

Defining and Detecting Toxicity on Social Media: Context and Knowledge are Key

Amit Sheth¹, Valerie L. Shalin^{2,1}, Ugur Kursuncu^{1,3}

¹ *AI Institute, University of South Carolina*

² *Department of Psychology, Wright State University*

³ *J. Mack Robinson College of Business, Georgia State University*

Abstract

Online platforms have become an increasingly prominent means of communication. Despite the obvious benefits to the expanded distribution of content, the last decade has resulted in disturbing toxic communication, such as cyberbullying and harassment. Nevertheless, detecting online toxicity is challenging due to its multi-dimensional, context sensitive nature. As exposure to online toxicity can have serious social consequences, reliable models and algorithms are required for detecting and analyzing such communication across the vast and growing space of social media. In this paper, we draw on psychological and social theory to define toxicity. Then, we provide an approach that identifies multiple dimensions of toxicity and incorporates explicit knowledge in a statistical learning algorithm to resolve ambiguity across such dimensions.

Keywords: toxicity, cursing, harassment, extremism, radicalization, context

1. Introduction

Online social media platforms are arguably among the most culturally significant technological innovations of the 21st century. The numerous benefits include the wide distribution of content crossing geographic boundaries, and enabling interaction and exchanges that are nearly free of physical constraints except for infrastructure. Communities have emerged around every conceivable special interest from science to travel, from politics to child-rearing. The easy

spread of data, information, and knowledge was expected to foster informed decision-making, cultural exchanges, and the coordination of activities online and in the physical world. Unfortunately, social media has also significantly enhanced the reach and scale of harmful content including disinformation, conspiracies, extremism, harassment, violence, and other forms of socially toxic material. While social media platforms attempt to counter such harmful content, their efforts are largely ineffective and as such themselves have the potential for unintended adverse impact. The effectiveness of moderation is potentially biased by the platforms’ economic interest or political and regulatory considerations. Or failure may simply be due to the lack of effective tools and sufficient investment. Irrespective of the reason, human content moderation has resulted in relatively unsatisfactory outcomes [1]. Although the political and public health climate of 2020 encouraged society to adopt technological and specifically AI-based solutions, success was also limited. A prominent reason is the lack of understanding of the challenging nature of toxicity, which fundamentally requires context outside of the explicit content. The detection of toxicity demands an interdisciplinary perspective with empirical approaches. Consistent with our people-content-network framework for the characterization of social media exchange [2, 3, 4, 5], we assert the more general role of context, and in particular cultural context, in the interpretation of content. This paper has three goals:

- identify the psychological and social dimensions of the problem
- identify the limitations of contemporary computational approaches, and
- outline an advanced technical approach founded on knowledge-driven context-based analysis.

2. A PsychoSocial problem Meets Computation

Our view of toxic content extends beyond the current classifications that focus on “threats, obscenity, insults, and identity-based hate”¹. We also include harassment and socially disruptive persuasion, such as misinformation, radicalization and gender-based violence [6, 7, 8, 9, 10]. While the cultural foundations of toxicity are readily apparent in misinformation and radicalization, we contend that culture provides essential context to the determination of any toxic content. Figure 1 guides this section, starting with conventional content analysis and expanding the psychological, social, and cultural scope of the required analysis.

2.1. Content Analysis

By far the most common approach to toxicity detection focuses on the content of exchanges. Offensive keywords, often so-called “coarse language” are easy to tabulate in a lexicon. More sophisticated analyses employ lexicons specific to intelligence, appearance, race, sexual preference, etc. [8]. Keyword-based content analysis encounters a number of challenges. An evolving culture conveys an insulting connotation to otherwise apparently banal language, e.g., basic, cancel, Karen, shade, snowflake, and thirsty. This not only requires constant maintenance of the lexicon, but




Figure 1: Conventional toxicity analysis examines the content exchanged between individuals in a community (1). Often external observers impose their own culturally biased decision rules (e.g., gender-based violence) (2). Detecting toxic sources (3) expands analysis, but still fails to acknowledge the reaction of the target (4), which is likely tempered by common group membership (5).

¹<https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge/overview>

context to disambiguate the slang usage from general usage [11].

A second problem is that content analysis based on isolated lexical items does not necessarily confer toxicity [10, 9, 12]. For example, North American teenagers readily employ language among themselves that adults would consider offensive. More worrisome, the word “jihad” may be readily interpreted as a radical content by a Westerner, but a more culturally sensitive analysis reveals that this term also appears in benign religious text. The scope of toxic topics, general knowledge, and cultural foundations required for interpretation is virtually unbounded. “Dressing like your grandmother” directed to a teen is laden with cultural deprecation founded on both ageism and consumerism but contains no single offensive word in isolation. Irony, humor, and teasing between friends precludes simple sentiment analysis.

