

Extracting Age-Related Stereotypes from Social Media Texts

Kathleen C. Fraser, Svetlana Kiritchenko, Isar Nejadgholi

National Research Council Canada

Ottawa, Ontario, Canada

{kathleen.fraser, svetlana.kiritchenko, isar.nejadgholi}@nrc-cnrc.gc.ca

Abstract

Age-related stereotypes are pervasive in our society, and yet have been under-studied in the NLP community. Here, we present a method for extracting age-related stereotypes from English language Twitter data, generating a corpus of 300,000 over-generalizations about four contemporary generations (baby boomers, generation X, millennials, and generation Z), as well as “old” and “young” people more generally. By employing word-association metrics, semi-supervised topic modelling, and density-based clustering, we uncover many common stereotypes as reported in the media and in the psychological literature, as well as some more novel findings. We also observe trends consistent with the existing literature, namely that definitions of “young” and “old” age appear to be context-dependent, stereotypes for different generations vary across different topics (e.g., work versus family life), and some age-based stereotypes are distinct from generational stereotypes. The method easily extends to other social group labels, and therefore can be used in future work to study stereotypes of different social categories. By better understanding how stereotypes are formed and spread, and by tracking emerging stereotypes, we hope to eventually develop mitigating measures against such biased statements.

Keywords: stereotypes, bias, ageism, social media, computational social science

1. Introduction

Stereotypes are over-generalizations about the characteristics of a group of people, such that an individual is assumed to have these characteristics simply based on their perceived membership in the group. Stereotypes can lead to prejudicial behaviour against members of a group, as well as psychological harm. Not all stereotypes are overtly negative; however, even stereotypes that might appear on the surface to be positive can have negative effects (for example, the stereotype that women are warm and nurturing can lead to backlash against women who violate this stereotype by acting powerful and assertive).

Of particular concern is stereotyping on the basis of protected characteristics, such as race, sex, religion, or age. The latter has been particularly under-studied in natural language processing (NLP), despite the fact that ageism is widespread in North American society and can lead to age-based bias in the workplace (Perry et al., 2013), media representation (Lichtenstein, 2021), and the healthcare system (Wyman et al., 2018). Furthermore, such stereotypes can become a self-fulfilling prophecy when they are internalized by people who self-identify as older adults, leading to isolation and health decline (Chan et al., 2020).

Unlike the traditional approach to studying stereotypes in psychology, where a small group of people is directly asked to provide their own or common perceptions of a group of interest via survey questions, in this work we analyze pervasive stereotypes from naturally-occurring data in Twitter. While we lack the benefits of a well-controlled laboratory environment and need to apply various steps to collect, filter, and clean the data, we have access to spontaneously expressed opinions

of thousands of individuals. This allows us to detect less common or emerging stereotypes, without being limited by our *a priori* expectations of social stereotypes. Studying stereotypes occurring in real social interactions can contribute to our understanding of how stereotypes are formed and spread, and their impact on target groups and inter-group relations. Identifying and monitoring the dynamics of currently pervasive group perceptions is a necessary first step before intervening with educational, counter-narrative, and other mitigating measures.

While age can be subjectively described in relative terms (e.g., “old” and “young”), another common method of cohorting age groups is by *generation*, defined according to birth year. Here, we focus on both age-based and generational stereotypes and create a corpus of 300,000 English sentences that express opinions about six age-related groups: *baby boomers* (born between 1946 and 1964), *generation X* (born between 1965 and 1980), *millennials* or *generation Y* (born between 1981 and 1996), *generation Z* (born between 1997 and 2012)¹, *young* people, and *old* people. The data have been collected from Twitter over the period of three months. Unlike previous NLP works, which primarily used supervised methods to detect negative stereotypes, we apply a range of unsupervised and semi-supervised language analysis techniques to explore the data, and find various favorable and unfavorable over-generalizing statements about the groups.² Using word-association metrics at the group

¹<https://www.beresfordresearch.com/age-range-by-generation/>

²Note that while some of these stereotypes are reported here, the authors in no way endorse or support these views.

level to determine the words which occur more frequently with one age group than the others, we uncover common stereotypes as reported in the popular media. We observe only partial overlap between age stereotypes (young/old) compared to generational stereotypes, as predicted by psychological studies (Perry et al., 2013). We then use semi-supervised topic modelling and density-based clustering to determine the most highly-frequent opinions about each group across various areas of life (family and friends, finance, work, politics, technology, and health). We observe that age and generational stereotypes are context-dependent and vary across different life domains. The data suggest that what counts as “old” varies across different topics, also confirming previous questionnaire-based studies (Kornadt and Rothermund, 2011). Thus, our analyses demonstrate our ability to detect age-related stereotypes “in-the-wild” from naturally-occurring free text, and raise the possibility that similar methods may be used for future work on stereotyping of other social groups.

2. Background and Related Work

We first discuss findings on age-related stereotyping in psychology, and then overview work on discovering and mitigating stereotypical biases in NLP applications.

2.1. Studies of Age-Related Stereotypes in Psychology

Toomey and Rudolph (2017) define age-related stereotypes as “overgeneralized expectations and beliefs about the characteristics and traits of individuals on the basis of age.” Age-based categorization is a cognitive decision made quickly based on physical appearance, similar to gender- and race-based categorization (Blaine and Brechley, 2017; Brewer and Lui, 1989). Older adults are usually associated with more negative and more strongly activated stereotypes than younger people (Hummert et al., 1995; Chasteen et al., 2002). A common system of age-based categorization is the idea of social *generations*. Traditionally, generations are operationalized as “birth cohorts,” grouping people into generational categories based on their year of birth. Belonging to the same birth cohort often means living through similar social and historical processes and events and sharing common experiences, which shapes collective memories and common value systems and behaviors (Mannheim, 1952; Lyons and Kuron, 2014). However, not all individuals from the same birth cohort are influenced in the same way by the historical and social events of their formative years. A generational identity is formed when an individual socially identifies with their generation through a shared value system (Joshi et al., 2010). Social self-categorization and categorization of others may lead to stereotyping, often favoring the in-group at the expense of the out-group (Tajfel and Turner, 1979).

