

Neil Fasching

Quantitative Researcher | Computational Social Scientist | PhD
neilfasching@gmail.com | (818) 577-3205 | www.neilfasching.com | New York, NY

PROFESSIONAL SUMMARY

Quantitative researcher with industry experience at Meta and a PhD from the University of Pennsylvania. Specializes in survey design and analysis, experimental and quasi-experimental methods, and statistical inference (causal, Bayesian, and machine learning) applied to large-scale survey and behavioral datasets. At Meta, designed and fielded on- and off-platform surveys related to Facebook's Friending Intent and People You May Know (PYMK) algorithm. First-author on an *ACL* 2025 paper measuring demographic bias and reliability in LLM-based systems at scale. Additional first-author publications in top-tier journals including *Science Advances*, *Journal of Personality*, and *British Journal of Political Science*.

PROFESSIONAL EXPERIENCE

Quantitative UX Researcher

Mar 2026 – May 2026

Meta – Central Growth

- Updated the “New Connections” off-platform survey, then used Claude Code and OpenAI Codex to modernize the analysis, building automated longitudinal workflows across previously siloed survey waves
- Designed a quantitative on-platform Friending Intent survey, with results to directly inform the Friending team's decision to update Facebook's People You May Know (PYMK) algorithm
- Smoke-tested the *uxr-toolkit* research infrastructure (Claude Code slash command), helping to design quantitative research workflows that could drive efficiency gains across Meta's UXR Org
- Helped build and deploy multiple internal AI agents to accelerate research workflows, including a personal knowledge agent (“Second Brain”), a shared team agent (“Team Brain”) for the Central Growth Friending team, and a custom MyClaw instance (Meta's internal agentic framework)

Computational Research Fellow

Sep 2021 – Mar 2026

University of Pennsylvania

- Collaborated with professors Dr. Yphtach Lelkes and Dr. Duncan J. Watts on projects investigating the influence of news media and social media on human behavior
- Designed and executed large-scale quantitative studies including surveys, behavioral experiments, and NLP-based analysis to understand how users interact with media, AI systems, and digital platforms
- Built AI-powered measurement pipelines analyzing over 28,000 podcast episodes and 46.7 million social media posts, delivering findings to interdisciplinary research teams and public-facing stakeholders
- Shipped a public-facing research product (mediabiasdetector.seas.upenn.edu) for external audiences

Co-Instructor, Modern Data Mining (PhD Level)

Jul 2022 – Mar 2026

The Wharton School, University of Pennsylvania

- Taught PhD-level machine learning at Wharton covering LLMs, neural networks, and ensemble methods, translating complex statistical concepts for technical and non-technical audiences.

EDUCATION

Ph.D., Communication (Computational Social Science)

2026

University of Pennsylvania

Dissertation: Leveraging LLMs to Evaluate Topics, Misinformation, and Toxicity on Political Podcasts

Master of Science, Statistics and Data Science

2023

The Wharton School, University of Pennsylvania

SELECTED RESEARCH PROJECTS

Persistent Polarization: The Unexpected Durability of Political Animosity Around US Elections (*Science Advances* 2024)

- Analyzed 66,000 cross-sectional and panel survey interviews to quantify the effect of election proximity on political animosity
- Employed Interrupted Time Series (ITS) models to causally demonstrate that political animosity remains persistently high and stable around elections, challenging conventional wisdom about polarization

Model-Dependent Moderation: Inconsistencies in Hate Speech Detection Across LLM-based Systems (*ACL* 2025)

- Designed a 1.3-million-sentence factorial experiment to systematically measure how seven LLM-based content moderation systems (including Claude and GPT) produce inconsistent outputs across demographic groups
- Developed novel measurement metrics for AI system reliability and bias, establishing a reusable framework for auditing AI product behavior and quality at scale

Triangulating Political Toxicity on Twitter ([Link to Paper](#))

- Examined 46.7 million tweets (2012-2022) for toxicity (including harassing, hateful, and violent speech)
- Developed scalable measurement pipelines for toxicity using OpenAI and Mistral moderation systems
- Identified patterns in toxicity across time, demographics, and topics using advanced ML models

Partisan Airways: Leveraging LLMs to Evaluate Topics, Misinformation, and Toxicity on Political Podcasts (Dissertation)

- Analyzed over 28,000 podcast episodes for topics, misinformation, and toxicity
- Developed two novel measurement frameworks for assessing the prevalence of misinformation and hate speech at scale
- Benchmarked LLM-based model performance for transcription, topic segmentation, misinformation identification and hate speech classification

Automated Annotation with Generative AI Requires Validation (*arXiv* 2023)

- Built and validated a 5-step human-in-the-loop workflow for LLM text annotation, classifying 200,000+ text samples across 27 tasks and 11 datasets
- Introduced the “consistency score” metric to identify LLM failure modes before deployment as a quality and accuracy check, now adopted as a best practice for AI-assisted research

TECHNICAL SKILLS

Research Methods: Survey design and fielding, experimental design, A/B testing, causal inference (ITS, DiD, RDD), behavioral data analysis, mixed methods

Statistical Analysis: Regression (linear, logistic, multilevel), SEM, factor analysis, time-series analysis, Bayesian methods

Programming Languages: Python (Expert), R (Expert), SQL (Expert), JavaScript (Proficient)

Data Processing: Pandas, NumPy, PySpark, PyArrow, dplyr/tidyverse

AI Coding & Agent Tools: Claude Code, OpenAI Codex, OpenClaw, Hermes agents

Platforms: AWS, Microsoft Azure, Google Colab, Posit Workbench, Git/GitHub