



Toxicity in online platforms and AI systems: A survey of needs, challenges, mitigations, and future directions

Smita Khapre  ^{a,*}, Melkamu Abay Mersha  ^a, Hassan Shakil  ^a, Jonali Baruah  ^b,
Jugal Kalita  ^a

^a College of Engineering and Applied Science, University of Colorado Colorado Springs, CO, 80918, USA

^b Tarleton State University A Member of The Texas A & M University System, Stephenville, TX, 76401, USA

ARTICLE INFO

Keywords:
 Toxicity
 Fallacy
 Fake news
 Misinformation
 Social media
 Large language models
 Hate speech
 LLM safety personality
 Traits

ABSTRACT

The evolution of digital communication systems and the designs of online platforms have inadvertently facilitated the subconscious propagation of toxic behavior. Giving rise to reactive responses to toxic behavior. Toxicity in online content and Artificial Intelligence Systems has become a serious challenge to individual and collective well-being around the world. It is more detrimental to society than we realize. Toxicity, expressed in language, image, and video, can be interpreted in various ways depending on the context of usage. Therefore, a comprehensive taxonomy is crucial to detect and mitigate toxicity in online content, Artificial Intelligence systems, and/or Large Language Models in a proactive manner. A comprehensive understanding of toxicity is likely to facilitate the design of practical solutions for toxicity detection and mitigation. The classification in published literature has focused on only a limited number of aspects of this very complex issue, with a pattern of reactive strategies in response to toxicity. This survey attempts to generate a comprehensive taxonomy of toxicity from various perspectives. It presents a holistic approach to explain the toxicity by understanding the context and environment that society is facing in the Artificial Intelligence era. This survey summarizes the toxicity-related datasets and research on toxicity detection and mitigation for Large Language Models, social media platforms, and other online platforms, detailing their attributes in textual mode, focused on the English language. Finally, we suggest the research gaps in toxicity mitigation based on datasets, mitigation strategies, Large Language Models, adaptability, explainability, and evaluation.

1. Introduction

The rise of online platforms, social media, and Artificial Intelligence (AI) has revolutionized how people communicate, share information, and express opinions (Gao et al., 2020; Villate-Castillo et al., 2024). While these technologies provide significant benefits, they are also the source of toxicity, like hate speech, cyberbullying, misinformation, and harassment (Lin et al., 2024b). The pervasive nature of toxicity in online platforms and AI systems impacts individuals, society, and businesses, and has regulatory implications (Kiritchenko et al., 2021).

Many synonymous terms are used in industry and academia to refer to toxicity in digital media. A partial list includes words and phrases such as hate, hate speech, offensive language, flaming, incivility, risk, malicious, harmful, hateful sexual, illegal and abusive content, information hazards, discrimination, bias, fake news, misinformation, human-

computer-interaction harms, and cyberbullying (Kirk et al., 2022; Wang et al., 2024d). Toxicity is further categorized into implicit and explicit, based on target groups, political and social bias, illegal, violent, aggressive, and more (Pachinger et al., 2023). These are most prevalent on social media and other online platforms. Toxicity and harm have been interpreted as distinct by some research (Wang et al., 2024d). In this survey, we treat toxicity and harm as synonyms, with similar impacts. Markov et al. (2023) highlighted the challenges in real-world toxicity detection in Generative AI (GenAI). They emphasized the importance of a well-motivated taxonomy of toxic content being fundamental to their detection.

The research for detecting the toxicity in online platforms and AI Systems has picked up pace due to its ubiquitous and adverse psychological impacts on society. Misinformation and fake news have a dramatic effect on public opinion and are likely to move society toward extremism

* Corresponding author.

E-mail addresses: skhapre@uccs.edu (S. Khapre), mmersha@uccs.edu (M.A. Mersha), hshakil@uccs.edu (H. Shakil), baruah@tarleton.edu (J. Baruah), jkalita@uccs.edu (J. Kalita).



Fig. 1. LLMs depicted as Kids, trained on Toxicity Ignorant Datasets, unable to distinguish toxic language from normal language. Just as kids grow up in their toxic environments unaware that it is toxic, when grow up as adults start displaying toxic behavior same as their environments.

(Gjerazi & Skana, 2023; Lazer et al., 2018). Although individuals of all generations are negatively impacted, the influence on adolescents' faculties seems to be most pronounced (Lemaire et al., 2025).

Platform economy has taken over communication, and most platforms either are beset with toxicity issues or are prone to harbor toxicity in various ways (Villate-Castillo et al., 2024). Large Language Models (LLMs) can be leveraged to detect toxicity in online platforms and detoxify them (Gallegos et al., 2024). LLMs are trained on very large language corpora. The datasets are usually annotated, and not filtered, containing all types of implicit toxic words, explicit toxic words, and social bias. Thus, LLMs trained on such a dataset also propagate the inherent toxicity of its dataset. An analogical depiction is shown in Fig. 1, where LLMs are depicted as little kids unaware of toxicity, existing in an environment where toxicity abounds. Kids grow up and become an integral part of society with inherent toxic behavior, unaware that it is detrimental and that it should be avoided. LLMs are the underlying technology in GenAI, which generates human-like responses to given prompts. Toxicity is seen in their responses too. Moreover, using prompt engineering, prompt injection, and jail-breaking attacks, toxicity can be deliberately introduced (Liu et al., 2024; Wei et al., 2024; Xu et al., 2024). While there has been tremendous improvement in GenAI, there is a need to implement improved toxicity detection and detoxification (Kim et al., 2024).

There is significant research on detection and detoxification of explicit, implicit, and biased toxicities. With current technological advancements, dealing with toxicity requires a careful understanding of the environment and platform. Derave et al. developed a generic module of digital platform (DP) ontology, based on DP attributes, DP attribute values, DP software functionalities, DP types, and DP supported communication (Derave et al., 2021). It provides DP design gaps in DPs. The advent of Web 2.0, linguistic toxicity was predominant in communication. Most research dealing with toxicity still focuses only on linguistic expressions, and is usually classified as implicit and explicit toxicity (Lewandowska-Tomaszczyk et al., 2023). Their semantics of the words or phrases are used for classification as implicit or explicit. To mitigate the impacts of toxicity comprehensively, we need to understand the environment in which it appears. *Toxicity is a fatal disease that not only requires diagnosis but also cure and prevention.* This metaphor implies that the *Toxicity* is a disease, *diagnosis* is detection, *cure* is detoxification, and *prevention* requires an understanding of the environment and its source and impacts on society.

While there is tremendous effort in the research community to generate a taxonomy for toxicity in digital communication, with every passing year and the advent of new technology, it is never complete. Our survey provides a structured overview of toxicity in online platforms and AI systems. Based on the current state-of-the-art, we examine key areas such as detection methods, mitigation strategies, detoxification techniques, and bias-related concerns. Beyond technical approaches, we also explore the psychological and societal consequences of toxic content, along with the limitations of existing datasets. This work aims to inform future research and support the development of safer, more inclusive digital ecosystems. The main contributions of our survey are as follows.

- We present a holistic taxonomy of toxicity, capturing diverse forms of harmful behavior, language, and content, in the context of various platforms, including insights from human psychology.
- We systematically review over 200 studies addressing toxicity detection, mitigation, detoxification, associated biases, and the evolution of digital communication and online platforms.
- We evaluate major toxicity datasets identifying key limitations such as bias, annotation inconsistency, and coverage gaps.
- We examine the broader psychological and societal impacts of toxic content on individuals and communities.
- We identify open challenges and outline future research directions to advance responsible and inclusive online and AI systems.

The remainder of this paper is structured as follows. In Section 2, we discuss *Background and Motivation* for this survey. It includes the basic terminologies used, why the study of toxicity matters, and outlines the stakeholders of the toxicity ecosystem. In Sections 4 and 5, we discuss a taxonomy of toxicity we have developed (Fig. 3). In Section 6, we examine the impacts of toxicities on human psychology. In Section 7, we collate and discuss the toxicity datasets, followed by detection and detoxification research. In Section 8, we present the open challenges that still need to be addressed and future directions. Finally, in Section 9, we draw the conclusion of this survey.

2. Background and motivation

Modes of communication have evolved through history. Since 1792, there has been rapid technological progress in communication. Dhingra and Mudgal (2019) navigate the evolution of communication technology from the telegraph in 1792, the telephone in 1876, email (ARPANET) in 1971, the World Wide Web in 1989, and Social Media started in 2003. The emergence of Social Media Platforms (SMP) such as Facebook in 2004, YouTube in 2005, Twitter in 2006, WhatsApp in 2009, Pinterest, and Instagram in 2010 have drastically altered the digital communication landscape. Currently, AI-powered Social Media have taken over channels of communication. Vayadande et al. (2024) and Shen et al. (2024) navigate through the evolution and impacts of Web 1.0 (1991–2004), Web 2.0 (2004–2024) to Web 3.0 (2020 onwards). The use of AI improves user experience during enhanced communication, and the use of blockchain technology achieves safety, trust, privacy, and decentralization. Wang et al. (2024c) point to the increase in probability of toxicity in Web 3.0 content, requiring effective governance policies while we move toward decentralization in Web 3.0. We describe the basic terminologies used before getting into our motivation of this survey.

2.1. Selected basic terminologies

- **Toxicity** is a term we use to signify a substance or language construct or media element that highly impact an individual in a negative manner. In online context, it is often derogatory, rude, harmful, and disrespectful. Toxic content induces stress, anxiety, suicidal thoughts, depression, and loneliness. These are the least among the impacts on society (Akar, 2025; Maleki et al., 2022).

- **Implicit Toxicity** is the use of language or media elements that cannot be straightforwardly interpreted as toxic, but which suggest an implied toxic meaning. It introduces toxicity indirectly, which often targets one or more individuals, races, or genders in context (Bączkowska et al., 2025). It often involves use of figures-of-speech like metaphor, simile, irony or sarcasm.
- **Explicit Toxicity** is the direct and intentional use of toxic language or media elements. It is aggressive and offensive. Bansal et al. describe it as direct and intentional flaming (Bansal et al., 2011; Lewandowska-Tomaszczyk et al., 2023).
- **Fake News** is not authentic news, created intentionally and propagated via diverse users (Hu et al., 2025). It mimics real news construct.
- **Misinformation** is incorrect information whose authenticity can be verified as false but is generated unintentionally (Scheufele & Krause, 2019). It does not mimic real news. Disinformation is similar to misinformation, but is intentionally created false information. A rumor is unverified information that could be true or false.
- **Fallacy** (Helwe et al., 2024) provides logical reasoning behind untrue construct or falsehood. It is divided into three main types: ethos (assuming false credibility), logos (false logic), and pathos (appealing to emotions). Fallacies provide support to fake news, misinformation, and disinformation dissemination.
- **Illegal Content** is comprised of unauthorized or authorized users selling goods to evade taxes or the sale of smuggled goods, controlled substances, weapons, wildlife, sexually explicit content or trafficking in human, goods, or animals. Illegal content is unlawful and is regulated by law (Dangsawang & Nuchitprasitchai, 2024; Mademlis et al., 2024; Roy & Kumar, 2024; Singh & Nambiar, 2024).
- **Cyberbullying** is bullying or threatening behavior on digital technology platforms (Ray et al., 2024).
- **Modes** are the media used to communicate ideas, information, opinions, etc. They include text, audio, and video (Chandler & Munday, 2011).
- **Engagement** refers to the ability of digital media to keep the user highly interested or engaged.
- **Bias** is a false pre-conceived notion based on skewed data or information or expectation assessment or evaluation impacting a target population (Delgado-Rodriguez & Llorca, 2004).
- **Implicit Bias** is an unconscious inclination of one's behavior towards or against others' race, identity, culture, background, color, ability, community, gender, age, sexuality, nationality, status, religion, and so on (Banaji et al., 2015).
- **Social Bias** is an occurrence of societal discrimination against a target group, person, ideas, or beliefs based on prejudices and stereotypes prevalent in the society (Webster et al., 2022). Sap et al. (2020a) developed so-called social bias frames to generate the Social Bias Inference (SBIC) Dataset.
- **Algorithmic Bias** occurs when an algorithm outputs favors or ignores particular target groups or individuals for the wrong reasons, causing imbalanced impacts (Kordzadeh & Ghasemaghaei, 2022). It can be divided in two categories. One is when Social Bias, as explained earlier, influences the training dataset and generates the same bias in outputs (Johnson, 2021). The second is an intentional bias built into the platforms to gear the users' engagement and revenue generation. The Netflix Documentary "The Social Dilemma," directed by Jeff Orlowski in 2020, has explained it exceptionally well.

2.2. Why does toxicity matters?

Understanding the implications of online toxicity is crucial to address and develop effective mitigation strategies. The implications include psychological harm, social division, proliferation of misinformation, ethical challenges, legal complexities, and economic disruptions (Kiritchenko et al., 2021).