Toxic content is also multimodal, often exploiting images and videos that are not well monitored or understandable using contemporary technology. Facial recognition algorithms fail miserably for dark-skinned females. More generally, the benchmark image databases heavily favor western culture, at the expense of eastern cultures (China is represented by only 1% of images in the Imagenet²). Apart from these obvious multi-modal processing challenges, text and image content must be aligned in a common framework at the appropriate level of abstraction. Finally, we cannot assume that text provides useful image processing guidance. A cheery “Have a Nice Day” text can easily pair with an embarrassing photo of the recipient.

2.2. Culturally Bound Decision Criteria

Even a simple content analysis requires a decision criterion. Framing toxicity as a standard signal detection problem acknowledges the potential for two overlapping distributions of potentially toxic content instances, one over relatively low toxicity values and another over higher values. The decision criterion is

²<https://devopedia.org/imagenet>

vulnerable to cultural considerations. All-too-common annotator disagreement is resolved by the vote of a small sample of annotators, while the foundations of disagreement remain unstudied. Annotators may hold unconscious stereotypes, for example associating religion with radicalization. Personal experience and cultural differences create variable interpretations of label semantics; disagreement over interpretation is the source of great concern in moderating internet content.

We have already noted the preponderance of coarse language in teenagers [13, 14, 15]. The population of truly non-toxic instances is much larger than the population of toxic instances [9, 16, 17]. Such class imbalance does not work well with contemporary machine learning algorithms. As we lower the toxicity decision rule we admit more false positives, potentially resulting in orders of magnitude more false positives than true positives. This creates both adverse impacts to the falsely accused along with a practical problem of follow up [7].

Finally, corpus assembly itself conveys cultural bias. The classification algorithm for one population does not generalize to different populations, creating a validity problem. As implied above, the relevant data set for detecting toxic exchange among adolescents differs from the data set for defining such exchanges among adults [9] due to different cultural practices.

2.3. Identifying the toxic source

The source’s intent to hurt or harm is the defining feature of bullying. Harm may employ the disclosure of sensitive facts, denigrate, be grossly offensive, be indecent or obscene, be threatening, make false allegations, deceive, spam, spread misinformation, mimic interest, clone profile or personal invade space. The motivation for detecting the toxic source is mitigation, but the false alarm risk is real. Moreover, one instance is unlikely to constitute sufficient evidence. Evaluation of the potentially toxic source requires a corpus of the candidate’s content, raises a challenge of corpus scope, and introduces the need for aggregation of evidence.

2.4. Identifying the Target

Experienced harm is distinct from intent to harm, from the recipient's perspective resulting in discrimination, deception, fraud, disinformation, loss of money, offense, loss of reputation, manipulation, embarrassment, distraction, loss of time. Moreover, harassment, by definition, refers to the special case of bullying with respect to a protected class [18]. The nefarious source takes advantage of the specific features of target vulnerability such as age, occupation, and public stature. Contemporary victims of bullying include Parkland High School Students ³. The motivation for detecting the target is protection. Both bullying and harassment are associated with cases of adolescent suicide [19, 20, 21]. Age most certainly matters in the assessment of target experience; adolescent brains are still developing an ability to process social feedback making adolescents particularly vulnerable to both negative feedback [22], and radicalization efforts [23].

2.5. Participants' Relationship and Group Membership

Friendship, power differentials, and social network membership provide essential context. Insults are common among adolescent friends. Participants with the same racial background readily exchange otherwise offensive racial epithets. Social network structure has at least two consequences founded on the distinction between in-group and out-group membership and the target's position within these groups. First, multiple negative messages from different participants in the in-group targeted to an out-group recipient are as potentially toxic as the same number of messages from a single source. Social network membership is therefore an important feature in the detection of toxicity. Second, the promise of group membership and the threat of exclusion is a known factor in the radicalization effort [24], of particular appeal to adolescent recruits. Thus, network-driven dynamics play an important role in the detection of toxicity, es-

³https://www.huffpost.com/entry/parkland-students-combat-cyberbullying_n_5aeeee86e4b033e5c3f03126

pecially when toxicity is amplified by networked efforts such as hate campaigns and organized misinformation spread.

Because surrounding benign conversation mitigates the single potentially toxic comment, exchange history informs the determination of toxicity. Hence, an exchange history surrounding the potentially toxic comment must be present in the corpus to enable evaluation. Crawling on sender and recipient identifiers is too limited. Victims are often targeted with mention tags in an exchange between a sender and what might be charitably called bystanders. These concerns illustrate that the scope of a purportedly toxic item influences the annotation task. A single episode may look quite different in the context of other exchanges, suggesting that the potentially benign or toxic instance should be annotated with respect to its broader historical and network context. This argues for new requirements for systematically defining and scoping the annotation task. Expanded context also raises the problem of conflicting indicators, the assessment of stale content, and the need for confidence estimates.

