockdowns on mental health by exploiting the fact that shelter-in-place measures went into effect at different times in different states.

§3.2 R21: Benchmarking Reddit and Twitter as an early warning system for the opioid epidemic. In a 2022 R21 grant from NIDA, we are exploring how well the use of drug-related keywords on Reddit and Twitter–as well as more general markers of despair–can serve as an early warning system for when and where opioids will be used (especially Fentanyl). We are collaborating with the California Department of Public Health to determine if this system can inform the allocation of scarce overdose medications, like Naloxone pens. This work is an example of how computational methods applied to large natural social data sets may yield a potential breakthrough technology for resource allocation in public health.

§4 Future Directions

Unobtrusive assessment through social media and other natural language is particularly valuable in contexts where no alternative to such measurements exists. This is the case for population monitoring in emerging economies and the Global South (§4.1) or when trying to monitor a broad range of psychological variables in patient populations that cannot be burdened with additional surveys (§4.2). Beyond passive measurement, the recent breakthroughs in generative language models point to a future of targeted interventions that leverage psychological theory at scale (§4.3).

§4.1 Enable population health surveillance in LMICs and the Global South. Social media can be a tool for cost-effectively monitoring population mental health and the social determinants of health in resource-poor countries. In joined work, the Mexican Institute of Statistics (INEGI) has implemented a population well-being monitoring system through Twitter (prototype). We are currently revising an NIH R01 grant to develop such systems for Mexico and India. A particularly pressing need for cost-effective mental health surveillance arises through **climate change**, which disproportionately impacts the Global South and emerging economies.

§4.2 Understand patients' social and psychological contexts using natural language linked to medical records. I am very interested in digital text linked with biomarkers, diagnoses, and health outcomes in medical records [7, 11]. The language patients naturally produce outside of health care systems is a complex digital phenotype that can be mined to understand the psychological factors driving adherence, substance use, and risk and resilience in general. In addition, through analyzing clinical notes, it may be possible to understand physician attitudes toward patients and additional factors impacting their care. I am looking forward to collaborations leveraging these datasets and to train Ph.D. students and postdocs to carry out such studies, from across psychology and the social and life sciences.

§4.3 Using generative language models (such as GPT-3) for scalable interventions. Beyond using NLP to measure mental health and psychology, generative language models are disrupting the knowledge economy with their ability to write coherent and original text and communicate interactively. In first studies, we are exploring the ability of GPT-3 to generate messages to change opinions on smoking and assault weapon bans which pilot data suggests may be more persuasive than those produced by humans. Further, we are using NLP to measure treatment adherence in psychotherapy and developing a new augmented digital chat treatment modality in which a GPT system interactively provides suggestions to a therapist reflecting different therapeutic skills (such as reframing).

References

[1] **Eichstaedt, J. C.**, Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., et al. (2015). "Psychological language on Twitter predicts county-level heart disease mortality". *Psychological science* 26.2, 159–169.

[2] **Eichstaedt, J. C.**, Sherman, G. T., Giorgi, S., Roberts, S. O., Reynolds, M. E., Ungar, L. H., and Guntuku, S. C. (2021a). "The emotional and mental health impact of the murder of George Floyd on the US population". *Proceedings of the National Academy of Sciences* 118.39, e2109139118.

[3] Ward, G., De Neve, J.-E., Ungar, L. H., and **Eichstaedt, J. C.** (2021). "(Un) happiness and voting in US presidential elections." *Journal of personality and social psychology* 120.2, 370.

[4] **Eichstaedt, J. C.**, Kern, M. L., Yaden, D. B., Schwartz, H., Giorgi, S., Park, G., Hagan, C. A., Tobolsky, V. A., Smith, L. K., Buffone, A., et al. (2021b). "Closed-and open-vocabulary approaches to text analysis: A review, quantitative comparison, and recommendations." *Psychological Methods* 26.4, 398.

[5] Kern, M. L., Park, G., **Eichstaedt, J. C.**, Schwartz, H. A., Sap, M., Smith, L. K., and Ungar, L. H. (2016). "Gaining insights from social media language: Methodologies and challenges." *Psychological methods* 21.4, 507.