While some studies treat *age* and *generation* interchangeably, other studies have found striking differences between age-related and generational stereotypes. For example, Perry et al. (2013) studied stereotypes of older and younger workers; they found moderate overlap between the stereotypes for “older” workers and “boomers”, but smaller overlap between the stereotypes of “younger” workers and “millennials”. Additionally, the valence of the stereotypes was quite different, with older workers stereotyped as less productive and motivated, while boomers were stereotyped as career-driven, hard-working, and competitive. Furthermore, while the definition of the generations is static, Kornadt and Rothermund (2011) showed that our perception of what counts as “old” and “young” depends on multiple factors, including the age of the perceiver and the context (e.g., the workplace versus the family). Thus, in our work here, we explore both generational stereotypes (for boomers, gen-X, millennials, and gen-Z), as well as for older and younger adults.

Several studies have examined the generational differences in personality, work values, leadership and teamwork preferences, and career patterns, yet the empirical findings provide mixed evidence about the *actual* differences among the generations (Macky et al., 2008; Wong et al., 2008; Costanza et al., 2012; Lyons and Kuron, 2014; Becton et al., 2014). Nevertheless, *perceived* differences and stereotypes about generational groups persist and sometimes are enacted as self-fulfilling prophecies (Perry et al., 2013; Van Rossem, 2019b; Van Rossem, 2021). Meta-stereotypes—individuals’ beliefs about what other generational groups think about their own group—further affect attitudes and influence generational interactions (Van Rossem, 2019a). Therefore, regardless of whether the perceived differences among the generations actually exist, the impact of generational stereotypes on inter-generational relations can be significant.

These psychological studies are typically based on surveys directly asking a limited number of participants about their own or commonly known perceptions of age-related groups. In this study, we take a different approach and analyze social media texts written by thousands of individuals for perceived views of different age groups. Using social media text to analyze social cognitive behaviour has received growing support in the psychological community, as it has the benefits of being more unobtrusive, unconstrained, and ecologically valid than many laboratory studies (Nicolas et al., 2021) and can enable large quantities of data to be collected over extended timeframes (Meshi et al., 2015). However, it is also much less controlled than a typical laboratory study; for example, we do not have demographic or geographic information for the set of Twitter users included here.

2.2. Studies of Stereotypical Associations and Biases in NLP

In the past few years, much effort in NLP has been devoted to detecting and mitigating stereotypical associations encoded in word embeddings (Bolukbasi et al., 2016; Caliskan et al., 2017), large-scale language models (May et al., 2019; Bartl et al., 2020; Liang et al., 2021), and full models developed for various NLP applications, such as machine translation (Stanovsky et al., 2019; Prates et al., 2020), coreference resolution (Rudinger et al., 2018; Zhao et al., 2018), sentiment analysis (Kiritchenko and Mohammad, 2018; Thelwall, 2018), and others. Since these models are learned from human-generated texts, they inevitably inherit, and sometimes amplify, human-like biases, which can result in harmful outcomes when the models are deployed in real-life settings. Gender and racial biases have received the most attention (Kurita et al., 2019; Manzini et al., 2019; Kirk et al., 2021), yet other types of biases, including biases related to age (Díaz et al., 2018) and disabilities (Hutchinson et al., 2020; Venkit and Wilson, 2021), have also been studied. For example, Díaz et al. (2018) examined widely-used unsupervised, lexicon-based and supervised, corpus-based sentiment analysis systems and found significant differences in their outputs for sentences containing words related to young and old age.

NLP studies of *human* stereotypes and biases (as opposed to *system* biases) have mostly focused on abusive language and hate speech detection (Schmidt and Wiegand, 2017; Vidgen et al., 2019). Stereotypical perceptions of members of marginalized groups can lead to subtle or overtly abusive behavior in online communication and have a detrimental effect on the target groups. Supervised classification techniques have been successful at identifying abusive language containing explicitly insulting and derogatory expressions, but have experienced difficulty at uncovering subtle forms of online abuse, such as stereotyping, microaggression, condescension, etc. (Breitfeller et al., 2019; Wang and Potts, 2019; Price et al., 2020; Wiegand et al., 2021). Further, subtle forms of abusive language are also challenging for human annotators who bring in their own biases into the annotation process (Binns et al., 2017). Sap et al. (2020) trained a generative model in an attempt to automatically uncover the stereotypes implied by abusive utterances. They used crowd-sourcing to annotate a large corpus of abusive comments for the implied stereotypical associations, and showed that contemporary language models cannot yet effectively reproduce human interpretations of implied meaning behind abusive expressions.

Various unsupervised techniques have also been applied on textual corpora to uncover human stereotypes. Fraser et al. (2022) proposed a computational method to project stereotypical sentences from social media onto the two theoretically-motivated dimensions of warmth and competence, as described by the Stereo-

type Content Model (Fiske, 2018). Through examining biases encoded in word embeddings trained on texts from different time periods, Garg et al. (2018) showed the evolution of gender and ethnic stereotypes in the US during the 20th and 21st centuries. In a similar way, Charlesworth et al. (2021) investigated gender stereotypes in various types of texts, including child and adult speech, TV shows, and books. Marzouki et al. (2020) employed sentiment and word co-occurrence analysis to study the evolution of online stereotypes related to Islam, while Fokkens et al. (2018) examined the language used to describe Muslims in Dutch media to analyse various aspects of stereotyping. Rao and Taboada (2021) used topic modeling on a stream of news articles to show unequal gender representation in respondents quoted for different topics. In contrast to some of these works, which train word embedding models on the collected data (thus necessitating very large corpora and significant computational resources), our method extracts inherently interpretable over-generalizations directly from the data.

Detecting stereotypes in text is the first step towards developing mitigating strategies. For example, recent NLP work has shown promise in automatically generating counter-speech or counter-narrative responses to online abusive remarks (Mathew et al., 2018; Chung et al., 2019). Specifically with respect to stereotypes, Fraser et al. (2021) demonstrated how positive ‘anti-stereotypes’ can be generated from negative stereotypes, in an effort to counter and reduce stereotypical biases expressed online. Furthermore, identifying and tracking societal biases can inform institutional policies to reduce the negative consequences of stereotyping (e.g., *stereotype threat*, see Spencer et al. (2016)), such as addressing environmental cues, offering diversity training, and changing organizational mindsets (Casad and Bryant, 2016).

3. Data Collection and Processing

We collected tweets mentioning six age-related groups: baby boomers, generation X, millennials, generation Z, older adults, and young people. The collection was performed using the Twitter API over a period of three months, from August 20, 2021 to November 20, 2021. As search queries for the four generations, we selected terms representing the variations on the group names (e.g., *boomers*, *gen xers*, *gen z*) as well as some common misspellings (e.g., *millennials*).³ For older adults, we used the terms *elderly*, *elderly people*, *elderly folks*, *elderly persons*, *old people*, *old folks*, *old persons*, and *senior citizens*. For young people, the terms included *young people*, *young folks*, and *young persons*.

