

A Survey of Online Hate Speech through the Causal Lens

Antigoni-Maria Founta and Lucia Specia

Department of Computing

Imperial College London

United Kingdom

{a.founta20,l.specia}@imperial.ac.uk

Abstract

The societal issue of digital hostility has previously attracted a lot of attention. The topic counts an ample body of literature, yet remains prominent and challenging as ever due to its subjective nature. We posit that a better understanding of this problem will require the use of causal inference frameworks. This survey summarises the relevant research that revolves around estimations of causal effects related to online hate speech. Initially, we provide an argumentation as to why re-establishing the exploration of hate speech in causal terms is of the essence. Following that, we give an overview of the leading studies classified with respect to the direction of their outcomes, as well as an outline of all related research, and a summary of open research problems that can influence future work on the topic.

1 Introduction

User-generated content has flourished with the evolution of social media platforms, and with it there has been an emergence of several social phenomena. Some of them can have a positive impact on mental health and eventually prove beneficial for societies; for example, it has been shown that the use of social networks from individuals of advanced age can lead to social benefits as well as have promising cognitive effects (Quinn, 2018). There is, however, a plethora of societal issues on the digital world which are proliferating in these platforms; such are several types of online misbehaviours.

Despite the clear interest of the research community to prevent, detect and filter harmful content (Schmidt and Wiegand, 2017; Fortuna and Nunes, 2018; MacAvaney et al., 2019; Mishra et al., 2019; Tontodimamma et al., 2021; Vidgen et al., 2019), the problem is very complex and still far from addressed. For instance, its multifactorial character poses already a large obstacle with manifold parameters to consider: diversity of platforms, abundance

of languages, various facets it can manifest on, and different forms of language, to name a few. Even at a higher level, there are some deeply philosophical issues to contemplate, such as the thin line between awry, unwelcome content and freedom of speech. Nonetheless, the impact of online communications on society is considered more beneficial than harmful, hence it is important to advance the current scientific work to address its limitations.

The matter of unhealthy discourse is widely context-dependent (Saleem et al., 2017) and subjective (Salminen et al., 2018; Aroyo et al., 2019). A racial slur, for instance, will be interpreted differently from individuals with separate ethnic or racial backgrounds. Conversely, however, most existing approaches operate on the individual post/comment level or aggregate isolated pieces of information – usually of textual nature – and seek simplistic, ostensibly universal, decisions. Even though this was a necessary course of action at the beginning, such nuanced task requires higher sophistication. To combat the issue to the core, it is of utmost importance to deeply understand the factors that drive it and determine them in a definite way.

An increasingly popular scientific approach to ensure certainty is via the application of causal methods. Causality suggests an area of study that is ubiquitous in many parts of modern science (Russo and Williamson, 2011; Young, 2016; Landes et al., 2018; Hicks et al., 1980), among which social sciences (Marini and Singer, 1988; Reiss, 2009). Especially with the development of methodologies to extract causation from observational, instead of experimental, studies, it is now possible to utilise the profusion of web content for the extraction of causal knowledge (De Choudhury and Kiciman, 2017; Saha et al., 2019b; Doan et al., 2019).

Driven from the necessity of contemplating hate speech from a causal prism, there exists a relevant, yet still narrow, body of literature. In this survey, we attempt to summarise the existing research

which analyse phenomena of digital misbehaviours from a causal perspective, quantifying the influence on fundamental sociological outcomes related to harmful language online. In addition, we highlight the important research gaps and make suggestions for future directions. The goal of the present paper is to call attention to the importance of causation in the context of the discussed issue. In addition, it is intended to act as a point of reference for aspiring practitioners in the topic.

The rest of the paper revolves around the compilation of relevant publications, first classified based on their focus of sociological impact (§2) and, subsequently, arranged according to their approach towards salient methodological settings. The latter section (§3) is organised into several prominent hurdles, which include information regarding how they are treated from the reviewed work and points of discussion about the research gaps.