[6] Schwartz, H. A., **Eichstaedt, J. C.**, Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M. E. P., and Ungar, L. H. (2013a). "Personality, Gender, and Age in the Language of Social Media: The Open-Vocabulary Approach". *PloS ONE* 8.9, e73791.

[7] **Eichstaedt, J. C.**, Smith, R. J., Merchant, R. M., Ungar, L. H., Crutchley, P., Preoţiuc-Pietro, D., Asch, D. A., and Schwartz, H. A. (2018a). "Facebook language predicts depression in medical records". *Proceedings of the National Academy of Sciences* 115.44, 11203–11208.

[8] Preoţiuc-Pietro, D., Schwartz, H. A., Park, G., **Eichstaedt, J. C.**, Kern, M., Ungar, L., and Shulman, E. (2016). "Modelling valence and arousal in facebook posts". In: *Proceedings of the 7th workshop on computational approaches to subjectivity, sentiment and social media analysis*, pp. 9–15.

[9] **Eichstaedt, J. C.** and Weidman, A. C. (2020). "Tracking fluctuations in psychological states using social media language: A case study of weekly emotion". *European Journal of Personality* 34.5, 845–858.

[10] Guntuku, S. C., Yaden, D. B., Kern, M. L., Ungar, L. H., and **Eichstaedt, J. C.** (2017). "Detecting depression and mental illness on social media: an integrative review". *Current Opinion in Behavioral Sciences* 18, 43–49.

[11] Merchant, R. M., Asch, D. A., Crutchley, P., Ungar, L. H., Guntuku, S. C., **Eichstaedt, J. C.**, Hill, S., Padrez, K., Smith, R. J., and Schwartz, H. A. (2019). "Evaluating the predictability of medical conditions from social media posts". *PloS one* 14.6, e0215476.

[12] **Eichstaedt, J. C.**, Guntuku, S. C., Stade, E., Reynolds, M., and Subrahmanya, S. (n.d.). "The Language of Depression". In Preparation.

[13] Kern, M. L., **Eichstaedt, J. C.**, Schwartz, H. A., Park, G., Ungar, L. H., Stillwell, D. J., Kosinski, M., Dziurzynski, L., and Seligman, M. E. (2014a). "From "Sooo excited!!!" to "So proud": Using language to study development." *Developmental psychology* 50.1, 178.

[14] Sap, M., Park, G., **Eichstaedt, J. C.**, Kern, M., Stillwell, D., Kosinski, M., Ungar, L., and Schwartz, H. A. (2014). "Developing age and gender predictive lexica over social media". In: *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1146–1151.

[15] Park, G., Yaden, D. B., Schwartz, H. A., Kern, M. L., **Eichstaedt, J. C.**, Kosinski, M., Stillwell, D., Ungar, L. H., and Seligman, M. E. (2016). "Women are warmer but no less assertive than men: Gender and language on Facebook". *PloS one* 11.5, e0155885.

[16] Park, G., Schwartz, H. A., **Eichstaedt, J. C.**, Kern, M. L., Stillwell, D. J., Kosinski, M., Ungar, L. H., and Seligman, M. E. (2014). "Automatic personality assessment through social media language". *Journal of Personality and Social Psychology* 108 (6), 934–952.

[17] Kern, M. L., **Eichstaedt, J. C.**, Schwartz, H. A., Dziurzynski, L., Ungar, L. H., Stillwell, D. J., Kosinski, M., Ramones, S. M., and Seligman, M. E. (2014b). "The online social self: An open vocabulary approach to personality". *Assessment* 21.2, 158–169.

[18] Schwartz, H. A., **Eichstaedt, J. C.**, Dziurzynski, L., Kern, M. L., Blanco, E., Kosinski, M., Stillwell, D., Seligman, M. E. P., and Ungar, L. H. (2013b). "Toward Personality Insights from Language Exploration in Social Media". In: *In Proceedings of the AAAI Spring Symposium Series*. Stanford, CA.

[19] Park, G., Schwartz, H. A., Sap, M., Kern, M. L., Weingarten, E., **Eichstaedt, J. C.**, Berger, J., Stillwell, D. J., Kosinski, M., Ungar, L. H., et al. (2015). "Living in the Past, Present, and Future: Measuring Temporal Orientation with Language". *Journal of personality*.