To focus on personal opinions of individual Twitter users, we discard re-tweets, tweets with URLs to external websites, and tweets with five or more hashtags. We also discard all tweets from Twitter accounts that

³The full list of query terms is available at <https://svkir.com/projects/age-stereotypes.html>.

Target group	Number of sentences	Avg. number of words per sentence
Boomers	56,266	15.71
Gen-X	8,986	14.70
Millennials	30,698	15.31
Gen-Z	31,995	15.78
Older adults	93,108	16.62
Young people	93,295	19.36

Table 1: The number of extracted sentences for each target group.

include the words *bot*, *boomer(s)*, *millennial(s)*, *millennial(s)*, *millennial(s)*, *millennial(s)*, *old*, *elderly*, or *senior* in the user name or screen name, and accounts with more than 3 posts per day on average for the posts mentioning the four generations and with more than 5 posts per day on average for posts mentioning old or young people query terms in our 3-month collection. These heuristics help reduce the amount of irrelevant texts generated by bots as well as ads, news headlines, and promotional campaigns written by organizations. Still, in our preliminary experiments, we noticed that a few irrelevant tweets repeated many times negatively affected the analysis. Therefore, we also remove identical and near-identical copies of tweets from the corpus. Often, tweets mentioning the group of interest actually express opinions not about the target group, but rather about other people or events. Therefore, we further filter our data and retrieve only sentences where the target group is a nominal subject of the main or a subordinate clause of the sentence. We separate tweets into sentences and perform dependency parsing using the spaCy library.⁴ Sentences where a target group is described with qualifiers referring only to some members of the group (e.g., *some*, *these*, *several*, *few*) are excluded from the analysis. Further, we discard sentences where different target groups (e.g., boomers and millennials) are discussed together. Finally, we remove URLs (to embedded images and videos) and user mentions, and mask the query terms (e.g., *boomers*, *millennials*) with the generic token ‘<target>’ in the remaining sentences to reduce any impact of the group name on the word and cluster analysis. Table 1 presents the statistics on the extracted sentences for our target groups of interest.

4. Methods

We first look at the full sets of sentences for each target group to see an overview of the themes that are characteristic to each group. Using a word association metric, we extract words that tend to be used when talking about one group and not the others. Then, we compare the groups along six major topics (friends and family, finance, work, politics, technology, and health) since psychology studies suggest that age-

⁴<https://spacy.io/>

related stereotypes can vary across life domains (Kornadt and Rothermund, 2011). For this, we identify sentences belonging to each topic using anchored topic modeling. Then, since we are interested in detecting stereotypes and other widely-held opinions, we focus on the highest-density areas of the sentence embedding space for each topic. To identify these areas, we employ HDBSCAN, a hierarchical density-based clustering algorithm (Campello et al., 2013). We describe these techniques in more detail below.

4.1. Word Association

We calculate an association score of a word w with the target corpus C_{tar} (sentences mentioning the target group) as compared to the reference corpus C_{ref} (texts mentioning any of the other groups), using Pointwise Mutual Information (PMI):

$$Score(w) = PMI(w, C_{tar}) - PMI(w, C_{ref}) \quad (1)$$

where PMI is calculated as follows:

$$PMI(w, C_{tar}) = \log_2 \frac{freq(w, C_{tar}) * N(C)}{freq(w, C) * N(C_{tar})} \quad (2)$$

where $freq(w, C_{tar})$ is the number of times the word w occurs in the target corpus, $freq(w, C)$ is the number of times the word w occurs in the full corpus $C = C_{tar} \cup C_{ref}$, $N(C_{tar})$ is the total number of words in the target corpus, and $N(C)$ is the total number of words in the full corpus. $PMI(w, C_{ref})$ is calculated in a similar way. Thus, Equation 1 is simplified to:

$$Score(w) = \log_2 \frac{freq(w, C_{tar}) * N(C_{ref})}{freq(w, C_{ref}) * N(C_{tar})} \quad (3)$$

Words with high $Score(w)$ tend to appear more often in the target corpus than in the reference corpus. We consider words (alpha-numeric sequences) with $Score(w) \geq 0.6$ (i.e., words whose frequency in the target corpus is at least $1.5 \times \frac{N(C_{tar})}{N(C_{ref})}$ times higher than in the reference corpus) as representing the texts for the corresponding target group. We ignore stopwords and low-frequency words. The frequency threshold is set up to 0.5% of the number of sentences in the corresponding target corpus. The PMI method of word association has been successfully applied in a number of similar NLP applications (Rudinger et al., 2017; Clark et al., 2016; Kiritchenko et al., 2020). However, we note that other methods to estimate the degree of association of a word with a category (e.g., cross entropy, Chi-squared test, and information gain) can be used instead.

4.2. Anchored Topic Modeling

While most topic modelling techniques are unsupervised (Blei et al., 2003; Bianchi et al., 2021), supervised and semi-supervised topic modelling methods have been developed to incorporate domain knowledge when some labels are available (Huang et al., 2018; Yazdavar et al., 2017). In this work, we use the Correlation Explanation (CorEx) topic model, introduced

by Gallagher et al. (2017). CorEx separates the hidden topics through an information-theoretic framework and can incorporate various levels of domain knowledge as an unsupervised, semi-supervised, or hierarchical topic model. Gallagher et al. (2017) showed that CorEx is able to extract coherent topics from tweets, despite the general difficulty of topic modeling for short texts.

We specifically use the anchoring strategy of CorEx where topics are predefined with a set of anchoring words, to ensure comparable topics across the groups. We analyze our age-related data across six common topics expected to occur for all age groups: family and friends, finance, work, politics, technology, and health. The topics are seeded with a set of anchor words related to the themes we wish to explore:

- **Family:** *family, friends, parents, kids, children, relatives*
- **Finance:** *finance, money, savings, housing, house, homeowners, rich, poor, real estate, market, bitcoin, crypto, wealth, debt, student loan*
- **Work:** *work, job, work ethic, retire, retirement, goals, entrepreneur, boss, wage, office, career, hustle, workers, business, labor*
- **Politics:** *conservative, liberal, republican, democratic, election, vote, news, protest, government, legal, social justice, climate change, environment, diversity*
- **Technology:** *social media, internet, computers, facebook, instagram, tiktok, twitter, smartphone, phone, texting, screen*
- **Health:** *health, wellness, fitness, exercise, illness, hospital, death, disease, mental health, stress, depression, anxiety, diet, strength, covid*

CorEx has a parameter, referred to as anchor strength, that controls the mutual information between the anchor words and their respective topics. We set this parameter to 5 to enforce topics highly associated with the anchor words. We also let the topic model extract ten topics in total, leaving room for other topics that might naturally appear in our data (these additional four topics are discarded, and not included in the analysis).