2 Sociological Causal Impact

Digital profanity is, first and foremost, a sociological phenomenon that affects many aspects of the virtual world. There are some pronounced directions which require thorough examination, such as the consequences of such actions on the affected communities as well as their targets, or the impact of banning policies aiming at fighting against it. Moreover, there are several other underlying effects that could get in the spotlight, including the ramifications of interventions or the reasons that drive toxic actors. As a result, it is of paramount importance to examine hate speech in a holistic way and concretely quantify the drivers of these outcomes. However, it is impossible to achieve this task without considering causality, because non-causal inferences can never be conclusive. Surprisingly, despite the broad interest of the research community in this topic, very little work has been done on attempting causal links, even on the most prominent tasks related to online hate speech (OHS).

In the present survey, we classify fundamental sociological outcomes related to OHS and outline the most distinguished body of literature. The classification results into three major pillars, with respect to the following:

- **Digital misbehaviours versus the physical world:** we summarise studies concerning the propagation of online hate speech to real life (Müller and Schwarz, 2018, 2020) as well as the influence of offline events to the dissemi-

nation of the issue online (Olteanu et al., 2018; Thomas et al., 2021)¹.

- **Harmful content versus the individuals:** we outline research concerning the impact of toxic behaviours on the targets or passive readers (Saha et al., 2019a) as well as the by-products of web characteristics (such as anonymity) on hate speech producers (von Esen and Jansson, 2020).
- **Effect of interventions:** finally, we review works that revolve around quantifying the effect of combating strategies which various platforms adopt. Existing research has focused on limiting policies, censoring and counter-speaking (Álvarez-Benjumea and Winter, 2018), social sanctioning (Munger, 2017), quarantining (Chandrasekharan et al., 2020) and banning hateful communities (Chandrasekharan et al., 2017; Thomas et al., 2021).

In all the of the cases, we can consider the role of hate speech both from the position of a phenomenon which is the result of (potentially) multiple causes, but also from the position of the causal root and study its effects. Therefore, in the following sections we review both directions.

2.1 Digital and Physical World

The first and most apparent pillar concerns the interrelation of online hate speech with the physical world, in a range of simple dissemination to absolute influence. Towards this direction, Müller and Schwarz (2018, 2020) have conducted research on the causal effect of social media on the propagation of hate sentiments offline, whether anti-refugee (Müller and Schwarz, 2018) or more specifically anti-Muslim (Müller and Schwarz, 2020). In (Müller and Schwarz, 2018), the authors provide evidence that there is significant association between negative content against refugees existing on Facebook and offline hate crimes in Germany on a municipal level, while controlling for multiple potentially confounding factors such as German municipalities' characteristics and overall social media usage. To reach to this conclusion

¹This direction has also been studied by (Scharwächter and Müller, 2020), however they merely calculate correlations and do not control for potential confounders, hence are not included in the present survey.

they combine a variety of data sources; online anti-refugee sentiment is represented by content from a widespread Facebook page of a German right-wing party, which hosts plenty of far-right content; controlling for the network's popularity in Germany, the authors measure Facebook outages and internet disruptions; finally, to further measure user activity on the network and create controls based on a neutral subject, they explore another broadly popular page of a famous commercial product. The causal framework they implement is a fixed-effects regression model (inspired by [Bartik \(1991\)](#)), which considers the aforementioned panel data combined with a range of controls. They discover that the effect is stronger in areas with higher Facebook usage and demonstrate a robustly strong connection between the activity of the right-wing group and severe hate crimes.