Mental Health Impact on an individual: Toxicity emanating from online and AI systems can significantly harm an individual's mental health (Si et al., 2022). Exposure to online cyberbullying, harassment, and hate speech has been linked to increased anxiety, depression, stress, self-harm, and suicidal ideation, particularly among adolescents and marginalized communities (Twenge, 2020). Excessive exposure to online toxic content can disrupt sleep patterns and elevate self-harm risks (Twenge, 2020). The World Health Organization (WHO) has recognized problematic internet use as a growing public health concern, with online toxicity playing a significant role in reinforcing compulsive and harmful digital behaviors (Kumar & Mondal, 2018; Sandua, 2024).

Societal Impact: Toxicity arising from Online and AI systems extends beyond individual harm, fueling societal issues like polarization, proliferation of mental health crises, and unabated spread of misinformation (Fan et al., 2024a). AI algorithms amplify ideological divisions by reinforcing pre-existing beliefs, making individuals more resistant to opposing views (Biju & Gayathri, 2023; Guess & Munger, 2023).

Ethical Concerns: The ethical implications afflicting online and AI systems are deeply rooted in questions of algorithmic responsibility, bias, and digital rights (Cheng et al., 2021). AI-driven content moderation and recommendation systems often contain inherent biases that unequally impact certain demographic groups (Liu, 2024). Unregulated toxicity can distort public discourse, spreading misinformation, amplifying divisive narratives, and undermining trust in reliable information sources (Manheim & Kaplan, 2019).

Legal and Regulatory Implications: Governments and international bodies recognize the need for legal frameworks to regulate online toxicity and moderation of AI-driven content (Evangeline, 2025). The European Union's Digital Services Act (DSA) and General Data Protection Regulation (GDPR) have established basic content moderation and user protection standards, setting a global precedent for the regulation of digital platforms (Bradford, 2023). As digital communications continue to shape global discourse, addressing the challenges posed by toxicity in online platforms and AI systems is critical for ensuring a safer, more responsible, and equitable digital future (Stahl, 2021).

Business and Technological Risks: Toxicity is more than a social, ethical, legal, or psychological issue; it is a silent killer (George, 2023). It is also a significant risk to the integrity, reputation, and long-term sustainability of online platforms, businesses, and AI-driven technologies (Polemi et al., 2024). When users encounter abusive content, misinformation, or targeted harassment, they are more likely to disengage affecting user retention, lower engagement rates, and a weakening sense of community (Throuvala et al., 2021). The erosion of trust has far-reaching consequences. Toxicity can result in income loss, advertising reduction, and increased scrutiny and regulations (Beknazár-Yuzbashev et al., 2025). It damages brand credibility, discourages potential investors, and creates legal liabilities (Keller, 2018). It challenges the credibility of automated systems. Addressing toxicity is not just an ethical obligation, but a crucial business strategy. Ensuring a safe and trustworthy environment is essential for long-term growth, user satisfaction, resilience, and technological reliability (Polemi et al., 2024).

2.3. Stakeholders of the toxicity ecosystem

The prevalence of toxicity in online platforms and AI systems necessitates a collaborative effort among key stakeholders (Ognibene et al., 2023). Key stakeholders include individuals, platform owners, policy-makers, AI developers, researchers, clinicians, businesses, media and news organizations, and educational and psychological testing institutions. Online users and individuals should be responsible for spreading and moderating toxic content, and SMPs should implement AI-driven moderation strategies (Siapera, 2022). Online communities shape self-regulation norms, contributing to either the mitigation or amplification of toxicity (Zhang, 2024). AI developers create automated moderation systems and toxicity control tools to enhance online safety and user experience (Zhang et al., 2022). Researchers study toxicity patterns and

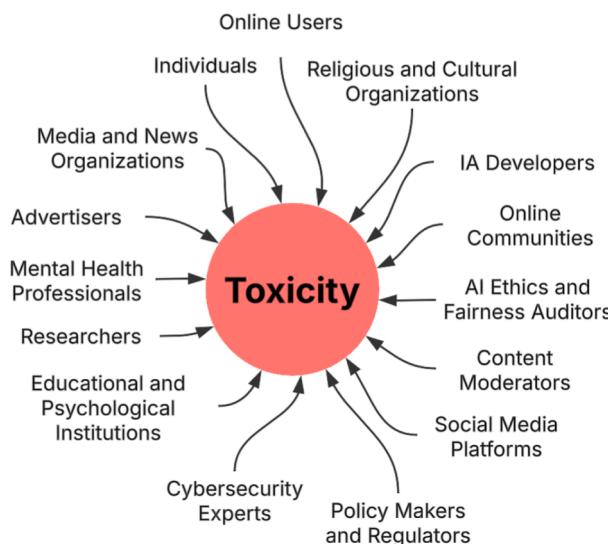


Fig. 2. Key Stakeholders in the ecosystem of online platforms and AI systems. Each stakeholder has a role in the generation, detection, moderation, and mitigation of toxic content.

biases to improve detection mechanisms (Garg et al., 2023). Policy-makers and regulators implement and enforce regulations to hold platforms accountable for toxic content, fostering a safer digital environment (Flew, 2021). Businesses and advertisers prioritize brand safety by advocating for faithful content policies and responsible platform practices (Brown, 2020). Media and news organizations influence public discourse on online platforms and toxicity of AI systems (Salminen et al., 2020). Educational institutions promote digital literacy and study the psychological impact of toxic interactions (Wardle & Derakhshan, 2017). Fig. 2 shows key stakeholders in the Online platforms and AI systems ecosystem.

3. Systematic review methodology

This section outlines the systematic methodology adopted to collect and analyze the literature on toxicity in online platforms and AI systems. Following established guidelines for systematic reviews, we employed the PRISMA framework (Moher et al., 2009) to ensure reproducibility and transparency. The review process consisted of three main phases: keyword formulation, source selection, and document filtering.

3.1. Keyword formulation

Given the interdisciplinary and evolving nature of toxicity in research on online platforms and AI systems, we designed a comprehensive keyword taxonomy spanning multiple conceptual and technical dimensions. Specifically, the search space was structured into six thematic categories, as shown in Table 1. This taxonomy ensured the retrieval of literature that addresses the phenomenon from various perspectives, including detection algorithms, social implications, model debiasing, explainability, and mitigation strategies. The complete set of keywords is presented in Table 1.

3.2. Search source

We performed a comprehensive literature search across four prominent academic repositories: *Google Scholar*, *ACM Digital Library*, *IEEE Xplore*, and the *ACL Anthology*. These platforms were selected for their extensive coverage of computational linguistics, artificial intelligence, and natural language processing research.

Keyword queries were formulated using logical OR operators within each category and, where appropriate, combined using AND to ensure

contextual specificity. For example: (“Toxicity” OR “Hate speech” OR “Cyberbullying”) AND (“Toxicity detection” OR “LLM detoxification”). The search covered works published between 2021 and 2025, with an English-language restriction applied to narrow the scope and exclude studies focused on multilingual and low-resource language settings. Duplicate records were removed before the screening process.

3.3. Document screening and inclusion criteria

An initial corpus of 4,094 documents was retrieved from the four selected databases. After removing duplicate records, a total of 2023 unique entries remained for screening. Our team members independently reviewed the titles and abstracts of these entries to assess their relevance to the scope of this review. During this phase, studies were excluded if they were non-peer-reviewed, had fewer citations, were outside the scope of computational sciences, or lacked methodological or empirical relevance to toxicity in online platforms and AI systems.

Following the initial screening, 545 papers were identified for full-text review. These articles were independently assessed by our team based on the following inclusion criteria: (1) the study addressed the detection or mitigation of hate speech, toxicity, offensive language, or cyberbullying; (2) it employed computational approaches relevant to online platforms or large language models (LLMs); and (3) it examined at least one key dimension such as bias mitigation, detoxification, or psychological impact. This structured evaluation resulted in the final inclusion of 247 peer-reviewed publications for in-depth analysis. Among these, 232 focused on tasks related to toxicity in online platforms and AI systems.

Our methodology ensures a comprehensive and diverse corpus, representative of current research in toxicity detection, and provides a solid foundation for the analyses presented in subsequent sections.

4. Modalities and sources of toxicity

4.1. Modalities of toxicity

Toxicity in online platforms and AI systems exists in various modalities, including text, images, audio, video, and multi-modal. The severity of harm caused depends on how the content is consumed, shared, and interpreted. The speed and reach of digital communication have accelerated the growth of modern online platforms, also amplifying their effects (Skiba, 2024). Each mode of toxicity presents unique challenges in detection, moderation, and mitigation. This section provides an in-depth discussion of the different modes of toxicity, their impact, and why they are considered toxic in online platforms and AI systems.

4.1.1. Text-based toxicity

Text-based toxicity is one of the most widespread and studied modalities of toxic and harmful content. The ability to quickly generate, modify, and distribute text across multiple platforms makes it a powerful medium for constructive and destructive discourse (Keen, 2011). Textual toxicity is particularly harmful due to its scalability. One of the primary reasons text toxicity is harmful is its direct psychological impact (Saha et al., 2019; Scheuerman et al., 2021). For instance, harassment, hate speech, and aggressive language can cause stress, anxiety, and depression (Bilewicz & Soral, 2020). Text-based toxicity is aggravated by anonymity and platform structures that enable rapid dissemination (Singh et al., 2022).

4.1.2. Image-based toxicity

Toxic images may display not only explicit content, but also spread misinformation and visual hate speech, and become viral (Pandiani et al., 2025). Images have an extraordinary power because they are processed faster by the human brain than text and are often more emotionally striking (Greenfield, 2015). A significant challenge with image-based toxicity is its ability to bypass traditional moderation methods

Table 1

Keyword categories used for the systematic literature review. Each column represents a thematic category, listing the search terms used to collect relevant literature on toxicity.

Toxicity & Hate Speech	Psychological & Social Impact	Mitigation & Detoxification	Detection Techniques	Bias and Fairness	Technical Terms
Toxicity taxonomy Implicit toxic language Hate speech Offensive language Offensive language taxonomy Abusive language Cyberbullying	Psychological impacts of language toxicity Mental health and online abuse Effect of toxic language on users Emotional impact of hate speech	Toxicity detection Language detoxification LLM detoxification Hate speech moderation Detoxifying large language models Mitigation of toxicity in online platforms Toxic content filtering Bias mitigation in LLMs Counterspeech detoxification	Toxicity detection Offensive language detection Hate speech detection Toxicity detection using LLMs Toxicity taxonomy Bias detection	Bias in language models Social bias detection in online platforms Fairness in large language models Bias in LLMs	Explainable AI in toxicity detection Evaluation metrics for toxicity Online platform evolution

