EMNLP 2021

The First Workshop on Causal Inference and Natural Language Processing

Proceedings of CI+NLP

November 10, 2021

Introduction

Establishing causal relationships is a fundamental goal of scientific research. Quantifying the effectiveness of a vaccine, the persuasive power of a public health ad, or the effectiveness of a lockdown policy, all boil down to questions of causality: how would the treatment (vaccine, ad or policy) affect the outcome (infection rates) compared to a counterfactual world with no treatment? The strength and direction of causal relationships, once formally identified, play a key role in the formulation of clinical treatments, public policy and other long-standing prescriptive strategies.

Text plays an increasingly important role in the study of causal relationships across domains. In the setting of crowd lending, communication researchers have examined the impact of the language of funding requests on receiving loans, where text is the treatment. In the context of immigration policy, political scientists have studied the impact of knowing an illegal immigrant's criminal history on written justifications of whether they should be jailed, where text is the outcome. A growing body of literature studies the interplay between causal inference and NLP.

Despite the interdisciplinary interest in causal inference with text, several challenges remain. For example, it is unclear how deficiencies in NLP methods (such as their inaccuracy with low-resource languages, and their tendency to propagate biases in the data) affect downstream causal estimates. In addition, hyperparameter selection and modeling assumptions in NLP are motivated by accuracy and tractability considerations; how these choices affect downstream causal estimates remains unexplored.

The task of causal identification (linking unobservable causal quantities to observable data) is inherently a process of argument based on domain expertise. Hence, interaction between substantive experts (in medical research, the social sciences and other domains) and methodological experts (in computational linguistics, statistics and other disciplines) is crucial to make valid causal inferences with text. With the existing research in this space spanning several different academic communities, opportunities for such interaction have been lacking thus far.

The **CI+NLP** workshop aims to bring together substantive and methodological experts across academic communities, with an interest in the intersection of causal inference and NLP. It will explore how current and new NLP methodology can contribute to establishing and evaluating causal relationships, and how substantive expertise from different domains can contribute to answering causal questions with text. The CI+NLP workshop differs from prior interdisciplinary NLP workshops (such as the *Text as Data* conference series) in its focus on *causality*, both substantively and methodologically.

Main Workshop Topics. The broad themes of the workshop include, but are not restricted to:

1. *Causal inference theory for language:* Includes methodology for causal inference and optimal policy learning, with experimental or observational data, and text as a treatment, outcome, mediator, moderator or other causal identification device.

2. *Language processing systems and causality:* Includes research on counterfactual predictions with text, causal model interpretation and explanation, causal model bias estimation and debiasing, and methods and datasets to benchmark and analyze the sensitivity and robustness of causal inference and NLP systems.

3. *Society, language, influence and causality:* Includes research on language as a mechanism of influence and provocation, and causal questions related to linguistic framing, medical notes and clinical outcomes, and political messaging and propaganda.

CI+NLP would not have been possible without the dedication of its program committee. We would like to thank them for their invaluable effort in providing timely and high-quality reviews on a short notice. We are also grateful to our invited speakers and panelists for contributing to our program.

The CI+NLP workshop organizers,
Organizing Committee: Amir Feder, Katherine A. Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty

Steering Committee: Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E. Roberts, Uri Shalit, Brandon M. Stewart, Victor Veitch, Diyi Yang

Organizing Committee

- Amir Feder
- Katherine A. Keith
- Emaad Manzoor
- Reid Pryzant
- Dhanya Sridhar
- Zach Wood-Doughty

Steering Committee

- Jacob Eisenstein
- Justin Grimmer
- Roi Reichart
- Margaret E. Roberts
- Uri Shalit
- Brandon M. Stewart
- Victor Veitch
- Diyi Yang

Program Committee

- David Adelani
- Tim Althoff
- Elliott Ash
- Yonatan Belinkov
- Dallas Card
- Alexander D'Amour
- Naoki Egami
- Amir Feder
- Anjalie Field
- Peter Henderson
- Zhijing Jin
- Yohan Jo

- Shalmali Joshi

- Sarvnaz Karimi

- Divyansh Kaushik

- Vivek Kulkarni

- Adyasha Maharana

- Emaad Manzoor

- Huyen Nguyen

- Jekaterina Novikova

- Nadav Oved

- Yuval Pinter

- Matt Post

- Reid Pryzant

- Roi Reichart

- Maarten Sap

- Zeerak Talat

- Chenhao Tan

- Kellie Webster

- Liuyi Yao

- Justine Zhang

Invited Speakers

- Susan Athey

- David Blei

- Cristian Danescu-Niculescu-Mizil

Table of Contents

Conference Program

9:00–9:10   **Opening Remarks**

9:10–10:00   **Keynote 1**

10:00–10:30   **Coffee Break**

10:30–12:00   **Short Talks Session 1**

10:50–11:00   *Causal Augmentation for Causal Sentence Classification*
Fiona Anting Tan, Devamanyu Hazarika, See-Kiong Ng, Soujanya Poria and Roger Zimmermann

11:00–11:10   *Text as Causal Mediators: Research Design for Causal Estimates of Differential Treatment of Social Groups via Language Aspects*
Katherine Keith, Douglas Rice and Brendan O'Connor

11:10–11:20   *Enhancing Model Robustness and Fairness with Causality: A Regularization Approach*
Zhao Wang, Kai Shu and Aron Culotta

11:20–11:30   *What Makes a Scientific Paper be Accepted for Publication?*
Panagiotis Fytas, Georgios Rizos and Lucia Specia

12:00–13:00   **Lunch Break**

13:00–14:30   **Short Talks Session 2**

13:10–13:20   *Sensitivity Analysis for Causal Mediation through Text: an Application to Political Polarization*
Graham Tierney and Alexander Volfovsky

13:20–13:30   *A Survey of Online Hate Speech through the Causal Lens*
Antigoni Founta and Lucia Specia

13:30–13:40   *Identifying Causal Influences on Publication Trends and Behavior: A Case Study of the Computational Linguistics Community*
Maria Glenski and Svitlana Volkova

13:40–13:50   *It's quality and quantity: the effect of the amount of comments on online suicidal posts*
Daniel Low, Kelly Zuromski, Daniel Kessler, Satrajit S. Ghosh, Matthew K. Nock and Walter Dempsey

14:30–14:45   **Mini Break**

14:45–15:35   **Invited Talk 2**

15:35–16:25   **Invited Talk 3**

16:25–16:45   **Coffee Break**

No Day Set (continued)

16:45–17:30   Panel Discussion

17:30–18:30   Poster Session