Similarly, in ([Müller and Schwarz, 2020](#)), they study the causal impact of Islamophobic social media content on registered crimes and overall negative sentiment against Muslims, and whether former US president Donald Trump's Twitter campaign has contributed to the propagation of Islamophobia. To ensure validity and robustness of their findings, they fuse a number of different data sources and employ a difference-in-differences approach. The data originate mainly from Twitter for the social media information, a survey by FBI to discover hate crimes, data from mass media, demographic information about US counties etc. Their findings provide evidence to associate a 38% larger increase in hate crimes, between 2010 and 2017, with higher exposure to social media. Moreover, consistently with previous research, they also provide evidence which shows a connection amid the start of Trump's presidential campaign with an increase of anti-Muslim sentiments in USA. Both projects illustrate there is strong evidence linking OHS with offline occurrences of hate-related crimes, with the former having a causal effect on the latter.²

Looking at the opposite direction, online hate speech is frequently affected by events taking place in the offline world. For example, following the September 11th, 2001 attack to the twin towers of

²It goes without saying that results from papers that employ causality need to be interpreted with great care. As [Müller and Schwarz \(2018\)](#) emphasise, for instance, their findings do not indicate that social media can cause crimes against minorities, but rather "social media can act as a propagating mechanism [...]" so that "shifts in exposure to anti-refugee sentiment on social media can increase the number of anti-refugee attacks".

New York City, there seems to be an increase in Islamic terrorist attacks as much as an increase in Islamophobia. It is very well expected that this will affect the online world and the way Muslims are perceived, which seems to actually be the case according to [Olteanu et al. \(2018\)](#). In ([Olteanu et al., 2018](#)) they analyse the influence of offline events on online hate speech. More specifically, they study the impact of several attacks related to the Islamic State – both Islamic terrorist attacks and Islamophobic ones – in terms of online hate speech and counter-hate speech. Their findings indicate that terrorist attacks show an increase in OHS, especially towards Muslims and Arabs.

To calculate the causal effect, they construct time series from Reddit and Twitter – representing the number of posts and unique users involved with the event – and then synthesize counterfactual time series for the same period of time, such as they would be produced had the attacks not happened. The counterfactual data are created by composing timelines of the same event, while adjusting for a temporal shift prior to the event, so that the time series will reflect a similar period of time but in different time windows. The produced timelines, put together with other external data sources, are then fit into a state space model ([Brodersen et al., 2015](#)) using maximum likelihood estimation, to predict the synthesized control time series. Comparing the treatment and control series, they calculate relative effects for a number of manually curated terms, which are also annotated across four hate speech dimensions (stance, targets, severity, and framing).

2.2 Actors and Targets

Beyond broadly looking at the overall impact of OHS, the next step would be to concentrate on the individuals and speculate how OHS affects and is affected by the participating members, whether these are at the producing or receiving end of such content. For example, in the aforementioned works, [Olteanu et al. \(2018\)](#) and [Müller and Schwarz \(2020\)](#) have made some general remarks regarding OHS, but in order to effectively focus on the task they narrow the type of hate speech to be racism and, more specifically, Islamophobia. Emphasis, however, is given on understanding the dynamics of diffusion and not on studying the impact on individuals who support Islam.

On these grounds, [Saha et al. \(2019a\)](#) study the effect of hate prevalence on the stress levels of

US college students, within Reddit college communities. To measure the levels of toxicity, they employ hate lexicons, with the help of which they compute the College Hate Index (CHX), as fractions of hateful keywords in each community compared to other subreddits banned for violating the hate-limiting policies of Reddit. Similarly, they quantify the exposure of users to hateful content based on the threads they have participated and to account for their stress levels they use a binary classifier based on existing models. They examine numerous observable confounders – such as the subreddit and user activity – and apply propensity score matching to calculate the causal effect whilst controlling for the covariates. Their results demonstrate an increase in stress expression caused by exposure to hateful speech.

Additionally, inherent characteristics of online environments, such as anonymity, ease of access, and size of audience (Brown, 2018) are highly likely to affect the behaviour of online social networks' users and sometimes make it easier for them to misbehave. von Essen and Jansson (2020) discuss the outcomes of the by-products of online world characteristics – in this case anonymity – on hate speech actors. In particular, they compare the degree of hatefulness before and after the identities of a large set of users from Flashback, an anonymous Swedish discussion platform similar to Reddit, have been publicly exposed. Their hypothesis suggests that once running the risk of exposure, users decrease the volume of hateful content they post. To detect hate, they implement a machine learning model and make predictions on the data, which will afterwards be used with a difference-in-differences approach to make causal claims. According to their estimates, the reduction of anonymity, as in risk of exposure, leads to a decline in general hate and overall activity, and even more on xenophobic content. Surprisingly, levels of misogyny increase. These empirical findings mostly support the author's original hypothesis.