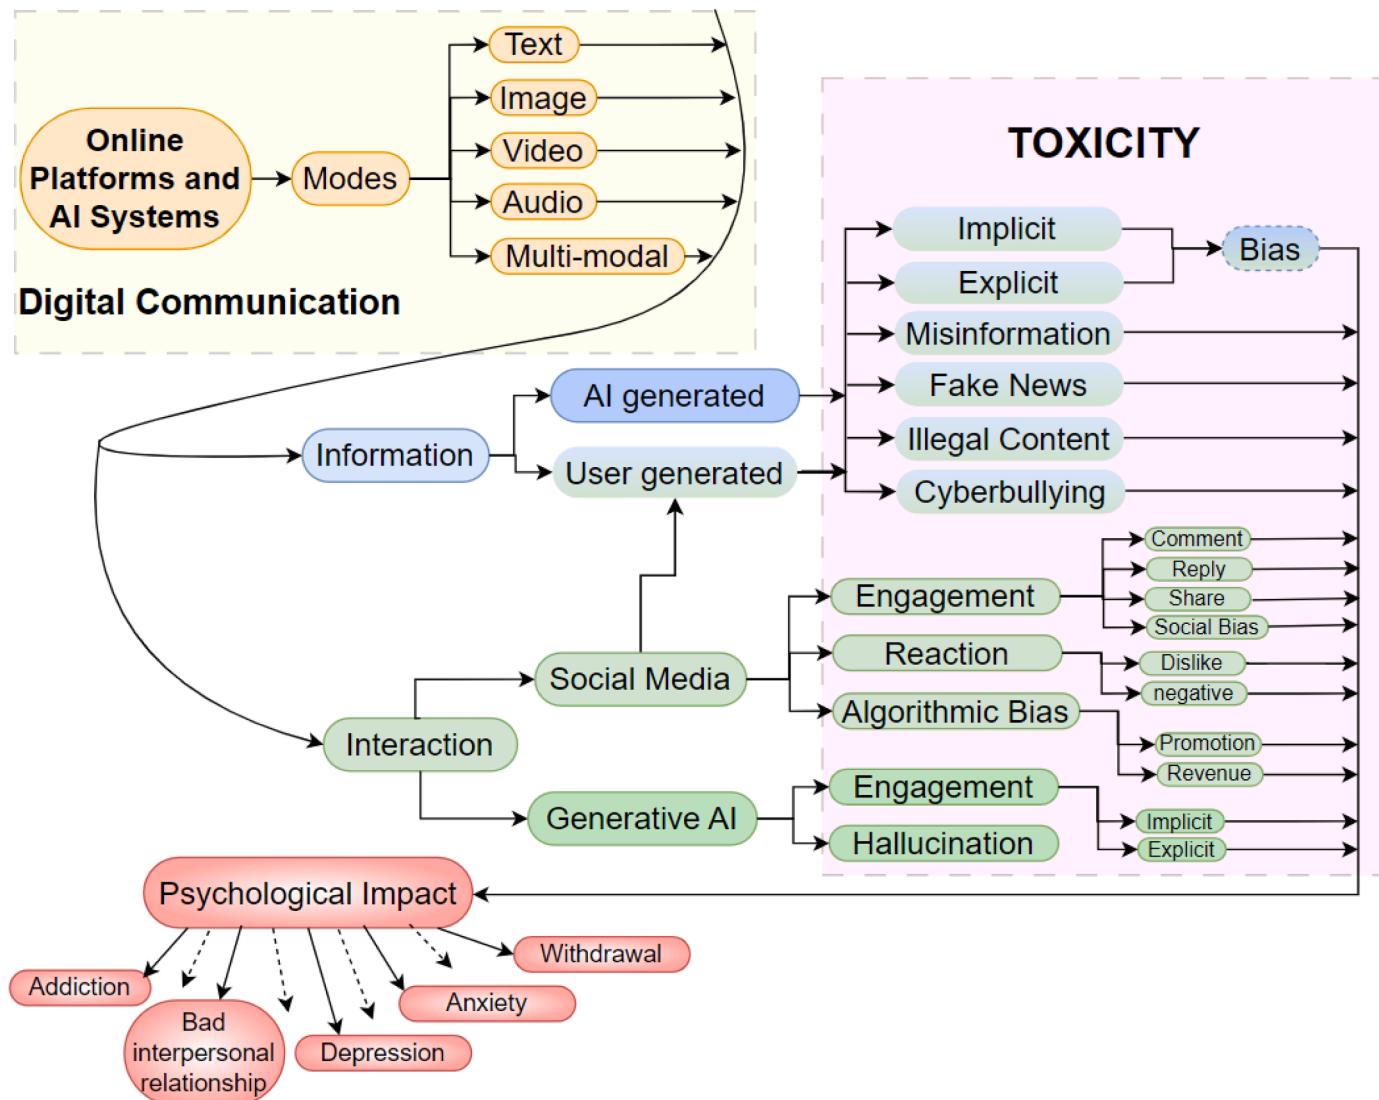


Fig. 3. A Taxonomy of Toxic Content based on origination of digital communication in online platforms and AI systems via text, image, audio, video, and/or multi-modal channel. Each mode of communication hosts toxicity divided mainly into information-based and interaction-based systems. Red color signifies the impact of toxicity, green is toxicity in interaction-based platforms, Blue represents other online platforms, and blue-green color gradient in both social media and other online platforms. The pink colored box encompasses the toxicities requiring the detection and mitigation effort.

(Li et al., 2025b). Explicit or violent images are relatively easy to detect. However, more indirect forms of image toxicity, such as hate symbols and deepfakes, are difficult to recognize (Aprin, 2024). Detecting these forms of toxic image content requires advanced mechanisms. Image toxicity is particularly harmful due to its viral nature.

Manipulated images, offensive visual content, and graphic violence can spread rapidly across online platforms (Duncombe, 2020). Such

widespread dissemination can have significant consequences in the real world, affecting social behavior, politics, and legal outcomes. Misleading toxic images can provoke violence and fuel political and social conflicts (Gamson et al., 1992). Toxic images generated by AI, such as fake content, create a more significant challenge. Advances in generative AI enable the creation of hyper-realistic images, which can be used for offensive purposes, including generation of misinformation (Wang, 2023).

Such content can seriously harm individuals and institutions. The legal and ethical implications of AI-generated imagery remain a topic of ongoing debate worldwide (Florindi et al., 2024).

4.1.3. Audio-based toxicity

Audio-based toxicity is becoming widespread with the rise of voice assistants, social audio platforms, and AI-generated speech technologies (Yu et al., 2023). Unlike text and images, audio-based toxicity presents a significant challenge due to its temporal nature (Yousefi & Emmanouilidou, 2021). The harm caused by toxic audio content often arises from the emotional intensity conveyed through spoken words (Soni & Singh, 2018). Elements like tone, pitch, and emphasis can significantly shape the meaning of a statement. As a result, spoken or audio toxicity can sometimes be more damaging than written words. Detecting and moderating audio toxicity is difficult due to linguistic and cultural diversity. Detection and moderation become more complex when it is processed in real-time. Slurs, coded language, and harmful audio toxicities may not always appear explicitly offensive. However, their impact can vary according to the speaker's intent, tone, and contextual nuances.

Another emerging challenge is the use of AI-generated voices for toxicity (Desai et al., 2024). Deepfake audio can realistically mimic real people, enabling impersonation, fraud, and targeted harassment. This technology has been exploited to create fake audio clips of political figures, spread misinformation, and manipulate public perception (Desai et al., 2024)). The consequences of synthetic audio manipulation extend across politics, journalism, and cybersecurity. Real-time audio toxicity, such as harassment in various communities in online communications, has been especially difficult to address (Blackwell et al., 2019). Toxic audio messages delivered in situations when one is engaged in activities such as driving can also be life-threatening. Live voice interactions are challenging to monitor and moderate effectively. In addition, victims often struggle to report toxic behavior, as conversations disappear once they end, leaving no evidence (Saarinen, 2017).

4.1.4. Video-based toxicity

Toxicity delivered via video is multidimensional combining text, audio, and visual elements. As a result, it is very difficult for online platforms and AI systems to recognize and mitigate. Toxic videos can contain explicit content, hate speech, misinformation, or violent material (Arora et al., 2023). Due to its immersive nature, video content can be more powerful and influential than other modalities (Tuong et al., 2014). This increases the impact of toxic video content, making moderation even more challenging (Gongane et al., 2022). The main concern with video toxicity is its ability to manipulate emotions and perceptions on a large scale. Videos are highly engaging and easily shareable, often reaching a broad audience before moderation (AlShami et al., 2024). Toxic content in videos has been linked to harassment campaigns and the spread of conspiratorial thinking. This makes video-based toxicity very dangerous and difficult to control (Maity et al., 2024a). The rise of AI-generated deepfake videos has made it easy to incorporate harmful content in seemingly realistic videos. Deepfakes can be highly realistic synthetic videos that can misrepresent individuals and fake events (Hancock & Bailenson, 2021). Deepfake-powered video toxicity presents serious threats to political stability, personal integrity, and legal systems. As deepfakes become more sophisticated, detecting and preventing video-based toxicity becomes a significant challenge (Nasri et al., 2018).

4.1.5. Multi-modal based toxicity

Multi-modal toxicity involves blending of text, image, audio, and video to create a more complex and often more dangerous form of toxic content (Maity et al., 2024b). It leverages multiple forms of inputs, making toxic and harmful content more engaging and impactful. The rise of toxic memes, which blend text and imagery to create offensive or toxic content is a perfect example. Memes spread toxicity effectively using humor and sharing to broadcast toxic content (Kiela et al., 2020). Due to

the interplay of different content types, multi-modal toxicity is challenging to detect and moderate (Liu et al., 2024). Traditional AI moderation tools struggle to analyze cross-modal interactions and toxicity (Nedungadi et al., 2025).

4.2. Information-based toxicity

Online platforms facilitate businesses, education, gaming, banking and financial services, healthcare, transportation, news media, legal, public services, and many more. The content on these platforms is either user-generated or AI-generated. It includes articles, reports, notifications, reviews, forms, chatbots, and many other formats.

When the content is generated by the subscriber or the user of a particular platform, it is referred to as user-generated content (Wyrwoll, 2014). Regardless of mode, such content in any mode is generated with or without the help of AI tools or by hand. The content that uses Generative AI models (Cao et al., 2025) is usually referred to as AI-generated content, discussed in detail in Section 4.4. Such content has the potential to host at least one or all of the forms of toxicity. Both user-generated and AI-generated content may host various fine-grained toxicities, explained in detail in Section 5.

4.3. Interaction-based toxicity

Various forms of digital communication like emails, messenger, and social media permeate our lives, providing information pertaining to news, education, business, healthcare, social engagements, government engagements, and so on. These enhance the interaction among people and entities. On various occasions, they host toxicity impacting social or physical well-being.

4.3.1. Social media platforms

Social Media Platforms (Dhingra & Mudgal, 2019), as the name suggests, are meant to bring people together. People share opinions, moments, pictures, videos, and stories informally. Over the last 2 decades, SMPs have grown to be a multi-billion-dollar business. SMPs have added various features like followers, reactions, comments, replies, streaming services, advertisements, feeds, and political and social influencers, to mention the least. The primary source of revenue for SMPs is user engagement, i.e., how long users are glued to it. The algorithms implemented track users' activities. Based on users' reactions, a platform keeps flooding their accounts with relevant feeds, which may impact users psychologically. Arguably, users get addicted to SMP and disconnected from mainstream communications. The 2020 Netflix Documentary *The Social Dilemma*, directed by Jeff Orlowski, explains this phenomenon, in particular how it impacts youth and adults, and how it moves people to extremism without realization. The main attributes of social media that engender toxicity in SMP are the following (Khapre, 2025).

- **Engagement:** Users engage in an SMP either by post or feed. They can comment, reply, and share posts or feeds. These acts, in turn, may exhibit their own biases (Cheng & Li, 2024; Shahbaznezhad et al., 2021)
- **Reactions:** Users react to the posts and feeds using emojis (Cheng & Li, 2024; Eberl et al., 2020; Lin & Utz, 2015). Negative reactions may introduce toxicity, depending on the context. Another potent form of reaction is to un-follow or un-friend a user. A reduction in the number of followers or friends may impact upon an individual's self-esteem or affect in other ways. Negative reactions may signal criticism, disapproval, or denial of relationships among the users, potentially impacting their well being.
- **Algorithmic Bias:** (Qiu, 2024; Thomas & Bhat, 2022) Two types of algorithms employed in an SMP promote user engagement in the addictive content and for revenue generation from advertisement feeds. The algorithms prioritize emotionally provocative or controversial

content to maximize user interaction (Vosoughi et al., 2018; Zollo & Quattrociocchi, 2018).

SMPs are afflicted with fine-grained toxicities discussed in Section 5. Although knowingly or unknowingly, an SMP may host a lot of these toxicities, with concomitant negative impacts on society and human psychology (Ali et al., 2025; Ali, 2024; Antypas & Camacho-Collados, 2023; Aprin, 2024; Beknazaryuzbashev et al., 2025; Dangswang & Nuchitprasitchai, 2024; Duncombe, 2020; Fan et al., 2024a; Gao et al., 2020; Gongane et al., 2022; Hall, 2025; Khapre, 2025; Ognibene et al., 2023; Przybyla et al., 2024; Qiu, 2024; Ray et al., 2024; Roy & Kumar, 2024; Shahbaznezhad et al., 2021; Sharma et al., 2023; Shu et al., 2020; Throuvala et al., 2021; Vaidya et al., 2020; Wyrwoll, 2014; Zhang, 2024). SMPs are popular in society for a lot of their positive contributions.

4.4. Generative AI

Generative AI has revolutionized digital content creation, enabling the rapid synthesis of text (Shakil et al., 2024a), images (Archana Balkrishna, 2024), audio (Mitra & Zualkernan, 2025), and video (Zhou et al., 2024). Large Language Models (LLMs) such as OpenAI's GPT (Shakil et al., 2024b), Google's Gemini (Imran & Almusharraf, 2024), and Meta's LLaMA (Li et al., 2025a) generate human-like text by leveraging vast datasets. However, these models are also prone to producing toxic content, including hate speech, misinformation, and offensive material, either unintentionally or as a byproduct of biased training data (Gehman et al., 2020; Vidgen et al., 2021a). The generative capability of AI introduces significant challenges in online toxicity management, as it can reinforce harmful biases, amplify misinformation, and propagate toxic content at scale (Mathew et al., 2021; Sap et al., 2020a).