Before performing the topic modeling, we pre-process the data as follows: we first lemmatize the words using the spaCy lemmatizer. We then use the gensim library to compute high-frequency bigrams that can be treated as single tokens in the topic modelling (e.g., *social media* and *real estate*). The text is then converted to a binary term-document matrix, removing stop words and words occurring less than 10 times.

4.3. Density-Based Clustering

We apply the HDBSCAN algorithm on the sentences associated with each topic, for each target group, separately. We first compute sentence embeddings using the

pretrained `all-mpnet-base-v2` model from SentenceTransformers.⁵ We chose this embedding model as it had the highest overall performance on the SentenceTransformers leaderboard at the time of development. The embedding model was trained on over 1 billion sentences from a variety of sources; it can process sentences containing up to 384 tokens (word pieces), and outputs 768-dimensional representations.

Following the recommendations in the HDBSCAN documentation,⁶ we set the ‘minimum samples’ parameter to the minimum value of 1, for the least conservative clusters possible (i.e., the fewest points will be discarded as noise). The epsilon parameter controls the distance threshold under which clusters may be merged; this is set to 0.1 so that sentence vectors which are close together will be merged into the same cluster. Finally, we set the minimum cluster size individually for each topic and target group. We expect that commonly-repeated opinions should result in large clusters; however, we observe that if this parameter is set too large, no clusters are returned. Following the guidance of Schubert et al. (2017) that the largest component should not contain more than 50% of the clustered points, we heuristically set the minimum cluster size to be as large as possible, such that the largest component does not violate that constraint. Clusters are then manually labelled with the most frequent opinion expressed by sentences included in the cluster. Clusters that are purely factual in nature (i.e., not representing over-generalizations) or irrelevant to the topic are removed.

5. Results

We apply the aforementioned methods to discover frequently appearing over-generalizations in our Twitter data concerning the four generations, and compare them with views on young and old people commonly expressed by Twitter users.

5.1. Comparing the Four Generations

We start by examining sentences collected for various generation groups and identify words that characterize each group, using Equation 3. Table 2 shows 30 words with the highest association with a target group corpus for each of the four generations. Various themes are discussed in relation to the four generations, and the most prevalent stereotypes immediately surface in this word analysis. For baby boomers, the topics include retirement, wealth, political and economic impacts of this generation on the US, and the current COVID-19 pandemic. Also, the words *facebook* and *fb* are prominent for boomers, reflecting the common stereotype that the platform is overrun with older generations, in particular

⁵https://www.sbert.net/docs/pretrained_models.html, accessed December 2021.

⁶https://hdbscan.readthedocs.io/en/latest/parameter_selection.html

Target group	Words with highest association with the target corpus
Boomers	<i>retiring, stream, retire, die, fb, facebook, selfish, dying, retirement, covid, mad, country, economy, dead, wealth, leave, news, pay, government, crypto, system, fucked, complain, war, rich, white, left, free, long, children</i>
Gen-X	<i>1965, forgotten, 1980, 80, raised, 80s, exist, knows, best, grew, fellow, learned, generational, 90s, music, parents, late, needs, early, 50, kind, save, middle, cool, hell, came, called, true, told, little</i>
Millennials	<i>geriatric, elder, avocado, toast, 30s, killing, industry, afford, 40, killed, student, lazy, buying, fellow, oldest, article, houses, homes, older, debt, home, woke, house, buy, eat, 30, housing, able, try, 5</i>
Gen-Z	<i>tik, alpha, jeans, tok, tiktok, thinks, fashion, sensitive, cancel, knows, needs, wants, bring, hope, cringe, 11, 9, grow, 90s, online, save, cool, different, honestly, fun, person, culture, makes, weird, word</i>

Table 2: Words associated with different generational groups. Top 30 words are shown for each group.

baby boomers,⁷ who often post absurd messages and believe everything they see on social media.⁸ For Gen-X, we see more factual information, describing it as a generation born between 1965 and 1980, and raised during the 80s and 90s. It is considered a forgotten and mostly ignored generation, lost between the two larger groups, boomers and millennials. This is also supported by the fact that Gen-X is mentioned less frequently in tweets, and the number of the collected sentences for Gen-X is 3–6 times smaller than the numbers of the sentences for the other generations.

For millennials, we note the common stereotypes portraying them as being lazy, blaming them for killing many industries from diamonds to McDonald’s, and accusing them of not being able to afford their own houses because they spend all their money on avocado toast.⁹ Gen-Z is portrayed as a generation that grew up with the Internet, who spend most of their time online,¹⁰ and who learn about important social and political issues on TikTok and other social media.¹¹

Next, we look at the results of the cluster analysis and compare the most prevalent stereotypes on Twitter for the four generations across six major topics. Figure 1 shows the largest coherent clusters that emerge for each topic, for each generation. In discussions on family relations, the older generations, boomers and Gen-X, are blamed for having been terrible parents, and the younger generations are anticipated to do better, if they have children of their own at all. Regarding finances, boomers are considered to control all the wealth while the other generations struggle in poverty and debt. At the same time, boomers are believed to have worked hard for their wealth whereas the younger generations

are often seen as lazy and not wanting to work. However, many users also think that it is time for boomers to retire and let younger people take their jobs. Boomers are believed to be actively participating in elections, and there are many calls for the younger generations to get involved and vote as well. Regarding political orientation, Gen-X and Gen-Z are seen as conservative or at least becoming more conservative while millennials keep liberal views. For boomers, we see contradictory statements claiming that all boomers are liberal and all boomers are conservative. As the oldest generation, boomers are portrayed as more resistant to change, clinging to old ways of doing things and having a difficult time adapting to new technologies (e.g., they don’t know how to use Internet or invest in cryptocurrency, and don’t believe in climate change). They are also concerned about their physical health, especially at the time of the global COVID-19 pandemic, and have no concept of “mental health”. The younger generations, on the other hand, are eager to adopt new ideas, from teleworking to bitcoin, quickly move from older social media platforms to emerging ones, and prioritize their health, including mental health, over work.

5.2. Comparing Young and Old

We again start by examining sentences collected for older and younger adult groups and identify words that tend to appear frequently for one group and less frequently for the other group. Table 3 shows the 30 words with the highest association for the two groups.