2.3 Interventions

Last but not least, perhaps the most extensively studied of all three pillars is the effect of interventions, possibly due to its relatively close connection to causality. Interventions here refer to actions that are taken towards the elimination of the phenomenon; such strategies include quarantining (Chandrasekharan et al., 2020), banning (Chan-

drasekharan et al., 2017; Thomas et al., 2021), censoring (Álvarez-Benjumea and Winter, 2018), sanctioning (Munger, 2017), and counter speaking (Álvarez-Benjumea and Winter, 2018). Reversions of each policy vary, depending on the platform and the methodology followed.

Chandrasekharan et al. (2017) first studied the effect of banning an entire hateful community using causal inference methods. More specifically, they investigate through a quasi-experiment how banning targeted communities influenced hate speech levels on Reddit. Initially, they investigate the activity of the participants, post-banning the examined subreddits, with respect to activity level and hateful content volume. To control for potential confounders related to user characteristics – such as the activity, popularity, or age of user accounts – they employ Mahalanobis Distance Matching (Rubin et al., 2006) between treatment and control users, which is then further enhanced with a difference-in-differences analysis (Abadie, 2005) of the two groups over time. Subsequently, they inspect the level of hate propagation to other *invaded* (as they call them) subreddits. In this case, matching needs to be applied on a subreddit level, rather than user level, so instead they employ the Interrupted Time Series approach (Bernal et al., 2017). Their results show that the ban worked *for Reddit*, meaning that having targeted a particular area of the platform, they have successfully eliminated hateful content without conveying the problem elsewhere within Reddit.

In addition to this work, the authors also studied the causal effect of *quarantining* Reddit communities, in a similar experimental framework (Chandrasekharan et al., 2020). Quarantining is a form of intervention that Reddit applies, where communities are indicated as potentially problematic and users have to deliberately choose to enter them, after being warned about toxicity levels within. This approach is less stern than banning, however according to the findings of this study, it is still effective to restrict the influx of users in hostile communities, while preserving the freedom of speech. In this case, the causal inference strategy used is the Interrupted Time Series regression (Bernal et al., 2017), which models interruptions caused by the treatment variable.

Paper	Domain	Treatment	Outcome	Controls	Causal Method
Chandrasekharan et al. (2017)	Reddit	Banning subreddits	Former members' activity, language & migration trends	Subreddits that could potentially have been banned	MDM, DiD & ITS Regression Analysis
Munger (2017)	Twitter	Social sanctioning, varying by influence & identity	Race-based harassment	Gender & race of harassers, and anonymity	Randomized Control Experiment
Olteanu et al. (2018)	Twitter & Reddit	Islamophobic & Islamist terrorist attacks in Western countries	Hateful content	Synthetic counterfactual time series	Comparison of observed vs counterfactual time series
Álvarez-Benjumea and Winter (2018)	Custom forum	Censoring & counter-speaking	Hate speech score	Number of comments	Experimental approach
Müller and Schwarz (2018)	Facebook	Anti-refugee Facebook group	Hate crimes against refugees in Germany	Characteristics of German municipalities	Fixed effects panel regressions
Saha et al. (2019a)	Reddit	Hateful speech in college subreddits	Online stress levels	Subreddit & User activity	Propensity score matching & DiD Regression Analysis
Müller and Schwarz (2020)	Twitter	High Twitter usage	Anti-Muslim hate crimes	Characteristics of US counties	DiD approach
von Essen and Jansson (2020)	Flashback	Web anonymity	Hateful posts	Risk of web exposure	DiD approach
Chandrasekharan et al. (2020)	Reddit	Quarantining subreddits	Member participation & language, and new member influx	Subreddits that could potentially have been quarantined	ITS & Bootstrapping tests
Thomas et al. (2021)	Reddit	External events & regulatory actions (ie. bans)	Member activity & attrition	User participation	MDM & DiD analysis & Multivariate Bayesian changepoint analysis
Ananthkrishnan and Tucker (2021)	YouTube	Presence of hate speech	Propagation (or "virality") of hate speech	Popularity, quality and posting time of videos	Instrumental Variables