Generative AI interacts with online toxicity in several ways. First, it learns from unfiltered internet data, which often contains implicit (Wen et al., 2023) and explicit toxicity (Gunturi et al., 2023), leading to biased or harmful outputs (Wei et al., 2025). For example, LLMs trained on datasets with racial or gender biases may inadvertently perpetuate these biases in their outputs, even when explicitly instructed to avoid toxic content (Bender et al., 2021). Second, generative AI enhances engagement by generating content that aligns with user preferences, which can sometimes exacerbate toxic behaviors (Garg et al., 2023). It is particularly true in SMP, where algorithms prioritize emotionally provocative or controversial content to maximize user interaction (Vosoughi et al., 2018; Zollo & Quattrociocchi, 2018). Third, generative AI hallucinates, i.e., it produces misleading or fabricated information, contributing to misinformation and digital harm (Huang et al., 2025). Hallucinations can lead to the spread of false narratives, which are often more engaging and emotionally charged than factual content, making them more likely to go viral (Shakil et al., 2024c).

Addressing these issues requires a nuanced understanding of how AI-generated toxicity manifests and impacts online discourse. For instance, the ethical implications of generative AI in content creation are profound, including questions about accountability, transparency, and the potential for misuse in spreading disinformation or hate speech (Dubber et al., 2020). Moreover, the challenges of detecting and mitigating toxicity in multi-modal content (e.g., toxic memes, deepfake videos) further complicate the landscape, as traditional text-based moderation tools are often insufficient to address these emerging threats (Desai et al., 2024).

4.4.1. Engagement

Engagement in digital platforms is driven by algorithms that maximize user interaction, often by prioritizing emotionally provocative or controversial content (Vosoughi et al., 2018; Zollo & Quattrociocchi, 2018). Generative AI plays a significant role in shaping these engagement patterns, as its outputs are designed to be compelling and personalized. While AI-generated content fosters meaningful interactions, it also risks reinforcing toxic discourse through both implicit and explicit engagement. This dual nature of engagement poses significant challenges

for online platforms, as it can amplify harmful behaviors while simultaneously driving user retention and revenue (Selvakumar et al., 2025).

Implicit engagement refers to AI's ability to subtly shape user interactions without overt toxicity. Generative AI models reinforce biases by personalizing content recommendations based on user behavior, which can lead to echo chambers and filter bubbles (Jawad et al., 2024; Siapera, 2022). These echo chambers intensify ideological divides, making users more resistant to opposing views and indirectly perpetuating toxicity (Mathew et al., 2021). For example, AI systems that recommend content based on user preferences may inadvertently reinforce harmful stereotypes or prejudiced language structures, even when the content itself is not explicitly toxic (Garg et al., 2023; Zellers et al., 2019).

Explicit engagement occurs when AI-generated content directly propagates toxic discourse, often through misinformation, hate speech, or inflammatory language (Fitzgerald et al., 2020; Gambín et al., 2024; Mersha et al., 2024c). AI systems trained on unfiltered datasets are prone to generating harmful content that fuels online harassment, cyberbullying, and hate speech (Ali, 2024). For instance, toxic memes, deepfake videos, and offensive chatbot responses illustrate how generative AI can contribute to explicit online toxicity (Desai et al., 2024). These forms of toxicity are particularly dangerous because they are designed to provoke strong emotional reactions, making them highly shareable and difficult to moderate (Lehtimäki, 2024).

The engagement-driven nature of online platforms exacerbates explicit toxicity by rewarding high-interaction content. Studies indicate that AI-generated misinformation spreads significantly faster than factual content due to its provocative and emotionally charged nature (Vosoughi et al., 2018). This presents a critical challenge for content moderation, as automated AI models often fail to distinguish between engaging content and harmful narratives (Siapera, 2022). For example, AI systems may prioritize content that generates high engagement, even if it contains hate speech or misinformation, due to the lack of robust mechanisms for detecting and filtering such content in real time (Wang et al., 2023).

Addressing explicit AI-driven toxicity requires improved content moderation strategies, enhanced model transparency, and the integration of bias-mitigation techniques such as adversarial training and reinforcement learning with human feedback (Liu, 2023). Additionally, platforms must adopt a multi-stakeholder approach, involving policymakers, researchers, and civil society organizations, to develop ethical guidelines and regulatory frameworks that balance user engagement with online safety (Dubber et al., 2020).

4.4.2. Hallucination

Hallucination in generative AI refers to the phenomenon where models generate incorrect, misleading, or fabricated content, often presented as factual information (Ji et al., 2023; Shakil et al., 2024c). This issue is particularly problematic in toxicity management, as hallucinated outputs can contribute to the spread of misinformation, hate speech, and harmful stereotypes (Tonmoy et al., 2024). Hallucinations arise from several factors, including gaps in training data, overgeneralization, and the probabilistic nature of language model generation (Guerreiro et al., 2023). When AI systems lack sufficient context or rely on incomplete datasets, they generate content that may appear plausible but is factually incorrect (Huang et al., 2024).

The impact of AI hallucinations extends beyond misinformation, influencing public opinion and exacerbating social biases. For instance, AI-generated content can reflect and amplify racial, gender, and political biases present in the training corpus, perpetuating toxic discourse (Fang et al., 2024). In the context of large-scale SMP, hallucinated toxicity can rapidly spread, leading to real-world consequences such as increased polarization, targeted harassment, and even violence (Latif, 2025). For example, AI-generated fake news or deepfake videos can manipulate public perception, erode trust in institutions, and destabilize democratic processes (Wang, 2023).

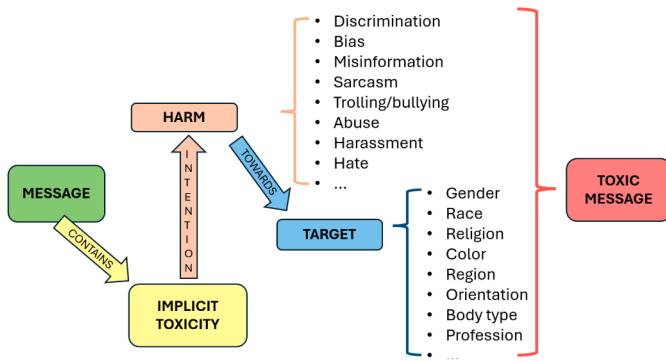


Fig. 4. Implicit Toxicity: The message seems benign or normal, but based on the context, the reader or viewer, or the environment, it intends to harm a particular target. Such messages have implicit toxicity.

Efforts to mitigate hallucinations in generative AI focus on refining training methodologies and improving model oversight. Techniques such as knowledge-grounded generation, fact-checking mechanisms, and adversarial training have shown promise in reducing AI hallucinations (Yu et al., 2024). Reinforcement learning with human feedback further enhances model reliability by aligning outputs with human values and factual accuracy (Ouyang et al., 2022). However, balancing creativity with factual correctness remains a challenge, particularly for models generating diverse and context-dependent content. For instance, while reinforcement learning with human feedback can reduce hallucinations, it may also limit the model's ability to generate creative or contextually rich responses, highlighting the need for a nuanced approach to hallucination mitigation (Yaprak, 2024).

Moreover, addressing hallucinations requires a multi-faceted approach that combines technical solutions with ethical and regulatory frameworks. For example, platforms must implement robust content moderation systems to detect and remove hallucinated content, while policymakers must establish guidelines to hold AI developers accountable for the societal impact of their systems (Bertонcini & Serafim, 2023). Collaborative efforts between researchers, developers, and policymakers are essential to ensure that generative AI systems are both innovative and responsible (Kim, 2024).

5. Fine-grained toxicities

The main types of toxicities are implicit toxicity, as depicted in Fig. 4, explicit toxicity, as depicted in Fig. 5, and fallacy. Explicit toxicity is also called fallacy by some. Fallacy (Helwe et al., 2024) refers to misinformation, disinformation, and fake news. Often, toxicity has intentions of causing harm and is targeted towards a particular target group (Mersha et al., 2024c). There is another important type of toxicity *without a target* which could be both implicit or explicit, containing profanity, negative, derogatory, disrespectful, and offensive language (Pachinger et al., 2023). Before 2021, this type was the focus of study and research as detailed by van Aken et al. (2018). Despot et al. highlighted that while analyzing implicit toxicity with figures of speech like metaphor, explicit toxic content without a target tends to be ignored (Despot et al., 2023; Štrkalj Despot, 2023). They have presented a typology of implicit toxicity such that explicit toxicity gets detected and not ignored. It includes caps lock and repetition as aggressive speech; irony, metaphor, circumlocution, hyperbole, rhetorical questions, (re)interpretation, euphemism, simile, contrast, and name-calling are fine-grained in discrediting, insulting, dehumanization, derogation, and discriminatory speech.

Pachinger et al. (2023) have analyzed various research and analyzed a granular enumeration of harms caused by toxicity, including threats, rudeness, mockery, name-calling, obscenity, pejorative content, aggres-

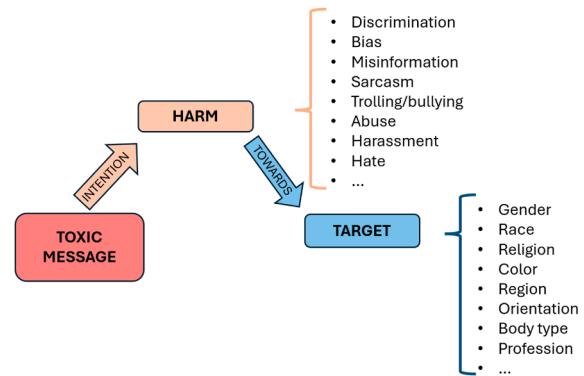


Fig. 5. Explicit Toxicity: The explicit toxic message intends to harm a particular target, which can be easily understood as toxic. Its implied meaning and actual content have the same meaning. Some explicit toxic messages that do not have a target are prevalent in the form of profanity, sarcasm, and other forms of offensive language.

siveness, profanity, insult, discrimination, discrediting, irony, sarcasm, derogation, hostility, hatefulness, vulgarity, and criminality.

Wang et al. (2024d) have come up with a taxonomy of LLM risks containing three levels with an extensive fine-grain at the third level. They have highlighted that the harms of toxicity may also be a result of illegal trade, and illegal activities such as abuse, cruelty, exploitation, intimidation, cyberbullying, harassment, baiting, stalking, and defamation of vulnerable individuals. Harm could also arise from the endorsement of unsafe items or practices, and the spread of information hazards such as personally identifiable information pertaining to life, health, finance, education, biometrics of an individual, confidential, classified, legal, and sensitive in nature. Such information may cause cybersecurity vulnerabilities, theft of technical data and internal communication, misinformation, disinformation, fake news, rumors, propaganda, misinterpretation, and unreliable experts giving legal, medical, and financial advice. Fine-grained target groups include those targeted for age, disability, gender, sexual orientation, race, ethnicity, religion, region, color, body type, language, class, nationality, profession, and others. Baczkowska et al. (2025) presented an implicitness model comprising indirectness in rhetorical questions, figures of speech like simile, metaphor, and irony (sarcasm), exaggerations in forms of over- or understatement; resulting in hostile, hateful, racist, homophobic, discredit, threat, misogynic and vulgarism. Guest et al. (2021) presented a detailed misogyny taxonomy based on online misogyny literature, mainly on social media. They have classified online misogyny into 4 major classes of derogatory, pejorative, treatment, and personal attacks; the pejorative and personal attacks in their taxonomy seem explicit toxicity; the other two have further fine-grained toxicities, like derogatory, which have intellectual and moral inferiority, sexual and physical limitations, and others. Treatment has threatening language involving physical and sexual violence, invasion of privacy, and disrespectful actions involving controlling, manipulation, seduction, conquest, and others.

Korre et al. (2024) have introduced a compelling perspective of the legal framework in understanding hate speech, emphasizing its prosecutable nature and significant variation across countries. They have highlighted three legal approaches: content-, intent-, and harm-based, having some target being impacted. The content-based approach has direct content to offend certain people in society; intent-based is communication to its audience to invoke hatred towards minorities or other hateful intentions, evidently not legitimate; and harm-based induces lower self-esteem, physiological, emotional, and psychological conditions, socio-economic distress, and societal withdrawal symptoms to its victims.

6. Convergence with human psychology

In human psychology, toxicity is usually referred to as *incivility*, *offensive behavior* (Lewandowska-Tomaszczyk et al., 2023), *flaming*, and digital communication as *computer-mediated communication* (Bansal et al., 2012; Lee, 2005). Flaming (aka Toxicity) can further be considered as direct or indirect, and straight or satirical (Bansal et al., 2012). Ali et al. (2025) demonstrated the relation of social media platform addiction with impacts like depression, anxiety, bad interpersonal relationships, hopelessness, guilt, sadness, lack of enthusiasm, sleep, and compulsive disorders, and called for mitigation strategies in digital platforms. See Section 2.2 for mental health impact. Papageorgiou et al. (2024) highlighted the need for safeguarding information integrity to ensure societal trust in fallacy detection. Pennycook and Rand (2021) investigated the psychology behind fallacy and why it spreads faster than real news.