First, older adults are repeatedly discussed in regards of their vulnerability and high rates of death from various causes, and in particular from COVID-19. Older people living in nursing homes, along with disabled people and the immune-compromised, are considered at risk and needing protection. Other frequently occurring themes for older adults are their allegedly annoying and rude behavior, poor driving habits, and inability to handle new technologies, that either amuse or anger younger people. So, we mostly see common negative stereotypes of older people being fragile, inflexible, and bitter (Chasteen et al., 2002). Younger people, on the other hand, are described in a somewhat more positive light. They question authorities on po-

⁷<https://www.insider.com/facebook-gen-z-teens-boomer-social-network-leaks-2021-10>

⁸<https://www.theguardian.com/technology/2019/jul/08/millennials-baby-boomers-roleplaying>

⁹<https://www.chicagotribune.com/opinion/commentary/ct-perspec-millennials-killing-economy-avocado-toast-ram-pell-1210-20181207-story.html>

¹⁰<https://www.businessinsider.com/gen-z-struggles-always-connected-online-2019-7>

¹¹<https://itp.live/article/10316-how-tiktok-is-more-than-just-an-entertainment-app-for-gen-z>

Topic	Boomers	Gen-X	Millennials	Gen-Z
Family & Friends	<ul style="list-style-type: none"> - were terrible parents - invented the idea of "participation trophies" for their children 	<ul style="list-style-type: none"> - are terrible parents - were latch-key kids - tried to do better than their parents 	<ul style="list-style-type: none"> - are not having children - are not kids anymore, have kids of their own - are the first generation to be worse off than their parents - give their kids weird names 	<ul style="list-style-type: none"> - are going to be the best parents
Finances	<ul style="list-style-type: none"> - control all the wealth - don't understand crypto - will pass their wealth to their kids, if they don't spend it first 	<ul style="list-style-type: none"> - are poor and in debt - accumulated wealth in the housing boom - are more wealthy than millennials, but less than boomers - need housing too 	<ul style="list-style-type: none"> - can't afford houses - invest in crypto - don't know how to manage money - are the poorest generation ever 	<ul style="list-style-type: none"> - invest in crypto - have no money - will never be able to afford houses
Work	<ul style="list-style-type: none"> - don't want to see an increase in minimum wage - worked hard for their wealth - are retiring soon, or should be - difficult to work with 	<ul style="list-style-type: none"> - are at the point of retirement, but can't afford it - like working from home 	<ul style="list-style-type: none"> - will never be able to afford to retire - don't like to work - are known as being lazy, but just want a living wage - are working class 	<ul style="list-style-type: none"> - are entering the workforce and redefining work - are entitled, spoiled, and lazy - are not missing out by not being in the office - don't want to work
Politics	<ul style="list-style-type: none"> - don't take climate change seriously - are liberal - are conservative - should not be running the government - do vote 	<ul style="list-style-type: none"> - are becoming more conservative - don't know how to vote 	<ul style="list-style-type: none"> - are liberals - don't vote enough - are concerned about climate change 	<ul style="list-style-type: none"> - are conservative - need to get out and vote - really care about climate change
Technology	<ul style="list-style-type: none"> - use Facebook - shouldn't be allowed on Twitter - don't know how to use the Internet - talk on a phone all the time 	<ul style="list-style-type: none"> - were the last generation to grow up without the Internet - use Twitter - use Facebook - don't like making phone calls 	<ul style="list-style-type: none"> - never answer phone calls - are moving to TikTok - were growing up just as the Internet was invented - don't use Facebook anymore - use Twitter 	<ul style="list-style-type: none"> - are on TikTok all the time - ruined Twitter - are afraid to talk on the phone - don't know how the Internet used to be - grew up with social media - don't use Facebook
Health	<ul style="list-style-type: none"> - are scared of dying from COVID - are retiring due to COVID concerns - are getting old and dying - have no concept of "mental health" 		<ul style="list-style-type: none"> - all have anxiety and other mental health issues - prioritize health before work 	<ul style="list-style-type: none"> - are serious about mental health - have high rates of anxiety and depression

Figure 1: Most frequently mentioned over-generalizations for the four generation groups pertinent to six topics, as revealed by the cluster analysis. The shown opinions are sorted by frequency in descending order.

litical, social, and environmental issues, and are seen as future leaders that can bring positive change. Yet, they are inexperienced and need support to get access to education and develop their skills.

We notice similar themes when examine the clustering results for the two groups across the six main topics (Figure 2). Older people are portrayed as more resistant and less able to adopt new technologies, such as smart phones, social media, and cryptocurrency. Younger people, even though more technology-savvy, are inexperienced in managing their finances. While considered passive in voting, young people are believed to be more socially and politically aware, raising concerns about climate change and protesting government dictatorships. Young people are believed to be mostly liberal while older people tend to be conservative. There are more concerns about physical health for older adults and more discussions around mental

health for young people. Many of these views replicate the common over-generalizations found for the older or younger generations, respectively (Sec. 5.1). In addition to these, we see other views common in the two datasets. For example, many users believe that older people, and in particular the boomer generation, need to retire now as young people need jobs. We also see lots of contention between parents (older people) and their children (young people). The noticeable exception to these similarities is the assumed wealth of the boomer generation while old people are perceived as poor, barely meeting their needs.

6. Discussion

Our analysis shows that we can extract many of the prevalent age-related stereotypes directly from social media. For example, much of the psychological literature on generational stereotypes has focused on how

Target group	Words with highest association with the target corpus
Older adults	<i>people, nursing, disabled, vulnerable, lmao, funny, facebook, phone, ass, homes, eat, mad, die, walk, flu, fucking, damn, younger, fuck, lol, weird, immune, shit, wanna, death, drive, kids, died, hate, sick</i>
Young people	<i>opportunities, skills, climate, opportunity, role, involved, leaders, interested, education, action, amazing, create, learning, future, mental, support, ensure, build, lead, schools, great, learn, important, leaving, today, hope, political, change, healthy, jobs</i>

Table 3: Words associated with the *older adults* and *young people* target groups. Top 30 words are shown for each group.