3. Technical Challenges to Automated Detection of Toxic Language

As described above, the toxicity detection problem is not a purely computer science or AI problem. To identify toxicity, it is necessary to understand the broader context beyond the situation and domain-specific content analysis, with reference to applicable human values, social norms, and culture, at the individual, group, and community levels [25]. Toxicity detection is an interdisciplinary problem founded on theory, empirical models, and knowledge to guide classification [26, 9, 12, 27, 28]. In contrast, conventional approaches for the identification of toxic exchange have been treated as a content processing problem [29, 30]. Researchers often rely on the post-level for building datasets and designing algorithms to detect toxicity between two individuals relying on the explicit language of insult [9]. The state-of-the-art algorithms used to model toxic content are mostly autoregressive models (e.g., BERT, GPT-2,3), designed to predict the next token given previous tokens from the dataset as input. As these models have been trained using data collected from the web, corpus bias

and incidentally confounded features result in models that can cause intentional or unintentional harm⁴ [31, 32, 12]. For example, recent studies [33, 34, 35] suggested that these state-of-the-art algorithms are prone to generating racist or sexist schemes. While these models can be retrained using transfer learning by fine-tuning model parameters, significant harmful bias will still carry over. Such models can be dangerous in highly consequential areas, such as online toxicity as well as health [36, 37, 38]. For instance, Google’s Perspective API⁵ designed for toxicity detection received criticism for biased scoring of content based on gender, sexual orientation, religion, or disability. Further, this model almost always assigned a high toxicity score if the content included insults or profanity, regardless of the intent or tone of the author [39]. Hence, policymakers and practitioners justifiably assert serious usability and safety concerns that constrain the adoption of poorly understood technologies [40].

Below, we discuss the technical consequences of our expanded approach to toxicity detection in three subsections: the need for empirical models, the need for a curated corpus, and the need for external knowledge.

3.1. Need for Empirical Models

Computational modeling of human behavior requires domain expertise to inform the classes and subclasses of toxicity. On the other hand, such domain expertise is scarce; hence, we require a conceptual model in a structured or semi-structured format that is readable by machines. The presently available models are impoverished. Below we consider the use cases for cursing, extremism, and harassment to demonstrate the need for better empirical models to guide content analysis.

Cursing. The intention of the parties of a conversation along with social context determines the meaning of their language. In [10], we studied the communica-

⁴<https://thenextweb.com/neural/2021/01/19/gpt-3-has-consistent-and-creative-anti-muslim-bias-study-finds/>

⁵<https://www.perspectiveapi.com/>

tions on Twitter concerning the use of cursing and its relations with intention and emotions. While we found around 8% of conversations contain profanity and curse words, the intention of users may not necessarily be toxic. We explored the role of emotions in identifying intention, as cursing may be associated with positive emotions as well as negative emotions, and these emotions may indicate the real intention. We identified three variables outside of the text that determines *when*, *where* and *how* cursing occurs. For example, we found that people curse more when they wake up, in relaxed virtual environments. Clearly, such features of social context should moderate the interpretation of cursing as a toxic indicator.

Extremism. In [7, 41, 42] we started with the political science notion that radicalization is a process employed by extremist groups, with systematic changes in persuasive content over time. For the particular use case of Islamist extremism, appropriate domain expertise is critical to distinguish the true extremist from non-extremist communication. Guided by an empirical model developed by a political scientist, we examined three dimensions to model this content: religion, ideology, and hate. Ambiguity is a significant challenge. For instance, the meaning of the keyword “jihad” in religion is referred to as a self-spiritual struggle, while it indicates intent to harm other individuals in the Islamist extremist ideology. As the same term has two different meanings for extremist and non-extremist content, it needs to be represented differently in a computational model for resolving such ambiguity. Hence, a multi-dimensional and contextually sensitive model of this content incorporating knowledge (in this example, religious knowledge) allows us to address ambiguity, reducing false alarms and mitigating unfairness. Such a socially responsible model with improved fairness would mitigate the adverse impacts on nearly 2 billion Muslims.

Harassment. Many early researchers defined harassment as a binary classification - a social media post (e.g., a tweet) is either harassing or not [43, 44, 45, 46]. As the context is crucial in capturing harassment, content classification will change based on the linguistic meaning, interpretation, and distribution. In [8]

we expanded the dimensions of harassment to include; (i) sexual, (ii) racial, (iii) appearance-related, (iv) intellectual, and (v) political content, and created a type-aware lexicon and annotated dataset [47]. Then we employed a multi-class classification algorithm based on these five dimensions. While coarse lexical items signal some of these, ambiguous common language (fat, dumb) and idioms are also relevant. A multi-class approach is required because perpetrators can exploit more than one subclass in targeting a victim. Critically, in the absence of a multi-class model, the victim’s experience of harassment over time will not surface.

