Victoria Lin

PERSONAL INFORMATION

Phone: +1 (408) 318-4905
Email: victoria@stat.cmu.edu
Citizenship: United States, Canada

Website
Google Scholar

EDUCATION

Carnegie Mellon University 2020 - 2026
Ph.D., Statistics and Machine Learning
Advisors: Louis-Philippe Morency and Eli Ben-Michael

Carnegie Mellon University 2020 - 2021
M.S., Statistics

Carnegie Mellon University 2018 - 2020
M.S., Computational Data Science

Harvard University 2013 - 2017
A.B. (cum laude), Statistics and Molecular & Cellular Biology

AWARDS

Meta Research PhD Fellowship 2023 - 2025
Two Sigma Diversity PhD Fellowship Finalist 2023
Best Paper Award Nominee, ICMI 2020 2020
Best Paper Award, AffCon Workshop at AAAI 2020 2020
Research Fellowship, Harvard College Program for Research in Science and Engineering 2016
Research Grant, Harvard College Pechet Family Research Fund 2016

EXPERIENCE

Carnegie Mellon University Jan 2019 - Aug 2026
Graduate Research Assistant (Advisors: Louis-Philippe Morency and Eli Ben-Michael)

I work on "causal NLP," or causal inference for natural language data and language models. Motivated by the need for models to be robust, generalizable, and aligned with social values, my recent work identifies failure modes in current LLM training and alignment practices, such as their susceptibility to bias and confounding, and proposes solutions built on formal ideas from causal theory.

My overarching research agenda aims to develop a family of causal inference methods that improve the ability to learn principled language models from imperfect data. Some other technical problems I find interesting include learning causal effects from high-dimensional unstructured data like language and learning language representations for valid causal inference.

Previously, I worked on affective computing and multimodal machine learning (with LP Morency and Jeffrey Girard) and mental health applications of causal inference (with Edward Kennedy).

Google DeepMind Feb 2026 - June 2026
Student Researcher (Host: Alexander D'Amour)

Microsoft Research Cambridge Jun 2025 - Aug 2025
Research Intern (Host: Javier González)

Microsoft Research Redmond & Microsoft Experiences+Devices Jun 2022 - Aug 2022

JEM Research Intern (Hosts: Srinagesh Sharma and Dimitrios Dimitriadis)

Harvard School of Public Health Sept 2017 - Aug 2018
Research Assistant (Advisor: Miguel Hernán)

Broad Institute of MIT and Harvard Jun 2016 - May 2017
Research Assistant (Advisor: Steven McCarroll)

Publications

* denotes equal contribution

Preprints

The Illusion of Intervention: Your LLM-Simulated Experiment is an Observational Study
Victoria Lin, Taedong Yun, Maja Matarić, John Canny, Arthur Gretton, Alexander D'Amour
arXiv:2605.20767

Peer-Reviewed Publications

Omitted Variable Bias of Language Models Under Distribution Shift
Victoria Lin, Louis-Philippe Morency, Eli Ben-Michael
ICML 2026

Better Think Thrice: Learning to Reason Causally with Double Counterfactual Consistency
Victoria Lin, Xinnuo Xu, Rachel Lawrence, Risa Ueno, Amit Sharma, Javier González, Niranjani Prasad
ICLR 2026 Workshop on Logical Reasoning of Large Language Models

Isolated Causal Effects of Natural Language [code]
Victoria Lin, Louis-Philippe Morency, Eli Ben-Michael
ICML 2025

Optimizing Language Models for Human Preferences is a Causal Inference Problem [code]
Victoria Lin, Eli Ben-Michael, Louis-Philippe Morency
UAI 2024

Text-Transport: Toward Learning Causal Effects of Natural Language [code]
Victoria Lin, Louis-Philippe Morency, Eli Ben-Michael
EMNLP 2023

Counterfactual Augmentation for Multimodal Learning Under Presentation Bias [code]
Victoria Lin, Louis-Philippe Morency, Dimitrios Dimitriadis, Srinagesh Sharma
EMNLP Findings 2023

SenteCon: Leveraging Lexicons to Learn Human-Interpretable Language Representations [code]
Victoria Lin, Louis-Philippe Morency
ACL Findings 2023

SeedBERT: Recovering Annotator Rating Distributions from an Aggregated Label
Aneesha Sampath, **Victoria Lin**, Louis-Philippe Morency
AAAI 2023 Workshop on Uncertainty Reasoning and Quantification in Decision Making

Toward Multimodal Modeling of Emotional Expressiveness [code]
Victoria Lin*, Jeffrey Girard*, Michael Sayette, Louis-Philippe Morency
ICMI 2020. Nominated for Best Paper Award.

gfoRmula: An R Package for Estimating the Effects of Sustained Treatment Strategies via the Parametric g-formula [code]
Sean McGrath*, **Victoria Lin***, Zilu Zhang, Lucia Petito, Roger Logan, Miguel Hernán, Jessica Young
Patterns 1(3), 100008

Context-Dependent Models for Predicting and Characterizing Facial Expressiveness
Victoria Lin, Jeffrey Girard, Louis-Philippe Morency
AAAI 2020 Workshop on Affective Content Analysis. Best Paper Award.

Software Libraries

`sentecon` (4,500+ downloads): Python library for interpretable language representations.
`gfoRmula` (44,000+ downloads): R package for causal effect estimation with the g-formula.

Presentations

Invited Talks

Carnegie Mellon University DeGroot Research Workshop 2023
Counterfactual Augmentation for Learning Under Presentation Bias.

Contributed Talks

American Causal Inference Conference (ACIC) 2025
Isolated Causal Effects of Natural Language.

American Causal Inference Conference (ACIC) 2023
Generalizing Text Experiments to Real-World Contexts with Large Language Models.

International Conference on Multimodal Interaction (ICMI) 2020
Toward Multimodal Modeling of Emotional Expressiveness.

AffCon Workshop at the AAAI Conference on Artificial Intelligence 2020
Context-Dependent Models for Predicting and Characterizing Facial Expressiveness.

Reviewing

NeurIPS (2026), ICML (2026), UAI (2025), ACL (2023), EMNLP (2022, 2023)

Teaching

Carnegie Mellon University *Graduate Teaching Assistant*

36-402 Advanced Methods for Data Analysis	Spring 2023
36-468 Text Analysis	Fall 2022
36-462 Methods of Statistical Learning	Spring 2022
36-309 Experimental Design for Behavioral & Social Sciences	Fall 2021
36-401 Modern Regression	Fall 2020
11-631 Data Science Seminar (*also as guest lecturer*)	Fall 2019

Harvard University *Undergraduate Teaching Fellow*

CS109A Data Science	Fall 2016

Mentorship

Aneesha Sampath (CMU B.S. → University of Michigan Ph.D.)

Skills

Programming & Frameworks: Python, PyTorch, HuggingFace, R, Bash, Git
Languages: English (native), Mandarin Chinese (heritage)