Topic	Older adults	Young people
Family & Friends	<ul style="list-style-type: none"> - are similar to children - love children - hate children 	<ul style="list-style-type: none"> - can't afford to start a family - have a difficult relationship with their parents
Finances	<ul style="list-style-type: none"> - are forced to sell their homes to pay for their care - are not all rich - don't understand crypto - can't afford to heat their homes - are seen as a financial burden on society 	<ul style="list-style-type: none"> - don't know how to manage their money - invest in crypto - can't afford houses
Work	<ul style="list-style-type: none"> - need to retire now - it's sad to see them having to work - have worked and paid taxes their whole lives 	<ul style="list-style-type: none"> - need jobs - don't want to work for low wages - it's great to see them thrive in their careers
Politics	<ul style="list-style-type: none"> - do vote - are right-wing / conservative - like to watch the news all day - should not be running the government 	<ul style="list-style-type: none"> - are left-wing / liberal - are concerned about climate change - don't vote - protest government dictatorships
Technology	<ul style="list-style-type: none"> - use Facebook - should get off Twitter - always shout on the phone - don't know how to use smart phones - write text messages like writing a letter - are taking over TikTok 	<ul style="list-style-type: none"> - are addicted to social media - don't use Facebook - are always on their phones - have moved to TikTok
Health	<ul style="list-style-type: none"> - are at most risk of dying from COVID - are taking up all the beds in hospital 	<ul style="list-style-type: none"> - are not at risk from COVID - care about mental health - are at higher risk from vaccine than COVID

Figure 2: Most frequently mentioned over-generalizations for the *older adults* and *young people* target groups pertinent to six topics, as revealed by the cluster analysis. The shown opinions are sorted by frequency in descending order.

workers from different birth cohorts are perceived in the workplace. Many of the stereotypes reported in these studies are mirrored in our analysis, such as that workers from the baby boomer generation are seen as hardworking, resistant to change, and not technology-savvy (Perry et al., 2013; Weeks et al., 2017; Van Rossem, 2019b), Gen-X are seen as preferring work flexibility and being more comfortable with technology than boomers (Perry et al., 2013; Weeks et al., 2017), and millennials are seen as entitled and lazy, valuing the monetary rewards of their job, and emphasizing work-life balance (Perry et al., 2013; Weeks et al., 2017; Van Rossem, 2019b).

However, our results also demonstrate that stereotypes depend on context: boomers are seen as powerful in terms of controlling wealth and dominating the housing market, but weak when it comes to physical health and ability to use technology. This is consistent with the findings of Kornadt and Rothermund (2011), who showed that the positivity/negativity of age-related

stereotypes varied across life domains. For example, in their study, older people were rated more positively in the domain of family and partnership than in the domain of mental and physical fitness.

Our results also confirm the hypothesis that in some cases, different stereotypes are associated with the four generations than with old and young people more generally. While some of the stereotypes in Section 5.1 and Section 5.2 are similar, we observe striking differences in how boomers are portrayed as wealthy and holding onto resources that should go to younger people, while older people are seen as frail, in danger of losing their homes, and the object of pity when seen working. One reasonable explanation is that users have different age cohorts in mind when they speak of *boomers* (by definition aged 58–76 years old) and *old folks*, *elderly*, or *senior citizens*. In fact, many of these terms have been criticized as connoting vulnerability and dependence, while also being ambiguous with respect to actual age (Berridge and Hooyman, 2020). Furthermore,

the age at which someone is considered “young” or “old” is also context-dependent. Kornadt and Rothermund (2011) found that the age of 60 was considered “old” in the work domain, whereas 70 was considered “old” in the family domain. Thus in our data, when users talk about old people “taking over TikTok” they are presumably thinking about a different age group than those old people entering long-term care. These fluid definitions of *old* and *young* lead to a complex and indirect mapping between age and generation.

Many of the stereotypes observed are quite negative, with a relative absence of typical “positive” stereotypes for older people as being wise, caring, and knowledgeable. This may be partly due to the nature of Twitter, where negative messages are more likely to gain traction than positive or neutral ones (Tsugawa and Ohsaki, 2015). Nonetheless, such negative views can be psychologically damaging. In the case of stereotypes of aging, negative stereotypes held in younger years, when “old people” are seen as the out-group, can become internalized and incorporated into one’s self-image when an individual reaches old age themselves (Kotter-Grühn and Hess, 2012).

We also see evidence of prescriptive stereotypes; i.e., beliefs about how members of a group *should* behave. These prescriptive ideas appear most frequently with respect to the succession of employment, power, and wealth, limitation of consumption of public resources, and avoidance of identity transgression (older people acting in ways traditionally associated with younger people) (North and Fiske, 2013). Such prescriptive stereotypes are dangerous, as they can lead to anger and resentment when they are perceived to be violated.

7. Limitations

This represents preliminary work, and we acknowledge several limitations. We have no access to any demographic information about the users, although it is known that stereotypical beliefs tend to vary with age, gender, culture, and so on. Additionally, we do not know the geographical location of the users whose data were included in the analysis. Our data collection is limited to three months, and only to the English language. While our results align well with the reported stereotypes in the academic literature and popular media, these observations are qualitative. More research, particularly in generating large labelled datasets, will be needed to quantitatively evaluate our approach against other (e.g., neural) models. The NLP tools used in the analysis (e.g., dependency parser, sentence embeddings) were not specifically developed for social media text and may not perform at the same level of accuracy on tweet data. Furthermore, despite our best efforts at filtering, some of the data may not represent over-generalizations about the target groups. We hope to address some of these limitations in future work.

8. Conclusion

While ageism is wide-spread in North American society and can have dramatic impact on the health and well-being of its targets, it has received little attention in NLP. To address this gap, we extracted and analyzed common over-generalizations expressed on social media about four generations and two generic age categories. We were able to uncover many of the well-known stereotypes as well as some novel views, and confirm trends and observations reported in the psychological literature. In future work, we hope to expand the analysis to explore different social groups and consider longitudinal analysis to assess how stereotypes change with time. Ultimately, our research goals involve a better understanding of stereotypes and how they are expressed in language, so that we can develop tools for education and bias mitigation.