3.2. Need for Curated Corpora

The analysis is only as good as the corpus. As previously noted, researchers often resort to the post-level for building datasets and designing algorithms to detect toxicity between two individuals, focusing on recognizing the explicit language of insult [9]. Keyword-based crawls in corpus assembly create misleading corpora, rife with ambiguity, e.g., a playful exchange between good friends with sarcastic content could be falsely flagged as harassment, or a religious reference to “jihad” could falsely flag a pious worshipper. We next consider the corpus assembly problem for extremism and harassment.

Extremism. For our extremism project, we relied upon a curated corpus [48] consisting of 538 verified extremist users, established by Twitter and the Lucky Troll club [49]. We balanced this with a corpus of 538 non-extremist users from an annotated Muslim religious dataset [50]. We make two points with this example. First, the set of positive cases reflected professional judgment. Second, the applicability of the resulting classification model depends on the quality of the distractor corpus. Here we were particularly concerned with adverse impact and therefore employed a distractor corpus that posed a significant false alarm opportunity. Nevertheless, this balanced corpus does not reflect the class imbalance in the uncurated data. Even very high precision results can produce a large number of false alarms in the natural, unbalanced corpus [7].

Harassment. We curated our own corpus for our high school harassment that addresses a number of corpus considerations [9], under an IRB-approved protocol requiring privacy protections through anonymization. This corpus, called ALONE, comprises of posts exchanged in interactions between pairs of participants to capture the appropriate contextual cues. Hence, each sample is an interaction that is an aggregate of posts between users along with other metadata. First, because the culture of the U.S. high school population is quite different from the general U.S. culture at large, we assured the identity of the participants. Starting with a seed set of known high school student names published in the newspaper as scholarship winners, we searched Twitter for unique matches to users with appropriate location indicators in their metadata. To grow the set, we searched on their Twitter contacts and then pruned the resulting list of candidates by requiring contacts with other members of the candidate list. Second, we make no assumptions regarding the nature of toxic content in assembling this corpus. Third, as we were concerned with capturing the full context for the individual post, we retrieved the history of exchanges between members and the multi-modal content of these exchanges including emoji and images which may also contain toxic content [9]. The diversity of modality enriches the interactions between humans and computers. Specifically, users create the context for their conversations using these modalities. As a result of our corpus assembly process, we can recover network structure [51] suitable for insider-outsider analysis. Finally, with the caveat of access restricted to public accounts, our corpus approximates a realistic class balance of benign and toxic content.

3.3. Need for Computationally Accessible External Knowledge

We advocate the use of relevant external knowledge in a variety of forms, including text sources and computationally accessible knowledge graphs (KGs). This assures attention to the different dimensions to account for subtle nuances in the semantics of toxic behavior. External knowledge constitutes a source of “ground truth” for evaluating message content. As we argue that toxic behavior

is multi-dimensional, leading to ambiguity and false alarms, we employ a multi-level and multi-dimensional approach that helps capture differences between various cultural and social senses of toxicity to resolve ambiguity. Our previous Person, Content, and Network (PCN) distinction [4, 52, 5, 53] functions at a higher (superficial) level, whereas the contextual dimensions of content (e.g., religion, ideology and violence) functions at the lower level [7], capturing the deep semantics of toxicity. Further, when incorporated into a classification algorithm, external knowledge enables opportunities to provide an explanation generally missing from contemporary deep learning approaches [54, 55, 56, 57].

Purohit et al. [25], provides a complementary framework for broader context to guide interpretation and evaluation. They identified three major dimensions of knowledge necessary to design humanity-inspired AI systems: personalization, social context, and intention. Here, we expand [25] to scope the relevant knowledge that they described in three dimensions: values, norms, and the domain. Each dimension is pegged by individual specificity and collective generality, and the perspective required to interpret the behavior of an individual is represented by the combination of all three dimensions. While [25] considered other actors as part of the environment, here, we consider them more explicitly (see Figure 2). They constitute a community, with norms and values. The concept of *personal semantics* for the target of toxicity covers much of what [25] intended in their analysis. Personal semantics includes knowledge about the targets’ language of insult, verbal abuse, and offensive language, involving sensitive topics specific to the individual and their social network. From the sources’ perspective, we require knowledge corresponding to their *intention*, particularly associated with indicators of power, truth, and trust [58]. Finally, the *history of interaction* such as duration and toxicity frequency between source and target requires knowledge about the structure of nominal conversation such as indicators of topic change and common ground that determine familiarity [59]. The target’s *emotional response* corresponds to the toxicity-specific emotion evoked in a recipient after reading messages, informed by conversation history and network membership. Toxicity detection requires a more sophisticated classification




Figure 2: Individuals are surrounded by the sources of data and knowledge required for computational analysis. Personalizing the analysis by incorporating personal semantics, intention, and emotion will help distinguish toxic behaviors from non-toxic. Further, infusing external knowledge will resolve ambiguity by better contextualizing multimodal data and providing a source of explanation.

scheme beyond binary toxicity, referring to knowledge related to the causes of experienced harm, embarrassment, loss of reputation, etc. as well as possible clinically relevant consequences such as depression and suicide [12]. While these sources of knowledge are typically not made explicit in toxicity analysis, the failure to make them explicit or acknowledge features corresponding to these contributes to disagreement among annotators and ultimately poor, and biased classification.

