

Johannes C. Eichstaedt, Ph.D.

Institute for Human-Centered AI & Department of Psychology | Stanford University | eichstaedt@stanford.edu
(215) 815-1640 | [LinkedIn](#) | [Website](#) | [Stanford Lab](#) | [Lab GitHub](#) | [Full Academic CV](#) | [Publications](#) | [Google Scholar](#)

Professional Summary

Research scientist and scientific leader with experience directing an interdisciplinary data science [research lab](#) focused on the impact of LLMs on society and 13 years of scientific project management experience leading interdisciplinary teams through the entire scientific life cycle, resulting in publishing [80+ peer-reviewed articles](#) cited 10,000+ times. Co-founded 3 tech startups that use NLP and AI for well-being and mental health. Passionate about deploying NLP, AI, and LLM technologies for social and psychological benefit, pushing the social sciences to track and predict specific societal impacts of these technologies, and communicating emerging impacts to broad stakeholder communities, including policy advocacy and PR/press engagement.

Relevant Skills

- Expert interdisciplinary computational social scientist trained in graduate physics, NLP, ML, and AI, as well as in psychology and the social sciences focused on deploying, evaluating, and testing NLP, ML, and LLM systems to benefit psychology and society, with [publications](#) spanning top computer sciences conferences as well as top journals in psychology, sociology, communication, public health, and public policy.
- Proven ability to design and effectively communicate research projects that attract broad scientific and public interest including media coverage spanning 170+ news articles [[link](#), [link](#)], including [6 stories](#) in The New York Times and a published article in the [Scientific American](#).
- Exceptional public speaker and expert advocate for positive technological impact in 70+ presentations to broad public, scientific, and policy audiences, including serving as an expert to governments and at the United Nations General Assembly, Organization for Economic Co-Operation and Development, and World Health Organization.

Technical skills

- Data science (expert R, basic Python [pandas, sklearn, NLTK] and matlab), including inferential statistics, structural equation modeling, dimensionality reduction (PCA, factor analysis), as well as databases (advanced SQL) and basic system administration (Linux servers, AWS and GCP)
- Deep expertise in NLP (including teaching a [PhD course](#)), wrote a [key introduction](#) to psychological language analysis, working with embeddings, LDA topic modeling, and co-developed Python text-analysis toolkit ([DLATK](#))
- ML, including support vector machines, random forests, and decision trees, elastic nets, etc.
- Strategies for LLM design (few-shot prompting, RAG, fine-tuning) and evaluation (user interviews, language-based assessments [[publication](#)])
- Expert data visualization (R/ggplot) and storytelling with data

Work History

Assistant Professor (Research) & Shiram Faculty Fellow, Institute for Human-Centered AI (HAI) & Department of Psychology, Stanford University, 2020 to present

Demonstrated excellence in leading a computational psychology and NLP [lab](#), supervising and mentoring [3 postdocs](#), [2 graduate students](#), [3 data scientists](#), and [3 staff](#) to publish 40+ research papers in high-impact venues, with 15+ ongoing research projects spanning data collection, NLP/ML/data analysis, and manuscript preparation

- Current research is scoping and delivering the staged, responsible roll-out of LLMs in mental health care ([publication](#)), including evaluation framework for effectiveness, equity, and safety ([publication](#)).
 - Conducting prompt engineering through iteration with expert clinicians (w/ few-shot prompting), and user interviews with therapists and patients to deploy LLMs to help train therapists with simulated patients, providing coaching based on therapist-patient transcripts, and designing “co-pilots” for therapy homework.
- Evaluated the (human-like) persuasiveness of LLM-generated messages for polarizing issues using randomized experimental design ([publication](#)).
- Evaluated how AI algorithms contribute to addictive platform/media consumption to provide guidance for policymakers (forthcoming).
- Argued for social scientists to resist simplistic narratives around LLM’s impact on society and instead to rely on specific domain expertise to make testable predictions and set up measurement/evaluation infrastructure to measure impacts *now* (for timely policy input) [[talk](#), publ. forthcoming]; wrote [policy brief](#) on depression detection through AI.
- Shared research on using AI for well-being and mental health in 25+ invited talks, including keynotes for the *American Psychological Association*, the *National Academy of Sciences*, *Engineering and Medicine*, as well as led several industry partner workshops on LLMs for the *Institute for Human-Centered AI*.
- Recognized teaching excellence delivering a comprehensive PhD-level [course](#) about psychological text analysis with NLP, ML, and LLMs with teaching effectiveness evaluations at the top of the school distribution.