6.1. Biases in LLMs

The bias embedded during training significantly affects the development of fair AI systems. A growing number of studies (Rao et al., 2023) indicate that LLMs exhibit reasoning and self-improvement abilities that resemble human cognitive processes, raising the possibility that these models may develop virtual personalities and psychological characteristics (Pan & Zeng, 2023). In particular, gender- and race-based discriminatory content is often introduced during training (Cabrera et al., 2023), which can embed harmful biases into the model's responses and behavior. These biases can manifest in stereotyping, preference for dominant social norms, or assumptions based on patterns that do not reflect the complexity of real-world experiences (Ferrara, 2023). The high parameterization of modern LLMs, while enhancing capacity, introduces multidimensional interactions that are difficult to predict and manage. Bai et al. (2024) reported that the most substantial biases are found in categories related to race, followed by gender, health, and religion. One illustrative example of gender bias is the reinforcement of exclusionary language norms, such as referring to "both genders," which implicitly excludes non-binary identities (Bender et al., 2021). Despite increased attention to these challenges, much of the current bias research in LLMs falls short in clearly identifying who is harmed, why the behavior is harmful, and how it reflects and reinforces social hierarchies (Blodgett et al., 2020). While all LLMs exhibit some level of bias, there are significant variations among them. Models with more parameters, such as GPT-4, GPT-3.5-Turbo, Claude3-Opus, and Claude3-Sonnet, as well as LLaMA2Chat-70B and 13B, tend to display more pronounced implicit biases. In contrast, smaller models like LLaMA2Chat-7B and Alpaca7 B show substantially lower bias levels (Kumar et al., 2024).

6.2. Discrimination in LLMs

Discrimination in LLMs can manifest in both direct and indirect forms. Direct discrimination involves overtly unequal treatment based on group membership. For instance, when LLM-powered resume-screening tools reinforce existing hiring inequities (Ferrara, 2023). Indirect discrimination occurs when model outputs, while seemingly neutral, rely on proxy variables that disproportionately disadvantage certain groups. This can be seen in LLM-assisted healthcare systems, where demographic proxies contribute to inequities in patient care (Ferrara, 2023). Many approaches, for example, rely on the implicit assumption that model outputs should be independent of social group identifiers, yet they often fail to explicitly articulate the underlying social or ethical values that support this assumption. Furthermore, inconsistency in how bias is defined and a lack of meaningful engagement with structural power dynamics hinder the creation of robust and equitable solutions (Blodgett et al., 2020).

6.3. Personality Types and Toxicity

Wang et al. (2025) showed in their experiments using the HEXACO framework on three LLMs, that the low levels of sensitivity of these six personality traits of *Honesty-Humility*, *Emotionality*, *eXtraversion*, *Agreeableness*, *Conscientiousness*, and *Openness to Experience*, produce higher levels of toxicity and bias in their responses. Serapio-García et al. (2023) through experiments have shown that the LLMs have a synthetic personality and can be fine-tuned to mimic desired personality traits. Pan and Zeng (2023) used another psychological test MBTI (Myers-Briggs Type Indicator), with dichotomies of extraversion, introversion, sensing, thinking, intuition, feeling, judging, and perceiving. They have shown that LLMs innately possess some personality traits. In a similar study, Jiang et al. (2024) used *Big Five Inventory* test on LLMs with personality profiles, and their self-reported evaluations were consistent with the LLMs' personality profiles. All these psychological evaluations of LLMs' personality traits provide a mitigation strategy that can be built into the LLMs by fine-tuning them with higher sensitivity, reducing bias and toxicity in LLM-generated content. At the same time, one should be careful during LLMs training, so that it avoids less sensitive personality traits which might otherwise make it generate toxic contents.

6.4. Issue-based discussions and political polarization

Although digital platforms allow open discourse among diverse populations, they are also breeding grounds for toxicity in politically charged conversations. Gao et al. (2024) analyzed over 40 million Reddit comments and found that incivility, including vulgarity, asperion, and name-calling, is more prevalent in politically biased forums. People are more likely to respond uncivilly to those with opposing views. Moreover, toxic comment threads tend to attract more participation but also escalate in hostility. Exposure to personal attacks often leads users to disengage from political spaces. If LLMs are trained on such content, they may inadvertently learn and replicate these patterns, reflecting partisan toxicity and asymmetrical ideological bias. For instance, certain forms of incivility may be overrepresented by specific political groups, which in turn can influence how models respond to politically sensitive prompts. Gao et al. (2024) also found that both left-leaning and right-leaning individuals became more uncivil when engaging with opposing views, though right-leaning users exhibited higher levels of all incivility subtypes except third-party attacks. The most prominent difference-interpersonal name-calling-was 177 percent higher among right-leaning users. In a related study, Rozado (Rozado, 2023) administered 15 political orientation assessments to ChatGPT and reported a tendency for the model to favor left-leaning perspectives. However, this finding should be interpreted cautiously due to the limited number of tests and the exploratory nature of the analysis.

In conclusion, LLMs not only mirror the toxic, biased, and polarized content present in online human communication but also risk reinforcing these tendencies when not properly constrained. Understanding these interactions is crucial for developing more equitable, responsible, and psychologically aware AI systems.

7. Toxicity mitigation

Toxicity mitigation mechanism requires a dataset, detections, and detoxification mechanisms (framework, model, or techniques). All the surveyed literature is summarized in Fig. 6 based on broadly classified domains as Web or SMP, Fake News or Misinformation, and LLM. Most of the models developed or proposed can be tested in at least one of these domains. Web or SMP toxicity requires a dataset of text sequences taken from SMP, gaming, healthcare, education, or other web-based platforms for detection models, and counterspeech for moderation or detoxification. Fake news or misinformation requires a fallacy-based strategic detection model, where detoxification is quite challenging. The fallacy datasets are created using Gossipcop and Politifact web databases. And

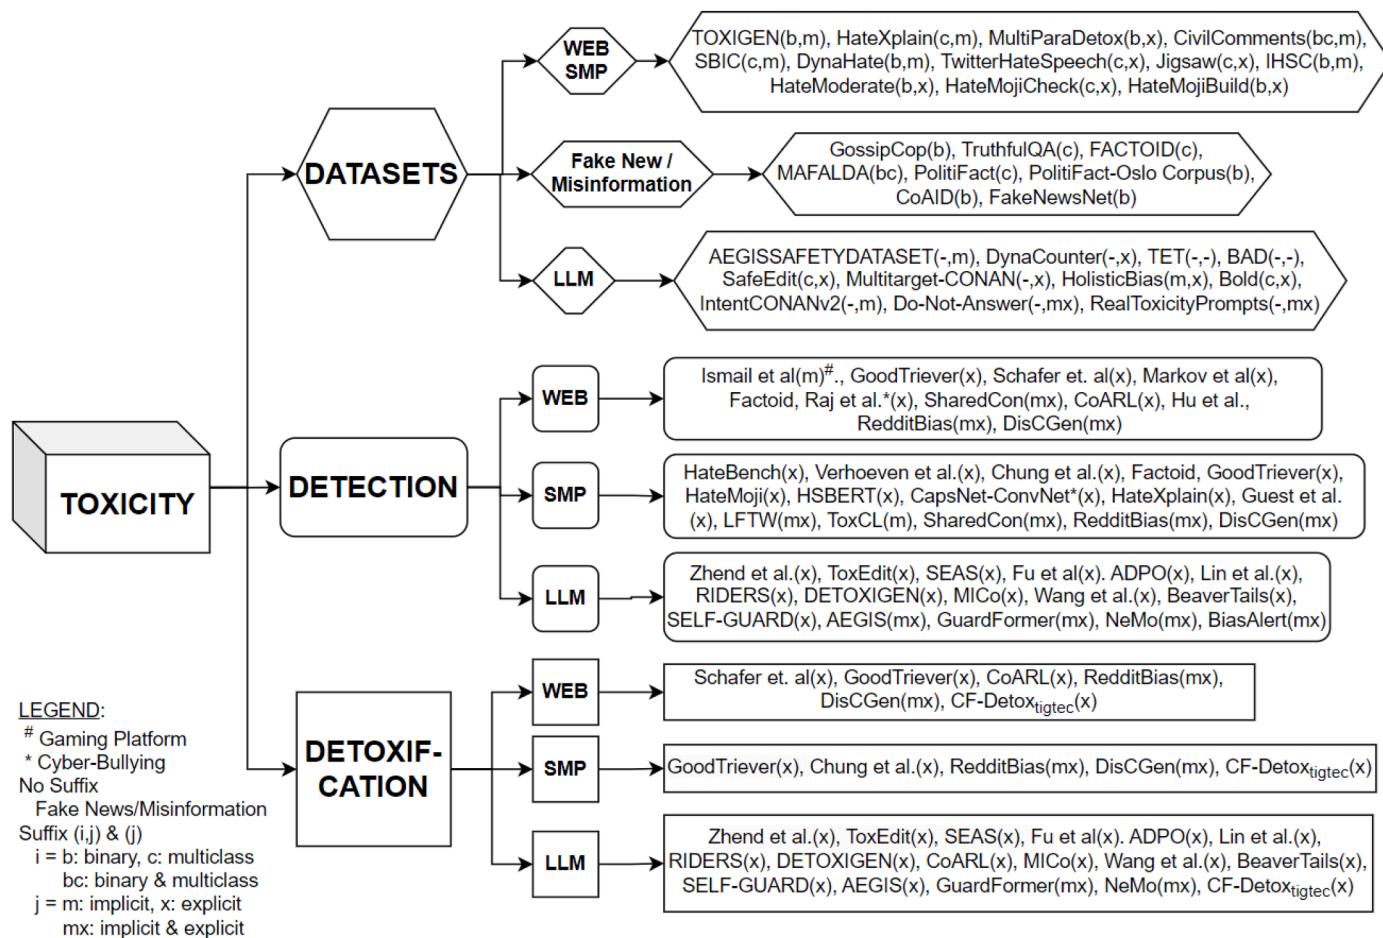


Fig. 6. Summary of Toxicity Mitigation mechanism, which includes dataset, detection, and/or detoxification mechanisms. Their references are listed in the Tables 2 and 3.

LLMs are the model itself and require design modification within the model to allow detection and detoxification. Toxicity datasets for LLMs are datasets of question-answer pairs. In Fig. 6, each listed dataset is suffixed with two parameters in parentheses, 'i' and 'j'. Each listed detection and detoxification method has a suffix in parentheses 'j'. Where 'i' is either binary 'b' or multi-class 'c' or both 'bc'. The recent momentum in research papers published by year is shown in Fig. 9. The toxicity mitigation research has shown momentum in 2024 using LLMs as detailed in Table 3.

7.1. Toxicity datasets

Toxicity datasets used in surveyed research are collated in Table 2 and is summarized in Fig. 7. The datasets used vary by application domains and research objectives. The datasets are usually curated for specific application domains. Thus, some are binary datasets (Cui & Lee, 2020; Gossipcop, 2025; Hartvigsen et al., 2022; Pöldvere et al., 2023; Vidgen et al., 2021b; Zheng et al., 2024) whereas others are multiclass. Many datasets are multi-label where the examples have primary hate classifications like *hate* and *not hate*, secondary labels for the target class, and additional fine-grained labels (Barikeri et al., 2021; Dhamala et al., 2021; ElSherief et al., 2021; Garg & Sharma, 2020; Ghosh et al., 2024; Helwe et al., 2024; Ji et al., 2024; Kirk et al., 2022; Lin et al., 2022; Mathew et al., 2021; Sakketou et al., 2022; Sap et al., 2020b; Shaikh et al., 2023; Shen et al., 2025; Smith et al., 2022; van Aken et al., 2018; Verhoeven et al., 2024; Wang et al., 2024a,d). AEGISSAFETYDATASET (Ghosh et al., 2024), BeaverTails (Ji et al., 2024), Do-Not-Answer (Wang et al., 2024d), HolisticBias (Smith et al., 2022), and HarmfulQ

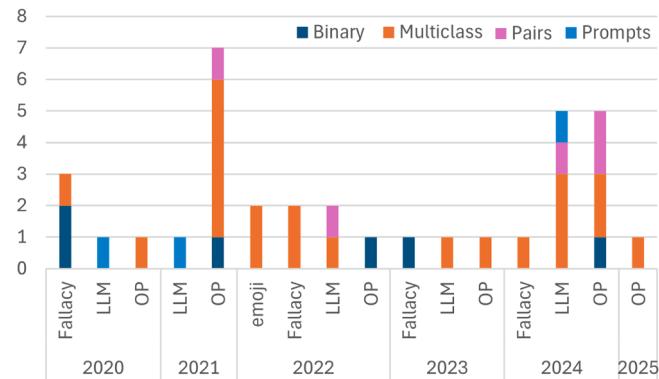


Fig. 7. Summary of Toxicity Datasets creation over the last five years for Fallacy, LLMs, Online Platforms (OP), and Emoji. The Y-axis is the number of datasets.