[20] Schwartz, H. A., Park, G. J., Sap, M., Weingarten, E., **Eichstaedt, J. C.**, Kern, M. L., Stillwell, D., Kosinski, M., Berger, J., Seligman, M., et al. (2015). "Extracting Human Temporal Orientation in Facebook Language". In: *NAACL-2015: Conference of the North American Chapter of the Association for Computational Linguistics.* Association for Computational Linguistics.

[21] Guntuku, S. C., Buffone, A., Jaidka, K., **Eichstaedt, J. C.**, and Ungar, L. H. (2019). "Understanding and measuring psychological stress using social media". In: *Proceedings of the international AAAI conference on web and social media*. Vol. 13, pp. 214–225.

[22] Jaidka, K., Buffone, A., **Eichstaedt, J. C.**, Rouhizadeh, M., and Ungar, L. H. (2018). "Modeling and visualizing locus of control with facebook language". In: *Twelfth International AAAI Conference on Web and Social Media*.

[23] Giorgi, S., Nguyen, K. L., **Eichstaedt, J. C.**, Kern, M. L., Yaden, D. B., Kosinski, M., Seligman, M. E., Ungar, L. H., Schwartz, H. A., and Park, G. (2022a). "Regional personality assessment through social media language". *Journal of personality* 90.3, 405–425.

[24] Rodríguez-Pose, A., Obschonka, M., Lee, N., **Eichstaedt, J. C.**, and Ebert, T. (2018). "Big data, artificial intelligence and the geography of entrepreneurship in the United States". *Artificial Intelligence and the Geography of Entrepreneurship in the United States (May 2018). CEPR Discussion Paper No. DP12949*.

[25] Yaden, M., Yaden, D., Buffone, A., **Eichstaedt, J. C.**, Crutchley, P., Smith, L., Cass, J., Callahan, C., Rosenthal, S., Ungar, L., et al. (2020). "Linguistic analysis of empathy in medical school admission essays". *International Journal of Medical Education* 11, 186.

[26] Son, Y., Clouston, S. A., Kotov, R., **Eichstaedt, J. C.**, Bromet, E. J., Luft, B. J., and Schwartz, H. A. (2021). "World Trade Center responders in their own words: predicting PTSD symptom trajectories with AI-based language analyses of interviews". *Psychological medicine*, 1–9.

[27] Liu, T., Meyerhoff, J., **Eichstaedt, J. C.**, Karr, C. J., Kaiser, S. M., Kording, K. P., Mohr, D. C., and Ungar, L. H. (2022). "The relationship between text message sentiment and self-reported depression". *Journal of affective disorders* 302, 7–14.

[28] Schwartz, H. A., **Eichstaedt, J. C.**, Kern, M. L., Dziurzynski, L., Lucas, R. E., Agrawal, M., Park, G. J., Lakshmikanth, S. K., Jha, S., Seligman, M. E., et al. (2013c). "Characterizing Geographic Variation in Well-Being Using Tweets." In: *ICWSM*, pp. 583–591.

[29] Giorgi, S., Preoţiuc-Pietro, D., Buffone, A., Rieman, D., Ungar, L., and Schwartz, H. A. (2018). "The Remarkable Benefit of User-Level Aggregation for Lexical-based Population-Level Predictions". In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1167–1172.

[30] Giorgi, S., Lynn, V. E., Gupta, K., Ahmed, F., Matz, S., Ungar, L. H., and Schwartz, H. A. (May 2022b). "Correcting Sociodemographic Selection Biases for Population Prediction from Social Media". *Proceedings of the International AAAI Conference on Web and Social Media* 16.1, 228–240.

[31] **Eichstaedt, J. C.**, Schwartz, H. A., Giorgi, S., Kern, M. L., Park, G., Sap, M., Labarthe, D. R., Larson, E. E., Seligman, M., and Ungar, L. H. (2018b). "More evidence that Twitter language predicts heart disease: a response and replication". *PsyArXiv Preprints*.

[32] Jaidka, K., Giorgi, S., Schwartz, H. A., Kern, M. L., Ungar, L. H., and **Eichstaedt, J. C.** (2020). "Estimating geographic subjective well-being from Twitter: A comparison of dictionary and data-driven language methods". *Proceedings of the National Academy of Sciences* 117.19, 10165–10171.