Table 1: Following the format of (Keith et al., 2020), this table summarises research papers related to causal effects of web hostility. The abbreviations used stand for the following; Difference-in-Differences (DiD) (Abadie, 2005); Interrupted Time Series (ITS) (Bernal et al., 2017); Mahalanobis Distance Matching (MDM) (Rubin et al., 2006).

3 Difficulties and Causal Methods

Experimentation for the sole purpose of causality is fairly costly and, in this case, can be largely unethical, considering that it would mean deliberate exposure of people to toxic material. As a result, it is crucial to exploit the available methodologies of observational causality, in combination with its intersection with natural language processing. There are, however, several open research problems which lie at the heart of textual causality and thus are inherited by this task. A prominent hurdle, for instance, exists in the conception of the causal diagram, contemplating all the possible covariates and determining their roles. Such graphs are domain-specific and can greatly deviate, due to differences introduced by the platforms or niche

types of speech. Furthermore, there are challenges such as the discovery of effective linguistic representations for the various parts where text can act as a surrogate.

In this section we discuss some prevailing difficulties, summarise the route of relevant literature, and finally, provide some suggested directions based on previous studies.

3.1 Confounding Bias

To establish true causation among a target variable X and an outcome Y , there should not be any indirect connections influencing the effect of X on Y . In reality it is rarely the case that this will hold organically true, hence being aware of any confounding factors is fundamental in order to intervene

accordingly and eliminate their impact. Many confounders can be detected after careful exploration of the domain and quantified through unadjusted versus adjusted estimates. These are called *observable* confounders and are largely discovered based on domain expertise or previous research. It is well possible, however, that not all such factors can be anticipated; there are cases where they might not be discoverable or their computations might not be feasible.

One very relevant work on treating *latent* confounders is the one of Cheng et al. (2019), who attempt to create a somewhat causally-aware cyberbullying detection model for observed data, making it robust on confounding bias by controlling for latent (or plausible, as they call them) confounders. To achieve that, they look for pairs of variables that demonstrate Simpson's paradox and then employ a clustering algorithm to group the data based on the discovered p-confounders. They afterwards perform classification within the clusters. Their proposed framework of detecting latent confounding factors is potentially generalisable outside the scope of this topic. Nevertheless, to the best of our knowledge, this is the only research project considering cyberbullying detection from a causal prism.

Suggestions While controlling for confounders, and in case they are categorical or numerical variables, it is possible to follow well established techniques such as matching (De Graaf et al., 2011). In simple terms, matching is a method to create pairs of samples from the same category and of different outcomes, to approximate randomised conditions and obliterate confounding bias so that any effect that remains must be realistic. For example, Zhang et al. (2018) perform their task while controlling for topical confounding, by creating pairs of "good" and "bad" conversations from the same Wikipedia page. This way, they ensure that any differences caused by the nature of the Wikipage topic are removed. However, it is not always possible to know the confounding factors. In that case, it is more useful to consider ways of eradicating confounding in a holistic way, such as the clustering approach of (Cheng et al., 2019).

Additionally, it is highly likely for features of this language-dominated topic to be textual. In that case, treating the covariates can be fairly challenging. Drawing inspiration from the textual causality literature (Keith et al., 2020), there are a number of

approaches both for observable and latent textual confounders. For the former, it is possible to follow the strategies that will be described in Section 3.2 and extract features that best represent each specific factor. The case of latent confounders, though, is more elaborate. One suggestion proposed by Landeiro et al. (2019) is the use of adversarial learning as a combating method to confounding shift, although it has yet to be examined within the context of OHS.