(Shaikh et al., 2023) are datasets used in toxicity detection and detoxification in LLMs. These are multiclass datasets curated for LLM safety. Thoroughly Engineered Toxicity (TET) (Luong et al., 2024), MultiParaDetox (Dementieva et al., 2024), ParaDetox (Logacheva et al., 2022), RealToxicityPrompts (Gehman et al., 2020), and RedditBias (Barikeri et al., 2021) are toxic and paraphrase pairs datasets for LLMs safety. TRam1/safe-guard-prompt-injection,¹ deepset/

¹ <https://huggingface.co/datasets/xTRam1/safe-guard-prompt-injection>

Table 2

Toxicity Datasets; [* Multi-class and Multi-label for each class and target group, # Binary with Multi-label for Hate class and Target group, @ Binary with Two Labels each, - Neither binary nor multiclass (Question-answer type dataset)].

Dataset	Year	Multiclass /Binary /Both	Number of classes	Size	Purpose
HateBenchSet (Shen et al., 2025)	2025	Multiclass	34	7838	benchmarking detectors for LLMs
AEGISSAFETYDATASET (Ghosh et al., 2024)	2024	Multiclass	15	26,000	-generated content Adaptive content safety in LLM
BeaverTails (Ji et al., 2024)	2024	Multiclass	14	360K	QA Pairs for Safety of LLMs
TwitterHateSpeech (Verhoeven et al., 2024)	2024	Multiclass	3	24.8K	Explicit hate detection
DynaCounter (Chung & Bright, 2024)	2024	-	-	1911	Abuse and reply pair with target group label specific to Football Premier league
IntentCONANv2 (Hengle et al., 2024)	2024	-	-	3488	4 counterspeech for each hate speech
Hatemoderate (Zheng et al., 2024)	2024	Binary	2	7.6K	instance of positive, informative, questioning and denouncing Hate speech policy or rules compliance
Do-Not-Answer (Wang et al., 2024d)	2024	Multiclass	6	939	Suitable for LLMs safety (referred as
Multilevel Annotated Fallacy (MAFALDA) (Helwe et al., 2024)	2024	Both	23	9745	toxicity in this paper) training 2 classes at Level 0, 3 classes at Level 1 and 23 classes at level 3
Thoroughly Engineered Toxicity (TET) (Luong et al., 2024)	2024	-	-	2546	Dataset of prompts to break the protective layers of LLMs
SafeEdit (Wang et al., 2024a)	2024	Multiclass	9	8100	Detection and detoxification in web or
MultiParaDetox (Dementieva et al., 2024)	2024	-	-	16K	SMP SafeEdit Extending ParaDetox, supports
CivilComments (Duchêne et al., 2023)	2023	Both	25	1.8M	Russian, Ukrainian and Spanish Unintended bias Classification
PolitiFact-Oslo Corpus (Pöldvere et al., 2023)	2023	Binary	2	2745	Fake news analysis and detection
HarmfulQ (Shaikh et al., 2023)	2023	Multiclass	6	200	Prompts, zero-shot chain of thoughts
TOXIGEN (Hartvigsen et al., 2022)	2022	Binary	2	274,186	increases bias and toxicity with model size Detection in online platforms
TruthfulQA (Lin et al., 2022)	2022	Multiclass	4	817	817 questions that span 38 categories
FACTuality & pOlitical bias (Dataset (FACTOID) (Sakketou et al., 2022))	2022	Multiclass	65	3.3M	4150 news-spreading users with
HateMojiCheck (Kirk et al., 2022)	2022	Multiclass	7	3930	3.3M Reddit posts Detection of emoji based Toxicity
HateMojiBuild (Kirk et al., 2022)	2022	Both	54	5912	Hate/non-hate; 54 target groups
HolisticBias (Smith et al., 2022)	2022	Multiclass	600	45K	Prompts dataset expanded to 769 descriptors
ParaDetox (Logacheva et al., 2022)	2022	-	-	10K	English toxic & paraphrase pairs in LLMs
Implicit Hate Speech Corpus (ElSherief et al., 2021)	2021	Both	7	22K	Implicit hate detection
HateXplain* (Mathew et al., 2021)	2021	Multiclass	3	20K	Implicit hate detection with 10
DynaHate# (Vidgen et al., 2021b)	2021	Binary	2	40K	target group labels Implicit hate detection; Hate class with
Bias in Open-Ended Language (Generation Dataset (BOLD) (Dhamala et al., 2021))	2021	Multiclass	5	23,679	6 labels and 29 target group labels 5 domains with total of 43 target
Multitarget-CONAN (Fanton et al., 2021)	2021	-	-	5000	groups Hate speech and counter Narrative pairs
Bot Adversarial Dialogue (BAD) (Xu et al. (2021))	2021	-	-	70K	ChatBot safety 5K dialogues and 70K utterances
RedditBias (Barikeri et al. (2021))	2021	Multiclass	4	11,873	Bias on 4 dimensions: religion, race,
Online Misogyny EACL2021 (Guest et al. (2021))	2021	Multiclass	5	6.567K	gender and orientation Online Misogyny collected from Reddit
Social Bias Inference	2020	Multiclass	7	150K	posts, annotated by trained annotators Implicit hate with 7 variables; up to 3
Corpus* (Sap et al., 2020b)					multi-labels; target group having 16 labels

Table 2
Continued.

Dataset	Year	Multiclass /Binary /Both	Number of classes	Size	Purpose
CoAID @ (Cui & Lee, 2020)	2020	Binary	2	5.2K	Misinformation regarding COVID with fake or real classes
RealToxicityPrompts (Gehman et al., 2020)	2020	-	-	100K	Sentence paired with toxicity scores
PolitiFact (Garg & Sharma, 2020)	2020	Multiclass	6	21K	Dataset of political news
FakeNewsNet (Shu et al., 2020)	2020	Binary	2	2.67B	Repository of Politifact and GossipCop
Jigsaw Toxic Comment Classification Challenge (van Aken et al., 2018)	2018	Multiclass	6	159K	Curation of Wikipedia Comments for Kaggle hosted competition in 2018
GossipCop (Gossipcop, 2025)	2009	Binary	2	22.152K	Fake or Real; subset from a fact-checking website created in 2009

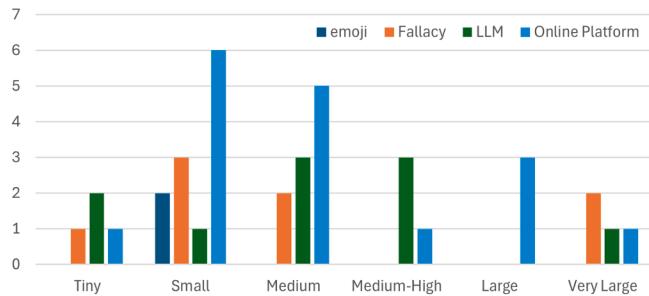


Fig. 8. Datasets size distribution: Y-axis is number of datasets, and X-axis is dataset size range where Tiny is <1000, Small is between 1000 and 9999, Medium is between 10,000 and 39999, Medium-High is between 40,000 and 129999, Large is between 130,000 to 299999, and Very Large is 300,000 and above.

prompt-injections² are Prompt Injection datasets, and lmsys/toxic-chat,³ SetFit/toxic_conversations_50k⁴ are toxicity datasets from Huggingface. TruthfulQA (Lin et al., 2022) is a dataset used to check truthful answers generated by LLM for the questions. BAD (Xu et al., 2021) is a dataset for chatbot safety. TOXIGEN (Hartvigsen et al., 2022), Implicit Hate Speech Corpus (ElSherief et al., 2021), Social Bias Inference Corpus (Sap et al., 2020b), HateXplain (Mathew et al., 2021), DynaHate (Vidgen et al., 2021b), TwitterHateSpeech (Verhoeven et al., 2024), HateMojiCheck (Kirk et al., 2022), HateMojiBuild (Kirk et al., 2022), CivilComments (Duchêne et al., 2023), and Jigsaw Toxic Comment Classification Challenge (van Aken et al., 2018) are the datasets for toxicity detection in the web/SMP context. DynaCounter (Chung & Bright, 2024), IntentCONANv2 (Hengle et al., 2024), HateModerate (Zheng et al., 2024), BOLD (Dhamala et al., 2021), SafeEdit (Wang et al., 2024a), Multitarget-CONAN (Fanton et al., 2021), and HateBenchSet (Shen et al., 2025) are datasets for toxicity detection and detoxification in web/SMP. GossipCop,⁵ CoAID (Cui & Lee, 2020), FACTOID (Sakkou et al., 2022) MAFALDA (Helwe et al., 2024), PolitiFact (Garg & Sharma, 2020) PolitiFact-Oslo Corpus (Pöldvere et al., 2023), FakeNewsNet (Shu et al., 2020) are datasets for fallacy detection. As shown in Fig. 8, most of the datasets are small and cannot be used in a real-world application. Only 11 datasets namely, RealToxicityPrompts, Bot Adversarial Dialogue, HolisticBias, DynaHate, Jigsaw Toxic Comment Classification Challenge, Social Bias Inference Corpus, TOXIGEN, FakeNewsNet, FACTuality & political bias, BeaverTails, and CivilComments, can be used for future

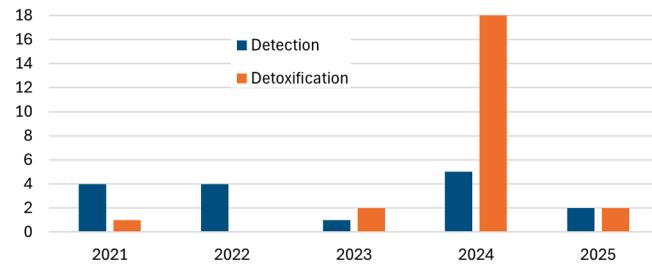


Fig. 9. Toxicity mitigation research since 2021. Y-axis is the number of research papers published.



Fig. 10. The involvement of LLMs in Toxicity Detection and Mitigation has been increasing rapidly since 2023, compared with other methods.

research, providing generalization. Out of these, only two of them, *RealToxicityPrompt* include implicit toxicity and *Beavertails* for explicit toxicity, can be used for detoxification. The remaining smaller datasets are mainly curated to prove the toxicity problems and call for further research.

7.2. Toxicity detection

Toxicity detection is a process of implementing a mechanism to identify a toxicity incident in communication in a given environment. Detection techniques may or may not be LLM-based, depending on the platforms, although the trend seems to move towards LLM-based approaches as shown in Fig. 10. In this paper, both LLM-based and non-LLM-based detection techniques are reviewed and presented in Table 3. It is evident that different LLMs have different definitions of toxicity and approaches for detection (Shen et al., 2025). Koh et al. (2024) proposed a framework for evaluating toxicity with variable definitions, emphasizing the flexibility of the toxicity metric and demonstrating the adaptability of LLMs to prompts. To train toxicity detectors, some researchers have generated datasets like TOXIGEN, AEGISAFETYDATASET, HateBenchSet, and HateXplain (refer to

² <https://huggingface.co/datasets/deeppset/prompt-injections>

³ <https://huggingface.co/datasets/lmsys/toxic-chat>

⁴ https://huggingface.co/datasets/SetFit/toxic_conversations_50k

⁵ <https://huggingface.co/datasets/osusume/Gossipcop>

Table 3

Summary of toxicity detection and mitigation research presented in this survey. NDS = Novel Dataset; NFW = Novel Framework; NMT = Novel Model or Technique; UEM = Used Established Models/Benchmarks; WEB = Content hosting websites like Wikipedia; InfoG = Info-Graphics; GP = Gaming Platforms; DET = Detection; MTG = Mitigation; FN&M* = Fake News and Misinformation.