4. A Knowledge-enhanced Socio-technical Approach to Toxicity Detection

Toxicity detection takes the form of two problems: detection of the toxic source(s) and identification of the vulnerable victim. For both problems, we require more sophisticated Natural Language Processing (NLP) and Machine Learning (ML) methods to detect and use the features indicative of toxicity. Because the meaning of the content is personalized based on the belief system of the source and target, the semantic meaning needs to be computationally represented separately. Such personalized belief systems are critical for understanding how toxic behavior is interpreted differently by different individuals [60]. These inter-related concepts and beliefs also evolve over time upon exposure to new information [61]. The question here is how one can computationally

tionally model the evolution of such complex social exchange. We advocate a Knowledge-infused Learning (K-iL) framework [41, 62, 63] where the model learns to recognize patterns of different meanings of toxic concepts from different perspectives to reduce ambiguity. However, the knowledge sources are not necessarily at the same level of granularity and abstraction. Accordingly, we categorized knowledge infusion [41, 62, 7] as shallow, semi-deep, and deep infusion to resolve the impedance mismatch due to different representational forms and abstractions [63]. Infusing knowledge is particularly important for overcoming the inescapable limitations and biases of data-driven processing [64].

We propose a framework that will account for Purohit et al.’s dimensions [4] to generate richer representations including personal semantics, intention, emotion, history of interaction, and social context. This collection of information will require dynamic hybrid models for different modalities of data and knowledge representation. As behavioral models are dynamic and evolve, this framework should also allow for change. Further, validation of such an approach is also challenging and likely requires some form of experimentally controlled data collection to support supervised learning. The framework must address multiple levels of analysis, such as content, individual, and community, ensuring that the individual level details are changing as a function of interacting with their network. Communities form around various topics of interest through network interactions, where the shared content displays an intent attached with emotions. As learning concepts and grasping causal relations go beyond the data available, conceptual and probabilistic models can perform inference over hierarchies of structured representations [65].

Among the Purohit et al. dimensions [4], personal semantics, interactions, and social context can be represented using both conceptual (e.g., knowledge graphs) and probabilistic models (e.g., language, image). External knowledge can be represented in structured (e.g., knowledge graph) and semi-structured forms (e.g., JSON) to inform computation. While knowledge can be acquired from data through various methods, dependence on data significantly limits the search space and extraction of the complete knowledge that is required to

represent the complex nature of toxicity [66]. Explicit structural relations in a knowledge graph constitute context and capture the intrinsic characteristics of this problem, which can be incorporated into a statistical learning algorithm (e.g., neural networks) to enhance the latent contextual space. This incorporation will adjust emphasis on sparse-but-essential and irrelevant-but-frequent terms and concepts, boosting recall without reducing precision [7, 41, 62]. While probabilistic models (e.g., BERT, GPT-3, ResNet, Inception) have advanced in recent years, generating knowledge representations from knowledge graphs or similar structured forms of knowledge remains an open area for advancement. However, a knowledge graph can be represented as embedding vectors including structural information of the graph, such as relationships. Existing methods, such as TRANS-E [67], TRANS-H [68], and HOLE [69], can generate embeddings from a knowledge graph. The generated knowledge representation can then be infused within a probabilistic model.

In a learning architecture, represented knowledge can be infused through an attention mechanism and knowledge-based constraints or dependency relations between words in a sentence [63]. Deep infusion of knowledge is still an open area of research, as we described in [62, 41]. Deep infusion of knowledge combines the representation of structural knowledge graph content with a latent representation of data, quantifying the information loss and identifying the level of abstraction. The infusion of knowledge can take place after each epoch optimizing the loss function. In this architecture, for deep infusion, related functions add an additional layer that takes the latent vectors of the previous layers, and the knowledge embedding, merging them to output a knowledge-infused representation. In this framework, as we utilize multiple dimensions to represent toxic behavior, an appropriate infusion of knowledge will form connections within the data resulting in better-contextualized representation. As our prior work suggests, infusion of knowledge mitigates unfair outcomes by reducing false positives that would lead to adverse social outcomes. [7, 62, 12].

5. Conclusion

In this paper, we identified the multiple influences on the detection of toxic exchange beyond conventional content analysis. Our goal was to provide a framework that identifies and utilizes the multiple dimensions of toxicity and incorporates explicit knowledge in a statistical learning algorithm to resolve ambiguity. For toxicity detection, we provided a framework founded on behavioral and social theory. Specifically, we highlighted the significance of multi-level analysis of data, namely, content, individual, and community, and the numerous features necessary to determine toxicity. Knowledge representation and its infusion in a learning algorithm is an emergent solution for toxicity detection and related sets of similar problems.