9. Bibliographical References

- Bartl, M., Nissim, M., and Gatt, A. (2020). Unmasking contextual stereotypes: Measuring and mitigating BERT’s gender bias. In *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*, pages 1–16, December.
- Becton, J. B., Walker, H. J., and Jones-Farmer, A. (2014). Generational differences in workplace behavior. *Journal of Applied Social Psychology*, 44(3):175–189.
- Berridge, C. and Hooyman, N. (2020). The consequences of ageist language are upon us. *Journal of Gerontological Social Work*, 63(6-7):508–512.
- Bianchi, F., Terragni, S., and Hovy, D. (2021). Pre-training is a hot topic: Contextualized document embeddings improve topic coherence. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 759–766, Online, August. Association for Computational Linguistics.
- Binns, R., Veale, M., Van Kleek, M., and Shadbolt, N. (2017). Like trainer, like bot? Inheritance of bias in algorithmic content moderation. In *Proceedings of the International Conference on Social Informatics*, pages 405–415. Springer.
- Blaine, B. E. and Brechley, K. J. M., (2017). *Understanding the psychology of diversity*, chapter 9. SAGE publications.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems*, pages 4349–4357.
- Breitfeller, L., Ahn, E., Jurgens, D., and Tsvetkov, Y. (2019). Finding microaggressions in the wild: A

- case for locating elusive phenomena in social media posts. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1664–1674, Hong Kong, China, November. Association for Computational Linguistics.
- Brewer, M. B. and Lui, L. N. (1989). The primacy of age and sex in the structure of person categories. *Social Cognition*, 7(3):262–274.
- Caliskan, A., Bryson, J. J., and Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Campello, R. J., Moulavi, D., and Sander, J. (2013). Density-based clustering based on hierarchical density estimates. In *Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 160–172. Springer.
- Casad, B. J. and Bryant, W. J. (2016). Addressing stereotype threat is critical to diversity and inclusion in organizational psychology. *Frontiers in Psychology*, 7:8.
- Chan, S. C., Au, A. M., and Lai, S. M. (2020). The detrimental impacts of negative age stereotypes on the episodic memory of older adults: does social participation moderate the effects? *BMC Geriatrics*, 20(1):1–12.
- Charlesworth, T. E., Yang, V., Mann, T. C., Kurdi, B., and Banaji, M. R. (2021). Gender stereotypes in natural language: Word embeddings show robust consistency across child and adult language corpora of more than 65 million words. *Psychological Science*, 32(2):218–240.
- Chasteen, A. L., Schwarz, N., and Park, D. C. (2002). The activation of aging stereotypes in younger and older adults. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 57(6):P540–P547.
- Chung, Y.-L., Kuzmenko, E., Tekiroglu, S. S., and Guerini, M. (2019). CONAN - COUNTER NARRATIVES THROUGH NICHE-SOURCING: A MULTILINGUAL DATASET OF RESPONSES TO FIGHT ONLINE HATE SPEECH. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2819–2829, Florence, Italy, July. Association for Computational Linguistics.
- Clark, P., Etzioni, O., Khot, T., Sabharwal, A., Tafjord, O., Turney, P., and Khashabi, D. (2016). Combining retrieval, statistics, and inference to answer elementary science questions. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Costanza, D. P., Badger, J. M., Fraser, R. L., Severt, J. B., and Gade, P. A. (2012). Generational differences in work-related attitudes: A meta-analysis. *Journal of Business and Psychology*, 27(4):375–394.
- Díaz, M., Johnson, I., Lazar, A., Piper, A. M., and Gergle, D. (2018). Addressing age-related bias in sentiment analysis. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–14.
- Fiske, S. T. (2018). Stereotype content: Warmth and competence endure. *Current directions in psychological science*, 27(2):67–73.
- Fokkens, A., Ruigrok, N., Beukeboom, C., Sarah, G., and Van Atteveldt, W. (2018). Studying Muslim stereotyping through microportrait extraction. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Fraser, K. C., Nejadgholi, I., and Kiritchenko, S. (2021). Understanding and countering stereotypes: A computational approach to the stereotype content model. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 600–616, Online, August. Association for Computational Linguistics.
- Fraser, K. C., Kiritchenko, S., and Nejadgholi, I. (2022). Computational modeling of stereotype content in text. *Frontiers in Artificial Intelligence*, 5.
- Gallagher, R. J., Reing, K., Kale, D., and Ver Steeg, G. (2017). Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics*, 5:529–542.
- Garg, N., Schiebinger, L., Jurafsky, D., and Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.
- Huang, M., Rao, Y., Liu, Y., Xie, H., and Wang, F. L. (2018). Siamese network-based supervised topic modeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4652–4662.
- Hummert, M. L., Garstka, T. A., Shaner, J. L., and Strahm, S. (1995). Judgments about stereotypes of the elderly: Attitudes, age associations, and typicality ratings of young, middle-aged, and elderly adults. *Research on aging*, 17(2):168–189.
- Hutchinson, B., Prabhakaran, V., Denton, E., Webster, K., Zhong, Y., and Denuyl, S. (2020). Social biases in NLP models as barriers for persons with disabilities. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5491–5501, Online, July. Association for Computational Linguistics.
- Joshi, A., Dencker, J. C., Franz, G., and Martocchio, J. J. (2010). Unpacking generational identities in organizations. *Academy of management review*, 35(3):392–414.
- Kiritchenko, S. and Mohammad, S. (2018). Examining gender and race bias in two hundred sentiment analysis systems. In *Proceedings of the Sev-*