Toxicity Research	Year	NDS	NFW	NMT	UEM	Domain	Mode	DET	MTG	Remarks
HateBench (Shen et al., 2025)	2025	Yes	Yes	No	Yes	WEB	Text	Yes	No	Explicit
Ismail et al. (2025)	2025	No	Yes	Yes	No	GP	Text	Yes	No	Implicit
Zheng et al. (2025)	2025	No	No	Yes	Yes	LLM	Text	Yes	Yes	Explicit
ToxEdit (Zhang, 2025)	2025	No	No	Yes	Yes	LLM	Text	Yes	Yes	Explicit
LATTE (Koh et al., 2024)	2024	No	Yes	No	Yes	LLM	Text	Yes	No	Both
SEAS (Diao et al., 2024)	2024	Yes	Yes	No	Yes	LLM	Text	Yes	Yes	Explicit
Goodtrifier (Ermis et al., 2024)	2024	No	Yes	No	Yes	Any	Text	Yes	Yes	Explicit
Fu et al. (2024)	2024	Yes	No	No	Yes	LLM	Text	Yes	Yes	Explicit
ADPO (Kim & Lee, 2024)	2024	No	No	Yes	Yes	LLM	Text	Yes	Yes	Explicit
CF-Detox _{rigid} (Bhan et al., 2024)	2024	No	No	Yes	No	Any	Text	Yes	Yes	Explicit
Lin et al. (2024a)	2024	No	Yes	No	Yes	LLM	Text	Yes	Yes	Explicit
RIDERS (Li et al., 2024)	2024	No	No	Yes	Yes	LLM	Text	Yes	Yes	Explicit
DETOXIGEN (Niu et al., 2024)	2024	No	Yes	Yes	No	LLM	Text	Yes	Yes	Explicit
CoARL (Hengle et al., 2024)	2024	No	Yes	No	Yes	Web	Text	Yes	Yes	Explicit
MICo (Siegelmann et al., 2024)	2024	Yes	No	Yes	Yes	LLM	Text	Yes	Yes	Explicit
Schäfer et al. (2024)	2024	No	No	Yes	No	Web	Text	Yes	Yes	Explicit
Wang et al. (2024b)	2024	Yes	No	Yes	Yes	LLM	Text	Yes	Yes	Explicit
ToxCL (Hoang et al., 2024)	2024	No	Yes	Yes	No	SMP	Text	Yes	No	Implicit
Verhoeven et al. (2024)	2024	No	No	Yes	No	SMP	Text	Yes	No	Explicit
BeaverTails (Ji et al., 2024)	2024	Yes	No	No	Yes	LLM	Text	Yes	Yes	Explicit
SELF-GUARD (Wang et al., 2024e)	2024	No	No	Yes	Yes	LLM	Text	Yes	Yes	Explicit
AEGIS (Ghosh et al., 2024)	2024	Yes	Yes	Yes	No	LLM	Text	Yes	Yes	Both
Chung and Bright (2024)	2024	Yes	No	Yes	No	SMP	Text	Yes	Yes	Explicit
GuardFormer (O'Neill, 2024)	2024	Yes	No	Yes	Yes	LLM	Text	Yes	Yes	Both
SharedCon (Ahn et al., 2024)	2024	No	No	Yes	No	WEB/SMP	Text	Yes	No	Implicit
Hu et al. (2024)	2024	No	No	Yes	No	LLM	Text	Yes	Yes	FN&M
BiasAlert (Fan et al., 2024b)	2024	No	No	No	Yes	LLM	Text	Yes	No	Both
DisCGen (Hassan & Alikhani, 2023)	2023	Yes	Yes	No	Yes	LLM	Text	Yes	Yes	Both
Markov et al. (2023)	2023	No	Yes	Yes	No	WEB	Text	Yes	No	Explicit
NeMo (Rebedea et al., 2023)	2023	No	Yes	Yes	No	LLM	Text	Yes	Yes	Both
HateMoji (Kirk et al., 2022)	2022	Yes	No	No	Yes	SMP	Text	Yes	No	Explicit
Factoid (Sakketou et al., 2022)	2022	Yes	No	No	Yes	WEB/SMP	Text	Yes	No	FN&M*
HSBERT (Toraman et al., 2022)	2022	Yes	No	No	Yes	SMP	Text	Yes	No	Explicit
CapsNet-ConvNet (Kumar & Sachdeva, 2022)	2022	Yes	No	Yes	No	SMP	Text Image/InfoG	Yes	No	Explicit
HateXplain (Mathew et al., 2021)	2021	Yes	No	Yes	Yes	SMP	Text	Yes	No	Explicit
LFTW (Vidgen et al., 2021c)	2021	No	Yes	Yes	No	SMP	Text	Yes	No	Both
Guest et al. (2021)	2021	Yes	No	No	Yes	SMP	Text	Yes	No	Explicit
Raj et al. (2021)	2021	No	Yes	No	No	Web	Text	Yes	No	Explicit
RedditBias (Barikeri et al., 2021)	2021	Yes	Yes	No	Yes	Web/SMP	Text	Yes	Yes	Both

Toxicity Dataset [Table 2](#)). Recently, many researchers have generated LLM-based techniques and frameworks. A notable approach called ToxCL ([Hoang et al., 2024](#)), brings a new perspective by not just detecting toxicity, but also explaining the detected instance. Explainable AI (XAI) is a research area to build transparency and trustworthiness of a model ([Mersha et al., 2025a,b](#)). It increases detection performance compared to developing separate models for detection and explanation.

In the following sections, we have listed the broader types of toxicity research.

7.2.1. Detection of explicit toxicity

A recent study conducted by [Shen et al. \(2025\)](#) on the detection of explicit toxicity in LLM-generated web content highlights limitations in the detection techniques that enable adversarial at-

tacks on LLMs with greater efficiency. HateBench (Shen et al., 2025) is a novel framework to create a dataset dynamically. It utilizes an ensemble of 8 models, out of which 4 are research (LFTW, TweetHate (Antypas & Camacho-Collados, 2023), HSBERT, and Bert-HateXplain) and 4 are commercial models (Perspective, Moderation, Detoxify-original, and Detoxify-unbiased). The study curated a dataset and proposed a detection mechanism for web content. The purpose of this research is to highlight that it is easy to launch hate campaigns without being detected, calling for a rigorous research effort in LLM Safety and detoxification. Although this framework cannot be used in real-world toxicity mitigation, it provides a base for future research.

Verhoeven et al. (2024) proposed another toxicity detection mechanism in SMP using Community models' generalization via graph meta-learning. They have trained the model on GossioCop and evaluated on CoAID and TwitterHateSpeech. They have pointed out that such models are either transductive (assumption: graphs are static and do not change) or inductive (assumption: graphs change all the time). They have pointed out the absence of datasets in inductive social graphs and achieved generalization using graph meta-learning. Although their results are encouraging, they have pointed out that further work is required to allow such a solution to be implemented in real-world applications.

Markov et al. (2023) proposed a detection mechanism (Moderation API) for web content on a 219K dataset of real-world, synthetic, and production content. They have acknowledged that their model suffers from bias, fairness, robustness, and multilingual support issues and requires further red-teaming and active learning experiments. Their dataset annotation is based on their taxonomy of eight categories and further sub-categories of undesired content or toxicity. Imbalance of data categories is another problem to solve; using data augmentation is suggested.

In developing HateMoji (Kirk et al., 2022) Kirk et al. curated two datasets, HateMojiCheck and HateMojiBuild, and proposed a detection mechanism for SMP where emojis are used in hateful language expressions. While they have curated a multiclass dataset, their experiments for hate detection detect binary level only. Their experiments only included six identities and had limited scope.

In creating HSBERT (Toraman et al., 2022) Toraman et al. curated a dataset of 100K instances for English and Turkish, each with scalability, and a proposed detection mechanism on SMP. These datasets cover only 5 subcategories of gender, race, religion, politics, and sports. Thus, limited in scope. In CapsNet-ConvNet (Kumar & Sachdeva, 2022), Kumar et al. curated a dataset of 10K instances from Twitter, YouTube, and Instagram and proposed a detection mechanism for cyberbullying on SMP. They have proposed a deep learning based model for toxicity detection for multi-modal data (text, image, and image with text). Due to the dynamic landscape of SMP, and Natural Language Processing (NLP) models' capability to outperform deep learning models, further research in cyberbullying detection is suggested using state-of-the-art methods.

While developing HateXplain (Mathew et al., 2021), Mathew et al. curated a multi-label dataset and proposed a detection mechanism using Ground Truth Attention with Filter for ten target groups on SMP. HateXplain dataset has three labels: one for Hate Speech, Offensive, Normal; the second for the target group; and the third for the rationale. This work shows that the models that perform great at first or second label classification cannot perform as well at third rationale explanation tasks.

Guest et al. (2021) curated a 6,567-post dataset for online misogyny, misogyny taxonomy, and proposed a binary detection mechanism for SMP. Their dataset is generated from 11 weeks of Reddit posts, annotated by trained annotators, and is publicly available. Their dataset is highly imbalanced in binary and across fine-grained misogyny, which prevents them from classifying the fine-grained using their mechanism, given the small size of their dataset. But their taxonomy and high-quality annotation pave the path for future work in detecting misogyny.

Raj et al. (2021) propose an architecture for detecting cyberbullying in online platforms. They have experimented with various machine learning algorithms and neural network models on two Wikipedia datasets. They found that a neural network bidirectional Gated Recurrent Unit with Global Vectors outperforms existing state-of-the-art techniques. This technique can be experimented on other datasets to compare their performance with the current state-of-the-art. Another observation of this work, which might be questionable, is their datasets' imbalance, having less than 9–12 % of toxic classifications. Neural networks' classification will tend to be towards the non-toxic side. If the performance parameters are shown for only the cyberbullying and toxic class would make their conclusion more reliable.

7.2.2. Detection of implicit toxicity

Ismail et al. (2025) have highlighted that the severe consequences of aggressive toxicity in online gaming platforms include suicide, and sensed the urgency for their detection. They have proposed "a novel embeddings-based valence lexicon approach" using natural language processing, to detect it on a Twitch Corpus, specifically for gaming platforms. They have also suggested that their approach should work on other online platforms as well. While this approach detects the implicit toxicity, its further target classification or explanation is not generated.

In ToxCL (Hoang et al., 2024), Hoang et al. proposed a framework of detection mechanism for an SMP with a Target Group Generator, an encoder-decoder for a detection teacher classifier, and a knowledge distillation approach for explanation. Thereby identifying the toxicity class, target group, and explaining their output. They have utilized target group labels in *Implicit Hate Corpus* and *SBIC* datasets, and nineteen fine-grained categories from HateXplain dataset. They have performed an effectiveness evaluation by comparing ToxCL with human evaluation. They have noted that their framework may not perform well on abbreviated texts, symbols, sarcasm, and irony; and may have more than one non-overlapping explanation.

In SharedCon (Ahn et al., 2024), Ahn et al. proposed a detection mechanism using shared semantics by clustering and contrastive learning for SMP and web content. They have used K-Means clustering and chosen a post randomly, which is close to the centroid as the shared semantic sentence in the cluster. While this method eliminates the human in evaluating the implicit toxic nature, the basis of the choice of clustering and the number of clusters is not discussed. Selecting the most suitable clustering algorithm for a specific dataset is crucial, as the effectiveness of this approach heavily depends on identifying optimal clusters, which serve as the core of this mechanism.

7.2.3. Detection of both implicit and explicit toxicity

Learning from the worst (LFTW) (Vidgen et al., 2021c) proposed by Vidgen et al. is a dynamic dataset creation mechanism for robust toxicity detection model training for SMP. It produces a balanced dataset across hate and non-hate categories. Their mechanism introduces perturbations after each round, for further challenging set creation, requiring trained annotators. They have pointed out that this technique requires specific infrastructure, cross-disciplinary domain expertise, and is time-consuming process.

Ghosh et al. (2024) presented a taxonomy with 13 critical and 9 sparse content safety risk categories in the human-LLMs conversation, curated AEGISAFETYDATASET, and developed an online adaptive content moderation technique called AEGISAFETYEXPERTS for LLM. During the publication, they acknowledged that their dataset was not complete, and further training of AEGIS is required. With a dataset size 26K, their performance outperformed existing state-of-the-art.

O'Neill (2024) proposed a synthetic data pipeline generation approach with multi-task learning and a guardrail classifier called GuardFormer for LLM. They have a synthetic data generation LLM for given prompts, which generates corresponding safe and unsafe prompts, rationale, and labels. Followed by custom policy guardrailing by specific model *PolicyGuard* and multi-policy guardrailing model *GuardFormer*

capable of being effective on all policies. It offers a smaller memory footprint under 512MB and is faster, outperforming state-of-the-art on prompt injection, toxicity classification (fine-tuned), and content-safety. While their performance on 8 datasets and 11 high-performing LLMs is superior, testing it on real-world scenarios would be further helpful.

Fan et al. proposed BiasAlert (Fan et al., 2024b), a plug-and-play tool using Retrieval-Augmented Generation (RAG) knowledge for bias detection in LLM. They have used their metric of efficacy score, classification score, attribution score, and overall score to evaluate and compare their performance. It is evaluated on the SBIC dataset only and requires an updated RAG database corresponding to the prompt-response generation.