Lead Research Scientist & Co-Principal Investigator, World Well-Being Project Consortium (originally at the University of Pennsylvania), 2011 to present

Co-founded and now co-direct a 5 university consortium of research labs focused on applying NLP, ML, and LLMs for the social good; co-raised \$9.3M in research funding and collectively trained 50+ graduate students and postdocs, now in impactful industry and academic positions.

- Coordinated collaborations with 138 colleagues across 48 universities and organizations to successfully complete high-impact research projects using NLP and ML/AI for the social good, resulting in a total of [80+ publications](#) across ML, NLP, HCI, AI conferences, and social and health journals.
 - Led and implemented all NLP/ML in the first research to demonstrate that patients' social media language can be used to predict future diagnoses of depression in medical records [[publication](#)], and the first research showing that geo-aggregated Twitter posts can be used to predict the heart-disease mortality of communities and characterize risk factors [[publication](#)].
 - Both papers resulted in 177 new stories ([link](#), [link](#)), including *The New York Times*, and reached the top 1% of public attention for scientific outputs.
- Co-designed and contributed to a shared scientific Python codebase (Differential Language Analysis ToolKit, [DLATK](#), build on top of sklearn, pandas, NLTK, etc.) that analyzes text for psychological and social insights used by 10+ research labs and in 150+ papers.
 - Designed and evaluated particular use cases to combine NLP operations (e.g., feature extraction, topic modeling) with statistical social science inference to identify language distinguishing psych. states and traits.
- Let collaboration and held workshops with *INEGI* (the Mexican National Institute for Statistics) on the measurement of population mental health using Twitter/X, resulting in a real-time dashboard for state-level well-being.
 - Helped the *World Health Organization* (WHO) develop a [consensus statement](#) on measuring the burden of misinformation during COVID-19.
- Experienced in management and HR functions for 13 years, including choosing 23 hires (Post-docs and PhD students, data scientists, engineers, project managers, finance admins, and research assistants) with a track record of championing diversity (63% of my hires reflected at least one kind of diversity).

Chief Scientist, Jimini Health, 2023 to present

Developed scientific and IP strategy for an AI start-up (current evaluation: \$34 million) using language model programming and multiple interacting LLMs to augment text-based therapy with therapist + LLM "care teams;" coordinated with cross-functional teams (science, marketing, legal, engineering, UX design) to meet growth and business objectives

- Developed strategic partnerships with scientific advisors/faculty across computer science/NLP, medicine, digital health, and clinical psychology to scope, implement, and evaluate LLMs in psychotherapy, including measuring the impact on patient symptom improvement and building out product roadmap.
- Supervised and designed LLM engineering strategies in high-stakes therapy contexts:
 - Used multiple interacting LLMs to deliver conversational turns to patients to meet a specific therapeutic or relational objective, either chosen by therapist or high-order LLM performing case conceptualization.
 - Relied on few-shot prompting and RAG to tailor LLMs for treatment goals, matching therapist's style, and maximizing therapeutic alliance between patient and therapist + LLM care team.
 - Designed pipeline for LLM-based real-time evaluation of conversations for quality of treatment delivery and potential need for escalation to a therapist (e.g., patient suggesting self-harm).

Education

- Ph.D., Psychology, University of Pennsylvania, 2017 ([thesis](#)). Graduate NLP, ML & clinical training in psychodiagnostic assessment, behavior therapy.
- M.A., Psychology, University of Pennsylvania, 2013
- M.A.P.P., Applied Positive Psychology, University of Pennsylvania, 2011
- M.S., Particle Physics, University of Chicago, 2010 ([thesis](#))
- B.S. (Hons.), Physics & Philosophy, King's College London, 2009, highest university distinction

Honors

- 2022: Rising Star, Association for Psychological Science (APS)
- 2021: Early Career Researcher Award, International Positive Psychology Association (IPPA)
- 2016: Scientist of the Year, Philadelphia Geek Awards
- 2014: Emerging Leader in Science & Society, American Association for the Advancement of Science (AAAS)

Languages

- English (quasi-native), German (native), Spanish (intermediate), French (fluent reader, basic speaker)