3.2 Linguistic Representations

Most of the body of literature described in Section 2 base their research on numerical features, regression, and time series analysis and do not look at the actual content of the social media platforms or question the validity of hatefulness. Contrary to traditional approaches to hate speech detection, which mostly work with text, these causal analyses largely overlook language. One possible explanation for this could be the inherent difficulty of constructing representative vectors that will capture the verbal essence. Language as a means of communication is intrinsically high-dimensional and reducing it in lower dimensions to allow for processing requires strenuous effort.

Saha et al. (2019a) and Cheng et al. (2019) are the only previous papers, to the best of our knowledge, to exploit the texts and attempt to represent them with some feature vectors, to produce causal representations of OHS: frequency of hate keywords in the first case, and Linguistic Inquiry and Word Count (LIWC –Tausczik and Pennebaker (2010)) in the second. There is great need for experimentation and improvement to attempt more holistic approaches to the problem.

Suggestions Since this task is significantly understudied in the context of online hate speech, there are a lot of possible directions to follow. For example, it might prove meaningful to extract high-level representations like sentiment features or toxicity scores. The latter is supported by von Essen and Jansson (2020), who claim that "hate begets hate", meaning that hateful content triggers replies of the same style. Furthermore, one could use the framework implemented by Pryzant et al. (2018), which automatically induces representative lexicons for social science tasks, while controlling for potential confounding factors. Having previously established some factors as such, it is possible to employ this framework and observe the performance of the

lexicon. Lastly, it could prove fruitful to explore pre-trained language representation models that are focused on capturing profanity (e.g. HateBERT (Caselli et al., 2020)) or causality (e.g. CausalBERT (Khetan et al., 2020; Veitch et al., 2020)), or even attempt to fine-tune a hate- and causal-specific version of the original BERT model (Devlin et al., 2018), for a combination of the two. The latter would probably be very demanding with respect to the need for training data, albeit promising.

3.3 Causal Transportability

Finally, a big benefit of true causal knowledge is the ability to transfer this knowledge across domains and achieve comparably optimal performance (Pearl and Bareinboim, 2011). Such an achievement can then work as an evaluation method, showing that the exhibited causal attributions are, in fact, correct and sufficient. For example, consider racist material, which is possibly similar across social networks. A detection model that takes into account any previously attributed causal factors, and controls for all discovered confounders, should be able to train with data from one network and exhibit satisfactory performance on a different network. (Cheng et al., 2019) is the only existing study that is attempting to address causal transferability in the context of this topic, by experimenting on one domain and testing on a different, switching between Twitter and Formspring. Despite the challenging task, it would be of great interest to apply this evaluation method across alternate platforms, various labels, or even multiple languages.

4 Conclusions

The task of hate speech mitigation can never be conclusively accomplished unless the research community seriously speculates about its roots. This would mean a methodological experimentation with the phenomenon, from a causal prism. Simply put, it would mean taking a step back from its detection and focusing on scientifically breaking it into its causes. If the causes are known, then any further steps can make meaningful change. Practitioners should then be able to take crucial actions towards either its efficient and interpretable detection or, even further, its *a priori* prevention. By understanding the most prominent and fundamental factors leading to the phenomenon, it is possible to conclusively build on its holistic solution.

The present survey is intended to initiate discussions towards this direction. By introducing some of the most relevant studies while simultaneously highlighting several points of interest, we aspire for this paper to work as a point of reference, as much as an inspiration and an open call for further research. To make any definite causal claims it is necessary to have a well defined system, where all influential factors are understood or known. All the existing research, including this survey, can only serve as a prelude to future discussions regarding this ubiquitous issue, which will consequently lead to more sophisticated and generalisable models.