- enth Joint Conference on Lexical and Computational Semantics, pages 43–53, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Kiritchenko, S., Hipson, W., Coplan, R., and Mohammad, S. M. (2020). SOLO: A corpus of tweets for examining the state of being alone. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1567–1577, Marseille, France, May. European Language Resources Association.
- Kirk, H. R., Volpin, F., Iqbal, H., Benussi, E., Dreyer, F., Shtedritski, A., Asano, Y., et al. (2021). Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models. *Advances in Neural Information Processing Systems*, 34.
- Kornadt, A. E. and Rothermund, K. (2011). Contexts of aging: Assessing evaluative age stereotypes in different life domains. *Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 66(5):547–556.
- Kotter-Grühn, D. and Hess, T. M. (2012). The impact of age stereotypes on self-perceptions of aging across the adult lifespan. *Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 67(5):563–571.
- Kurita, K., Vyas, N., Pareek, A., Black, A. W., and Tsvetkov, Y. (2019). Measuring bias in contextualized word representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172.
- Liang, P. P., Wu, C., Morency, L.-P., and Salakhutdinov, R. (2021). Towards understanding and mitigating social biases in language models. In *Proceedings of the International Conference on Machine Learning*, pages 6565–6576.
- Lichtenstein, B. (2021). From “coffin dodger” to “boomer remover”: Outbreaks of ageism in three countries with divergent approaches to coronavirus control. *The Journals of Gerontology: Series B*, 76(4):e206–e212.
- Lyons, S. and Kuron, L. (2014). Generational differences in the workplace: A review of the evidence and directions for future research. *Journal of Organizational Behavior*, 35(S1):S139–S157.
- Macky, K., Gardner, D., and Forsyth, S. (2008). Generational differences at work: Introduction and overview. *Journal of managerial psychology*.
- Mannheim, K. (1952). The problem of a sociology of knowledge. *Essays on the Sociology of Knowledge*, pages 134–190.
- Manzini, T., Chong, L. Y., Black, A. W., and Tsvetkov, Y. (2019). Black is to criminal as caucasian is to police: Detecting and removing multiclass bias in word embeddings. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 615–621.
- Marzouki, Y., Barach, E., Srinivasan, V., Shaikh, S., and Feldman, L. B. (2020). The dynamics of negative stereotypes as revealed by tweeting behavior in the aftermath of the Charlie Hebdo terrorist attack. *Heliyon*, 6(8):e04311.
- Mathew, B., Kumar, N., Goyal, P., Mukherjee, A., et al. (2018). Analyzing the hate and counter speech accounts on twitter. *arXiv preprint arXiv:1812.02712*.
- May, C., Wang, A., Bordia, S., Bowman, S., and Rudinger, R. (2019). On measuring social biases in sentence encoders. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 622–628.
- Meshi, D., Tamir, D. I., and Heekeren, H. R. (2015). The emerging neuroscience of social media. *Trends in cognitive sciences*, 19(12):771–782.
- Nicolas, G., Bai, X., and Fiske, S. T. (2021). Comprehensive stereotype content dictionaries using a semi-automated method. *European Journal of Social Psychology*, 51(1):178–196.
- North, M. S. and Fiske, S. T. (2013). Act your (old) age: Prescriptive, ageist biases over succession, consumption, and identity. *Personality and Social Psychology Bulletin*, 39(6):720–734.
- Perry, E. L., Hanvongse, A., and Casoinic, D. (2013). Making a case for the existence of generational stereotypes: A literature review and exploratory study. *The SAGE handbook of aging, work and society*, pages 416–442.
- Prates, M. O., Avelar, P. H., and Lamb, L. C. (2020). Assessing gender bias in machine translation: a case study with Google Translate. *Neural Computing and Applications*, 32(10):6363–6381.
- Price, I., Gifford-Moore, J., Flemming, J., Musker, S., Roichman, M., Sylvain, G., Thain, N., Dixon, L., and Sorensen, J. (2020). Six attributes of unhealthy conversations. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 114–124, Online, November. Association for Computational Linguistics.
- Rao, P. and Taboada, M. (2021). Gender bias in the news: A scalable topic modelling and visualization framework. *Frontiers in Artificial Intelligence*, 4.
- Rudinger, R., May, C., and Van Durme, B. (2017). Social bias in elicited natural language inferences. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 74–79.
- Rudinger, R., Naradowsky, J., Leonard, B., and Van Durme, B. (2018). Gender bias in coreference resolution. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 8–14, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Sap, M., Gabriel, S., Qin, L., Jurafsky, D., Smith, N. A., and Choi, Y. (2020). Social bias frames: Reasoning about social and power implications of lan-

- guage. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Schmidt, A. and Wiegand, M. (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.
- Schubert, E., Sander, J., Ester, M., Kriegel, H. P., and Xu, X. (2017). DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems (TODS)*, 42(3):1–21.
- Spencer, S. J., Logel, C., and Davies, P. G. (2016). Stereotype threat. *Annual Review of Psychology*, 67:415–437.
- Stanovsky, G., Smith, N. A., and Zettlemoyer, L. (2019). Evaluating gender bias in machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy, July. Association for Computational Linguistics.
- Tajfel, H. and Turner, J. C. (1979). An integrative theory of intergroup conflict. In William G. Austin et al., editors, *The Social Psychology of Intergroup Relations*, chapter 3, pages 33–47. Brooks/Cole Pub. Co.
- Thelwall, M. (2018). Gender bias in sentiment analysis. *Online Information Review*.
- Toomey, E. C. and Rudolph, C. W. (2017). Age stereotypes in the workplace. *Encyclopedia of Geropsychology*. Singapore: Springer, pages 89–95.
- Tsugawa, S. and Ohsaki, H. (2015). Negative messages spread rapidly and widely on social media. In *Proceedings of the 2015 ACM Conference on Online Social Networks*, pages 151–160.
- Van Rossem, A. (2019a). Generational meta-stereotypes: Generational stereotypes in action. In *Academy of Management Proceedings*, page 17490. Academy of Management Briarcliff Manor, NY 10510.
- Van Rossem, A. H. (2019b). Generations as social categories: An exploratory cognitive study of generational identity and generational stereotypes in a multigenerational workforce. *Journal of Organizational Behavior*, 40(4):434–455.
- Van Rossem, A. H. D. (2021). Nurses’ generational stereotypes and self-stereotypes: a cognitive study. *Journal of Health Organization and Management*.
- Venkit, P. N. and Wilson, S. (2021). Identification of bias against people with disabilities in sentiment analysis and toxicity detection models. *arXiv preprint arXiv:2111.13259*.
- Vidgen, B., Harris, A., Nguyen, D., Tromble, R., Hale, S., and Margetts, H. (2019). Challenges and frontiers in abusive content detection. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy.
- Wang, Z. and Potts, C. (2019). TalkDown: A corpus for condescension detection in context. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3711–3719, Hong Kong, China, November. Association for Computational Linguistics.
- Weeks, K. P., Weeks, M., and Long, N. (2017). Generational perceptions at work: In-group favoritism and out-group stereotypes. *Equality, Diversity and Inclusion: An International Journal*.
- Wiegand, M., Ruppenhofer, J., and Eder, E. (2021). Implicitly abusive language—what does it actually look like and why are we not getting there? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 576–587.
- Wong, M., Gardiner, E., Lang, W., and Coulon, L. (2008). Generational differences in personality and motivation: do they exist and what are the implications for the workplace? *Journal of Managerial Psychology*.
- Wyman, M. F., Shiovitz-Ezra, S., and Bengel, J. (2018). Ageism in the health care system: Providers, patients, and systems. In *Contemporary Perspectives on Ageism*, pages 193–212. Springer, Cham.
- Yazdavar, A. H., Al-Olimat, H. S., Ebrahimi, M., Bajaj, G., Banerjee, T., Thirunarayan, K., Pathak, J., and Sheth, A. (2017). Semi-supervised approach to monitoring depressive symptoms in social media. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pages 1191–1198.
- Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K.-W. (2018). Gender bias in coreference resolution: Evaluation and debiasing methods. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20, New Orleans, Louisiana, June. Association for Computational Linguistics.