Acknowledgement

This work is funded in part by National Science Foundation Award 1761931 Spokes: MEDIUM: MIDWEST: Collaborative: Community-Driven Data Engineering for Substance Abuse Prevention in the Rural Midwest. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] G. K. Young, How much is too much: the difficulties of social media content moderation, *Information & Communications Technology Law* (2021) 1–16.
- [2] M. Nagarajan, H. Purohit, A. Sheth, A qualitative examination of topical tweet and retweet practices, in: *Fourth International AAAI Conference on Weblogs and Social Media*, 2010.
- [3] A. P. Sheth, H. Purohit, A. S. Jadhav, P. Kapanipathi, L. Chen, Understanding events through analysis of social media.

- [4] H. Purohit, Y. Ruan, A. Joshi, S. Parthasarathy, A. Sheth, Understanding user-community engagement by multi-faceted features: A case study on twitter, in: WWW 2011 Workshop on Social Media Engagement (SoME), 2011.
- [5] U. Kursuncu, M. Gaur, U. Lokala, A. Illendula, K. Thirunarayan, R. Daniulaityte, A. Sheth, I. B. Arpinar, What's ur type? contextualized classification of user types in marijuana-related communications using compositional multiview embedding, in: 2018 IEEE/WIC/ACM International Conference on Web Intelligence (WI), IEEE, 2018, pp. 474–479.
- [6] H. Purohit, T. Banerjee, A. Hampton, V. L. Shalin, N. Bhandutia, A. P. Sheth, Gender-based violence in 140 characters or fewer: A# bigdata case study of twitter, arXiv preprint arXiv:1503.02086.
- [7] U. Kursuncu, M. Gaur, C. Castillo, A. Alambo, K. Thirunarayan, V. Shalin, D. Achilov, I. B. Arpinar, A. Sheth, Modeling islamist extremist communications on social media using contextual dimensions: religion, ideology, and hate, Proceedings of the ACM on Human-Computer Interaction 3 (CSCW) (2019) 1–22.
- [8] M. Rezvan, S. Shekarpour, F. Alshargi, K. Thirunarayan, V. L. Shalin, A. Sheth, Analyzing and learning the language for different types of harassment, Plos one 15 (3) (2020) e0227330.
- [9] T. Wijesiriwardene, H. Inan, U. Kursuncu, M. Gaur, V. L. Shalin, K. Thirunarayan, A. Sheth, I. B. Arpinar, Alone: A dataset for toxic behavior among adolescents on twitter, in: International Conference on Social Informatics, Springer, 2020, pp. 427–439.
- [10] W. Wang, L. Chen, K. Thirunarayan, A. P. Sheth, Cursing in english on twitter, in: Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing, 2014, pp. 415–425.

- [11] A. Sheth, P. Kapanipathi, Semantic filtering for social data, *IEEE Internet Computing* 20 (4) (2016) 74–78.
- [12] U. Kursuncu, H. Purohit, N. Agarwal, A. Sheth, When the bad is good and the good is bad: Understanding cyber social health through online behavioral change, *IEEE Internet Computing* 25 (01) (2021) 6–11.
- [13] T. Jay, *Cursing in America: A Psycholinguistic Study of Dirty Language in the Courts, in the Movies, in the Schoolyards, and on the Streets*, John Benjamins Publishing, 1992.
- [14] T. Jay, The utility and ubiquity of taboo words, *Perspectives on psychological science* 4 (2) (2009) 153–161.
- [15] M. R. Mehl, J. W. Pennebaker, The sounds of social life: a psychometric analysis of students’ daily social environments and natural conversations., *Journal of personality and social psychology* 84 (4) (2003) 857.
- [16] Z. Waseem, J. Thorne, J. Bingel, Bridging the gaps: Multi task learning for domain transfer of hate speech detection, in: *Online harassment*, Springer, 2018, pp. 29–55.
- [17] J. Golbeck, Z. Ashktorab, R. O. Banjo, A. Berlinger, S. Bhagwan, C. Buntain, P. Cheakalos, A. A. Geller, R. K. Gnanasekaran, R. R. Gunasekaran, et al., A large labeled corpus for online harassment research, in: *Proceedings of the 2017 ACM on web science conference*, 2017, pp. 229–233.
- [18] S. Einarsen, H. Hoel, D. Zapf, C. Cooper, *Bullying and harassment in the workplace: Developments in theory, research, and practice*, CRC press, 2010.
- [19] S. Hinduja, J. W. Patchin, Bullying, cyberbullying, and suicide, *Archives of suicide research* 14 (3) (2010) 206–221.
- [20] W. B. Roberts, D. H. Coursol, Strategies for intervention with childhood and adolescent victims of bullying, teasing, and intimidation in school settings, *Elementary School Guidance & Counseling* 30 (3) (1996) 204–212.