References

- Alberto Abadie. 2005. Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19.
- Amalia Álvarez-Benjumea and Fabian Winter. 2018. Normative change and culture of hate: An experiment in online environments. *European Sociological Review*, 34(3):223–237.
- Uttara M Ananthakrishnan and Catherine E Tucker. 2021. The drivers and virality of hate speech online. Available at SSRN 3793801.
- Lora Aroyo, Lucas Dixon, Nithum Thain, Olivia Redfield, and Rachel Rosen. 2019. Crowdsourcing subjective tasks: the case study of understanding toxicity in online discussions. In *Companion proceedings of the 2019 world wide web conference*, pages 1100–1105.
- Timothy J Bartik. 1991. Boon or boondoggle? the debate over state and local economic development policies.
- James Lopez Bernal, Steven Cummins, and Antonio Gasparri. 2017. Interrupted time series regression for the evaluation of public health interventions: a tutorial. *International journal of epidemiology*, 46(1):348–355.
- Kay H Brodersen, Fabian Gallusser, Jim Koehler, Nicolas Remy, and Steven L Scott. 2015. Inferring causal impact using bayesian structural time-series models. *The Annals of Applied Statistics*, 9(1):247–274.
- Alexander Brown. 2018. What is so special about online (as compared to offline) hate speech? *Ethnicities*, 18(3):297–326.
- Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2020. Hatebert: Retraining bert for abusive language detection in english. *arXiv preprint arXiv:2010.12472*.

- Eshwar Chandrasekharan, Shagun Jhaver, Amy Bruckman, and Eric Gilbert. 2020. Quarantined! examining the effects of a community-wide moderation intervention on reddit. *arXiv preprint arXiv:2009.11483*.
- Eshwar Chandrasekharan, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert. 2017. You can't stay here: The efficacy of reddit's 2015 ban examined through hate speech. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–22.
- Lu Cheng, Ruocheng Guo, and Huan Liu. 2019. Robust cyberbullying detection with causal interpretation. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 169–175.
- Munmun De Choudhury and Emre Kiciman. 2017. The language of social support in social media and its effect on suicidal ideation risk. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11.
- Michiel A De Graaf, Kitty J Jager, Carmine Zoccali, and Friedo W Dekker. 2011. Matching, an appealing method to avoid confounding? *Nephron Clinical Practice*, 118(4):c315–c318.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Son Doan, Elly W Yang, Sameer S Tilak, Peter W Li, Daniel S Zisook, and Manabu Torii. 2019. Extracting health-related causality from twitter messages using natural language processing. *BMC medical informatics and decision making*, 19(3):71–77.
- Paula Fortuna and Sérgio Nunes. 2018. A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30.
- John Hicks et al. 1980. *Causality in economics*. Australian National University Press.
- Katherine A Keith, David Jensen, and Brendan O'Connor. 2020. Text and causal inference: A review of using text to remove confounding from causal estimates. *arXiv preprint arXiv:2005.00649*.
- Vivek Khetan, Roshni Ramnani, Mayuresh Anand, Shubhashis Sengupta, and Andrew E Fano. 2020. Causal bert: Language models for causality detection between events expressed in text. *arXiv preprint arXiv:2012.05453*.
- Virgile Landeiro, Tuan Tran, and Aron Culotta. 2019. Discovering and controlling for latent confounds in text classification using adversarial domain adaptation. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pages 298–305. SIAM.
- Jürgen Landes, Barbara Osimani, and Roland Poellinger. 2018. Epistemology of causal inference in pharmacology. *European Journal for Philosophy of Science*, 8(1):3–49.
- Sean MacAvaney, Hao-Ren Yao, Eugene Yang, Katina Russell, Nazli Goharian, and Ophir Frieder. 2019. Hate speech detection: Challenges and solutions. *PloS one*, 14(8):e0221152.
- Margaret Mooney Marini and Burton Singer. 1988. Causality in the social sciences. *Sociological methodology*, 18:347–409.
- Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2019. Tackling online abuse: A survey of automated abuse detection methods. *arXiv preprint arXiv:1908.06024*.
- Karsten Müller and Carlo Schwarz. 2018. Fanning the flames of hate: Social media and hate crime. *Journal of the European Economic Association*.
- Karsten Müller and Carlo Schwarz. 2020. From hashtag to hate crime: Twitter and anti-minority sentiment. Available at SSRN 3149103.
- Kevin Munger. 2017. Tweetment effects on the tweeted: Experimentally reducing racist harassment. *Political Behavior*, 39(3):629–649.
- Alexandra Olteanu, Carlos Castillo, Jeremy Boy, and Kush Varshney. 2018. The effect of extremist violence on hateful speech online. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Judea Pearl and Elias Bareinboim. 2011. Transportability of causal and statistical relations: A formal approach. In *Twenty-fifth AAAI conference on artificial intelligence*.
- Reid Pryzant, Kelly Shen, Dan Jurafsky, and Stefan Wagner. 2018. Deconfounded lexicon induction for interpretable social science. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1615–1625.
- Kelly Quinn. 2018. Cognitive effects of social media use: A case of older adults. *Social Media+ Society*, 4(3):2056305118787203.
- Julian Reiss. 2009. Causation in the social sciences: Evidence, inference, and purpose. *Philosophy of the Social Sciences*, 39(1):20–40.
- Donald B Rubin, Elizabeth A Stuart, et al. 2006. Affinely invariant matching methods with discriminant mixtures of proportional ellipsoidally symmetric distributions. *The Annals of Statistics*, 34(4):1814–1826.
- Federica Russo and Jon Williamson. 2011. Epistemic causality and evidence-based medicine. *History and philosophy of the life sciences*, pages 563–581.

