

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 Aug 2020 - Aug 2026 (expected)
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 2019 -
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. Some technical problems I find interesting include using causal principles to improve language models, 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 2026
Student Researcher (Host: [Alexander D’Amour](#))

Microsoft Research Cambridge 2025
Research Intern (Host: [Javier González](#))

Microsoft Research Redmond & Microsoft Experiences+Devices 2022
JEM Research Intern (Hosts: [Srinagesh Sharma](#) and [Dimitrios Dimitriadis](#))

Harvard School of Public Health 2017 - 2018
Research Assistant (Advisor: [Miguel Hernán](#))

PUBLICATIONS

* denotes equal contribution

PREPRINTS

Victoria Lin, Louis-Philippe Morency, Eli Ben-Michael. “Omitted Variable Bias of Language Models Under Distribution Shift.” *Under review*.

Victoria Lin, Xinnuo Xu, Rachel Lawrence, Risa Ueno, Amit Sharma, Javier González, Niranjani Prasad. “Better Think Thrice: Learning to Reason Causally with Double Counterfactual Consistency.” *Under review*.

PEER-REVIEWED PUBLICATIONS

Victoria Lin, Louis-Philippe Morency, Eli Ben-Michael. “Isolated Causal Effects of Natural Language.” [\[PDF, code\]](#)
ICML 2025.

Victoria Lin, Eli Ben-Michael, Louis-Philippe Morency. “Optimizing Language Models for Human Preferences is a Causal Inference Problem.” [\[PDF, code\]](#)
UAI 2024.

Victoria Lin, Louis-Philippe Morency, Eli Ben-Michael. “TEXT-TRANSPORT: Toward Learning Causal Effects of Natural Language.” [\[PDF, code\]](#)
EMNLP 2023.

Victoria Lin, Louis-Philippe Morency, Dimitrios Dimitriadis, Srinagesh Sharma. “Counterfactual Augmentation for Multimodal Learning Under Presentation Bias.” [\[PDF, code\]](#)
EMNLP Findings 2023.

Victoria Lin, Louis-Philippe Morency. “SENTECON: Leveraging Lexicons to Learn Human-Interpretable Language Representations.” [\[PDF, code\]](#)
ACL Findings 2023.

Aneesha Sampath, **Victoria Lin**, Louis-Philippe Morency. “SEEDBERT: Recovering Annotator Rating Distributions from an Aggregated Label.” [\[PDF\]](#)
UDM Workshop at AAAI 2023.

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

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

Victoria Lin, Jeffrey Girard, Louis-Philippe Morency. “Context-Dependent Models for Predicting and Characterizing Facial Expressiveness.” [\[PDF\]](#)
AffCon Workshop at AAAI 2020. *Best Paper Award*.

Leslie Tong, Seo Yeon Yoon, Yaisa Andrews-Zwilling, Alyssa Yang, **Victoria Lin**, Hanci Lei, Yadong Huang. “Enhancing GABA Signaling during Middle Adulthood Prevents Age-Dependent GABAergic Interneuron Decline and Learning and Memory Deficits in ApoE4 Mice.” [\[PDF\]](#)
Journal of Neuroscience 36(7), 2316-2322.

SOFTWARE LIBRARIES [sentecon \(3,400+ downloads\)](#): Python library for interpretable language representations.
[gfoRmula \(38,000+ downloads\)](#): 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)