- [21] G. D. Cooper, P. T. Clements, K. E. Holt, Examining childhood bullying and adolescent suicide: Implications for school nurses, *The Journal of School Nursing* 28 (4) (2012) 275–283.
- [22] E. A. Crone, E. A. Konijn, Media use and brain development during adolescence, *Nature communications* 9 (1) (2018) 1–10.
- [23] W. Pedersen, V. Vestel, A. Bakken, At risk for radicalization and jihadism? a population-based study of norwegian adolescents, *Cooperation and conflict* 53 (1) (2018) 61–83.
- [24] S. Ozer, M. Obaidi, S. Pfattheicher, Group membership and radicalization: A cross-national investigation of collective self-esteem underlying extremism, *Group Processes & Intergroup Relations* 23 (8) (2020) 1230–1248.
- [25] H. Purohit, V. L. Shalin, A. P. Sheth, Knowledge graphs to empower humanity-inspired ai systems, *IEEE Internet Computing* 24 (4) (2020) 48–54.
- [26] S. Henry, School violence beyond columbine: A complex problem in need of an interdisciplinary analysis, *American Behavioral Scientist* 52 (9) (2009) 1246–1265.
- [27] U. Kursuncu, Y. Mejova, J. Blackburn, A. Sheth, Cyber social threats 2020 workshop meta-report: Covid-19, challenges, methodological and ethical considerations, *Workshop Proceedings of the 14th International AAAI Conference on Web and Social Media*.
- [28] U. Kursuncu, J. Blackburn, Y. Mejova, M. Squire, A. Sheth, Cyber social threats 2021: Ai, covid-19 vaccine, detection and countering strategies.
- [29] D. Noever, Machine learning suites for online toxicity detection, *arXiv preprint arXiv:1810.01869*.
- [30] J. Pavlopoulos, J. Sorensen, L. Dixon, N. Thain, I. Androutsopoulos, Toxicity detection: Does context really matter?, *arXiv preprint arXiv:2006.00998*.

- [31] K. McGuffie, A. Newhouse, The radicalization risks of gpt-3 and advanced neural language models, arXiv preprint arXiv:2009.06807.
- [32] A. Olteanu, C. Castillo, F. Diaz, E. Kıcıman, Social data: Biases, methodological pitfalls, and ethical boundaries, *Frontiers in Big Data* 2 (2019) 13.
- [33] S. Gehman, S. Gururangan, M. Sap, Y. Choi, N. A. Smith, Realtoxicityprompts: Evaluating neural toxic degeneration in language models, arXiv preprint arXiv:2009.11462.
- [34] S. Groenwold, L. Ou, A. Parekh, S. Honnavalli, S. Levy, D. Mirza, W. Y. Wang, Dats wassup!!: Investigating african-american vernacular english in transformer-based text generation, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 5877–5883.
- [35] E. Wallace, S. Feng, N. Kandpal, M. Gardner, S. Singh, Universal adversarial triggers for attacking and analyzing nlp, arXiv preprint arXiv:1908.07125.
- [36] H. Zhang, A. X. Lu, M. Abdalla, M. McDermott, M. Ghassemi, Hurtful words: quantifying biases in clinical contextual word embeddings, in: *proceedings of the ACM Conference on Health, Inference, and Learning*, 2020, pp. 110–120.
- [37] N. Garg, L. Schiebinger, D. Jurafsky, J. Zou, Word embeddings quantify 100 years of gender and ethnic stereotypes, *Proceedings of the National Academy of Sciences* 115 (16) (2018) E3635–E3644.
- [38] I. Y. Chen, P. Szolovits, M. Ghassemi, Can ai help reduce disparities in general medical and mental health care?, *AMA journal of ethics* 21 (2) (2019) 167–179.
- [39] L. Hanu, J. Thewlis, S. Haco, How ai is learning to identify toxic online content, *Scientific American*.

- [40] E. J. Topol, High-performance medicine: the convergence of human and artificial intelligence, *Nature medicine* 25 (1) (2019) 44–56.
- [41] U. Kursuncu, Modeling the persona in persuasive discourse on social media using context-aware and knowledge-driven learning, Ph.D. thesis, University of Georgia (2018).
- [42] I. B. Arpinar, U. Kursuncu, D. Achilov, Social media analytics to identify and counter islamist extremism: Systematic detection, evaluation, and challenging of extremist narratives online, in: 2016 International Conference on Collaboration Technologies and Systems (CTS), IEEE, 2016, pp. 611–612.
- [43] D. Yin, Z. Xue, L. Hong, B. D. Davison, A. Kontostathis, L. Edwards, Detection of harassment on web 2.0, *Proceedings of the Content Analysis in the WEB 2* (2009) 1–7.
- [44] G. Kennedy, A. McCollough, E. Dixon, A. Bastidas, J. Ryan, C. Loo, S. Sahay, Technology solutions to combat online harassment, in: *Proceedings of the first workshop on abusive language online*, 2017, pp. 73–77.
- [45] A. Bastidas, E. Dixon, C. Loo, J. Ryan, Harassment detection: a benchmark on the# hackharassment dataset, arXiv preprint arXiv:1609.02809.
- [46] M. Bugueño, M. Mendoza, Learning to detect online harassment on twitter with the transformer., *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*.
- [47] M. Rezvan, S. Shekarpour, L. Balasuriya, K. Thirunarayan, V. L. Shalin, A. Sheth, A quality type-aware annotated corpus and lexicon for harassment research, in: *Proceedings of the 10th ACM Conference on Web Science*, 2018, pp. 33–36.
- [48] M. Fernandez, M. Asif, H. Alani, Understanding the roots of radicalisation on twitter, in: *Proceedings of the 10th acm conference on web science*, 2018, pp. 1–10.