- Koustuv Saha, Eshwar Chandrasekharan, and Munmun De Choudhury. 2019a. Prevalence and psychological effects of hateful speech in online college communities. In *Proceedings of the 10th ACM Conference on Web Science*, pages 255–264.
- Koustuv Saha, Benjamin Sugar, John Torous, Bruno Abrahao, Emre Kiciman, and Munmun De Choudhury. 2019b. A social media study on the effects of psychiatric medication use. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, pages 440–451.
- Haji Mohammad Saleem, Kelly P Dillon, Susan Benesch, and Derek Ruths. 2017. A web of hate: Tackling hateful speech in online social spaces. *arXiv preprint arXiv:1709.10159*.
- Joni Salminen, Fabio Veronesi, Hind Almerakhi, Soon-Gvo Jung, and Bernard J Jansen. 2018. Online hate interpretation varies by country, but more by individual: A statistical analysis using crowdsourced ratings. In *2018 Fifth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, pages 88–94. IEEE.
- Erik Scharwächter and Emmanuel Müller. 2020. Does terrorism trigger online hate speech? on the association of events and time series. *The Annals of Applied Statistics*, 14(3):1285–1303.
- Anna Schmidt and Michael Wiegand. 2017. A survey on hate speech detection using natural language processing. In *Proceedings of the fifth international workshop on natural language processing for social media*, pages 1–10.
- Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54.
- Pamela Bilo Thomas, Daniel Riehm, Maria Glenski, and Tim Weninger. 2021. Behavior change in response to subreddit bans and external events. *arXiv preprint arXiv:2101.01793*.
- Alice Tontodimamma, Eugenia Nissi, Annalina Sarra, and Lara Fontanella. 2021. Thirty years of research into hate speech: topics of interest and their evolution. *Scientometrics*, 126(1):157–179.
- Victor Veitch, Dhanya Sridhar, and David Blei. 2020. Adapting text embeddings for causal inference. In *Conference on Uncertainty in Artificial Intelligence*, pages 919–928. PMLR.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. Challenges and frontiers in abusive content detection. Association for Computational Linguistics.
- Emma von Essen and Joakim Jansson. 2020. Misogynistic and xenophobic hate language online: a matter of anonymity.
- Gerald Young. 2016. Unifying causality and psychology. *Cham: Springer*.
- Justine Zhang, Jonathan P Chang, Cristian Danescu-Niculescu-Mizil, Lucas Dixon, Yiqing Hua, Nithum Thain, and Dario Taraborelli. 2018. Conversations gone awry: Detecting early signs of conversational failure. *arXiv preprint arXiv:1805.05345*.