- [49] E. Ferrara, W.-Q. Wang, O. Varol, A. Flammini, A. Galstyan, Predicting online extremism, content adopters, and interaction reciprocity, in: International conference on social informatics, Springer, 2016, pp. 22–39.
- [50] L. Chen, I. Weber, A. Okulicz-Kozaryn, Us religious landscape on twitter, in: International Conference on Social Informatics, Springer, 2014, pp. 544–560.
- [51] S. Bhatt, S. Padhee, A. Sheth, K. Chen, V. Shalin, D. Doran, B. Minnery, Knowledge graph enhanced community detection and characterization, in: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, 2019, pp. 51–59.
- [52] A. Sheth, A. Jadhav, P. Kapanipathi, C. Lu, H. Purohit, G. A. Smith, W. Wang, Twitris: A system for collective social intelligence, encyclopedia of social network analysis and mining, 2014J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 2 (2014) 68–73.
- [53] U. Kursuncu, M. Gaur, U. Lokala, K. Thirunarayan, A. Sheth, I. B. Arpinar, Predictive analysis on twitter: Techniques and applications, in: Emerging research challenges and opportunities in computational social network analysis and mining, Springer, 2019, pp. 67–104.
- [54] X. Wang, D. Wang, C. Xu, X. He, Y. Cao, T.-S. Chua, Explainable reasoning over knowledge graphs for recommendation, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33, 2019, pp. 5329–5336.
- [55] F. Lecue, On the role of knowledge graphs in explainable ai, Semantic Web 11 (1) (2020) 41–51.
- [56] M. Gaur, K. Faldu, A. Sheth, Semantics of the black-box: Can knowledge graphs help make deep learning systems more interpretable and explainable?, IEEE Internet Computing 25 (1) (2021) 51–59.
- [57] R. Bansal, W. S. Paka, S. Sengupta, T. Chakraborty, et al., Combining exogenous and endogenous signals with a semi-supervised co-

attention network for early detection of covid-19 fake tweets, arXiv preprint arXiv:2104.05321.

- [58] R. C. Mayer, J. H. Davis, F. D. Schoorman, An integrative model of organizational trust, *Academy of management review* 20 (3) (1995) 709–734.
- [59] I. Hutchby, Conversation analysis, *The Wiley-Blackwell Encyclopedia of Social Theory* (2017) 1–9.
- [60] N. E. Friedkin, A. V. Proskurnikov, R. Tempo, S. E. Parsegov, Network science on belief system dynamics under logic constraints, *Science* 354 (6310) (2016) 321–326.
- [61] J. L. Usó-Doménech, J. Nescolarde-Selva, What are belief systems?, *Foundations of Science* 21 (1) (2016) 147–152.
- [62] U. Kursuncu, M. Gaur, A. Sheth, Knowledge infused learning (k-il): Towards deep incorporation of knowledge in deep learning, *Proceedings of the AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice (AAAI-MAKE 2020)*.
- [63] A. Sheth, M. Gaur, U. Kursuncu, R. Wickramarachchi, Shades of knowledge-infused learning for enhancing deep learning, *IEEE Internet Computing* 23 (6) (2019) 54–63.
- [64] A. P. Sheth, K. Thirunarayan, *The inescapable duality of data and knowledge*, 2021.
- [65] J. B. Tenenbaum, C. Kemp, T. L. Griffiths, N. D. Goodman, How to grow a mind: Statistics, structure, and abstraction, *science* 331 (6022) (2011) 1279–1285.
- [66] L. G. Valiant, Robust logics, *Artificial Intelligence* 117 (2) (2000) 231–253.
- [67] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: *Neural Information Processing Systems (NIPS)*, 2013, pp. 1–9.

- [68] Z. Wang, J. Zhang, J. Feng, Z. Chen, Knowledge graph embedding by translating on hyperplanes, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 28, 2014.
- [69] M. Nickel, L. Rosasco, T. Poggio, Holographic embeddings of knowledge graphs, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 30, 2016.