Koh et al. proposed LATTE (Koh et al., 2024), a framework that demonstrates the dynamic adaptability of LLMs to various toxicity definitions, thereby allowing for their metric flexibility. It moves away from static definitions and metrics, leaning towards the socio-cultural aspects. They have noted that it may be subject to inductive bias. As discussed in Section 4.4, this makes LLMs vulnerable to toxicity detection issues.

7.2.4. Detection of bias and fallacy

Factoid (Sakketou et al., 2022) is a dataset of fake news and misinformation by Sakketou et al. They proposed a detection mechanism for SMP and web content using a Graph Attention Network (GAN). This dataset's users are binary annotated for misinformation spreaders or real news spreaders, and have four fine-grained labels with their ranges of factuality degree (-3, 3), political bias (-3, 3), scientific belief (-1, 1), and satire degree (0, 1). It is one such user-centric large dataset. Their proposed models' best performances range between 58.4 to 66.2% on F1-Scores, which have scope of improvement.

Hu et al. (2024) proposed an LLM-based technique called adaptive rationale guidance (ARG) network for fake news and misinformation detection, by encoding textual description and commonsense rationales for classification. They have evaluated LLM-only, small language model (SLM) only and LLM + SLM models. And found that LLM + SLM with ARG outperforms others. While their model does well in real news detection, it struggles in fake news detection. They have highlighted that their model relies on LLM-generated rationale only; there might be other perspectives outside of this rationale that should be considered.

For Bias detection, BiasAlert is proposed by Fan et al. is detailed in Section 7.2.3.

7.3. Toxicity detoxification

Detoxification is the mechanism to remove the toxicity. It is performed after toxicity detection. In online platforms, for implicit or explicit toxicity with or without bias. Detoxification either uses pairs of hate speech (toxic language) and counterspeech as a dataset to generate non-hate speech non-toxic language with the same semantics or intentions, thereby replacing toxic text, detoxifying the content. Corresponding research summarization is presented in Table 3.

7.3.1. Detoxification of explicit toxicity in LLM

Zheng et al. (2025) proposed a lightweight fine-tuned Sentence-BERT-based guardrails model for LLM safety. It reduces the parameter size to 67M compared to LLM-based (7B parameters) solutions while maintaining their performance comparable to state-of-the-art. But it is not customizable and flexible for few-shot topic filtering.

Another safeguarding model proposed by Zhang et al. is called ToxEdit (Zhang, 2025), which uses toxicity-aware knowledge editing. It detects toxicity dynamically while avoiding over-editing of non-toxic responses on SafeEdit. They have performed limited experiments on SafeEdit using an SVM classifier only, Neural network classification might yield better results. ToxEdit suffers from sentence repetition in its generation.

Diao et al. proposed a self-evolving adversarial safety optimization technique for LLMs called SEAS (Diao et al., 2024). They have gener-

ated a dataset of adversarial prompts and a pipeline for dynamic safety optimization in the SEAS framework, showing performance comparable to GPT-4. They have highlighted that the quality of their framework improves with the number of iterations; hence, initial dataset diversity may be limited. They have used Llama Guard 2 in their framework for safety classification; thus, human evaluation is advisable.

Ermis et al. proposed a framework called Goodtriever (Ermis et al., 2024) comparing RAG with fine-tuned models under evolving and static language with multilingual support using translation for text generation. It lays the foundation for multilingual toxicity mitigation. Due to a lack of low-resource languages, cultural differences, translation variation, and information loss during translation are evident.

Fu et al. (2024) curated an Instruction Tuning Safety-defense dataset with 2000 documents and proposed a fine-tuned model with Low-Rank Adaptation (LoRA). It identifies malicious documents, resulting in the refusal of processing, while processing benign documents effectively. This model is trained on a balanced dataset only.

Kim et al. proposed a training algorithm called adversarial direct performance optimization (ADPO). It assigns safe or unsafe responses using the toxic control token (Kim & Lee, 2024). This has limitations of harmfulness and bias in LLM's annotated labels. Human annotation is also prone to bias due to the same demographic profile. Further, only 16K dataset is used for this work; using 160K dataset is expected to enhance the performance.

Lin et al. (2024a) proposed a framework with four AI agents, an original chatbot, with user therapist, and critic, with reinforcement learning based LLM tuning to detect and mitigate toxicity in the behavior of chatbots. But this implementation highly relies on the quality of datasets during training to enhance generalization, such that chatbot's bias and user's privacy are as per ethical considerations. AI agent as a psychotherapist needs to ensure communication with empathy to make it effective. Quantification is necessary to be evaluated. Explainable AI and reinforcement learning techniques to be explored in complicated interactions to allow cooperation.

RIDERS (Li et al., 2024) is another mitigation technique proposed by Cao et al. that uses a toxic chain-of-thought (CoT) with residual decoding and the serial position swap. It reduces the toxic CoT problem and improves commonsense reasoning. They have focused on multiple-choice questions with open-ended commonsense reasoning, calling for benchmark-related research. Additionally, their reasoning task does not consider mathematical questions.

DETOXIGEN (Niu et al., 2024) is another detoxifier model proposed by Niu et al., which is an ensemble of a pre-trained language model as a generator that is parameter-efficient using contrastive decoding, and a detoxifier that produces toxic tokens for the generator. They are trained on *RealToxicityPrompts* dataset. This research lacks experimentation where both LLMs are trained on separate datasets and may not cover the complete toxicity scope. Additionally, their model's human evaluation and generalization study have not been conducted.

Hengle et al. proposed CoARL (Hengle et al., 2024), a framework using multiple instruction tuning using LoRA and RL using LLMs for counterspeech generation. While their model outperforms state-of-the-art LLMs on their metric, it is important to note that they have used a small dataset for their experiment and cannot be considered exhaustive. They have pointed out that their RL reward approach might add bias, non-alignment with human perceptions, and multiple feedback loops are not accounted for in the model, which may affect its effectiveness over time.

Sieglmann et al. proposed MICO (Sieglmann et al., 2024), Models with Inhibition Control for Preventative detoxification, and curated a dataset with 2850 entries. Due to the smaller size of their dataset, it is non-representative of diverse toxic content. Thus, this model can be trained on a larger dataset and may require fine-tuning and modifications.

Schäfer et al. (2024) proposed a technique for the mitigation of identity term bias in toxicity detection. They have defined a bias target in three levels, improving the generalization across targets. They have conducted a single run of experiments focused on religious bias and thus have not considered their scores to be robust.

Wang et al. (2024b) curated a dataset and proposed a technique with knowledge editing detoxification and an intra-operative Neural Monitoring mechanism. This work sets the groundwork for future research like ToxEdit (Zhang, 2025) described earlier.

Ji et al. proposed BeaverTails (Ji et al., 2024), curated a large dataset of QA pairs with meta-labels for safety, and evaluated the model with metrics of helpfulness and harmlessness. They have highlighted that the annotators' demographic diversity is limited, and there might be overlap in certain categories.

Wang et al. proposed an approach called SELF-GUARD (Wang et al., 2024e) to prevent jailbreaking attacks by performing content assessment and toxicity detection on prompts and responses. They have highlighted that this approach reduces harmful content generation probability but does not eliminate it.

7.3.2. Detoxification of both implicit and explicit toxicity in LLMs

DisCGen (Hassan & Alikhani, 2023) by Hassan et al. curated a counterspeech dataset and proposed a discourse-aware framework. Their classifier is trained on Multitarget-CONAN, which is a small dataset, exhibits bias. They have manually relabeled the wrongly tagged instances in their counterspeech dataset.

Rebedea et al. proposed NeMo (Rebedea et al., 2023), an open-source toolkit to add programmable rails to LLM-based applications for implicit and explicit toxicity. They acknowledge that it cannot be used as a standalone safety mechanism, but can be added to already existing embedded guardrails to prevent jailbreak attacks.

7.3.3. Detoxification explicit toxicity in SMP and web content

Chung and Bright (2024) curated a small dataset and proposed a detoxification technique for Twitter, where the football players were targeted during the UK Premier League championship. Barikaeri et al. proposed RedditBias (Barikaeri et al., 2021) and curated a dataset. Bhan et al. (2024) proposed CF-Detox_{topic}, a counterfactual-based approach for detoxifying text using local feature importance explanations. Their method preserves semantic content while reducing toxicity, demonstrating the potential of eXplainable AI to support language detoxification with interpretability and content faithfulness.

7.4. Model performance metrics and evaluations

Most toxicity detection and detoxification mechanisms are evaluated using the F1-Scores, accuracy, precision and recall metrics. To ensure that the model's performance is robust and fair, a generalized framework is essential. It should accommodate various fine-grained toxicities based on toxicity context, platform, and environment.

Nazir et al. proposed an open-source Python toolkit LangTest (Nazir et al., 2024), with 60 tests and claimed to help improve the performance of LLMs and NLP models in regards to toxicity recognition and mitigation. It dynamically generates the automated testing workflow based on real-world applications. It trains, evaluates, runs tests, augments the training datasets, and repeats the process. Test results before and after augmentation are compared for the "bias, robustness, accuracy, fairness, and security" aspects of the model.

For bias and fairness assessment of an LLM, Bouchard et al. proposed LangFair (Bouchard et al., 2025), a Python package. It claims to provide an evaluation framework tailored to LLM tasks, considering their responses. It generated toxicity score for the words and stereotypes based on occurrence and classification.

Rottger et al. proposed HateCheck (Röttger et al., 2021), consisting of 29 function tests, a test suite to evaluate toxicity detection models. They targeted the annotation process of toxicity datasets.

Verhoeven et al. (2024) proposed an evaluation setup for the generalization capability of community models for fallacy and hate-speech detection, claiming that social subgraphs and graph meta-learning have better generalizability compared to standard community models.

In most of the *toxicity mitigation* studies, where the dataset is curated or synthesized, a common problem is the lack of diverse annotators' backgrounds, causing bias. In many proposed techniques, the usage of a single small dataset for training and evaluation limits their scope, generalization, and robustness. Some studies using machine learning and deep learning techniques are conducted on an imbalanced dataset, which raises concerns about the reliability of their results. Many studies focused on providing the necessary context and demonstrating the problem's relevance to pave the way for future research.

8. Open challenges and future directions

Despite advances in toxicity detection, several challenges remain. These include the scarcity of representative datasets, difficulties in handling emojis and informal language, and the complexity of detecting multilingual and multimodal toxic content, implicit bias, and fairness issues in LLMs. Another major open challenge is the lack of interpretability in toxicity classifiers, which limits transparency and user trust. Addressing these challenges is essential for building inclusive and accountable online platforms and AI systems.

8.1. Datasets

Various datasets are being used in research in different domains to train machine learning models to detect toxicity. However, the dataset sizes are not large enough to generalize it to a large set of domains. There is an opportunity to merge many of the datasets and their classes into one large central dataset, so that research with high coverage and effectiveness can be performed. Such a dataset may serve as the starting point for various lines of research studies. Antypas and Camacho-Collados (2023) showed better performance in toxicity detection when 13 datasets of tweets were unified. During dataset unification, challenges like different formatting, class granularity such as binary or multi-class, imbalance among classes and languages were encountered. Torman et al. (2022) emphasized that the model detection performance highly depends on the quality of the dataset on which it is being trained. XAI techniques can support the construction and quality control of datasets. The labeling inconsistencies or annotation errors can be detected by analyzing model explanations. Further use of XAI-based highlighting to make more informed decisions, improving inter-annotator agreement. This is especially important given the documented biases and inconsistencies in commonly used toxicity datasets (Wiegand et al., 2019).

8.2. Multi-lingual multi-modal toxicity

This survey is limited to English content in textual mode. This taxonomy can be adapted to languages other than English, this remains unexplored in this survey paper. Research involving toxicity in most languages is sparse, and any toxicity taxonomy created for them could be expanded further. This survey has presented a thoughtful and detailed Toxicity Taxonomy, which should apply to them, generally speaking.

Toxicity detection and mitigation for modalities except text, such as audio, image, video, and multi-modal toxicity, are not surveyed in this work. There is scope to survey these modal toxicities based on the given taxonomy.

8.3. Implicit bias

Recently, research in implicit bias has picked up momentum. This survey is limited to unearthing the toxicity extant in platforms. Implicit bias exhibited during the dataset annotation process also besets datasets

with the bias inherent to experts or human annotators. Further, social and algorithmic bias are also not explored in this paper. [Gallegos et al. \(2024\)](#) presented a survey on LLM's bias and fairness and provided future directions for mitigation. Unintended bias in toxicity detection is challenging to address. [Lopez and Kibler \(2025\)](#) concluded that dependence on the annotators of the toxicity datasets makes it difficult to achieve reliability. Research has demonstrated bias in the toxicity detection process itself ([Gencoglu, 2020](#)) ([Vaidya et al., 2020](#)), a topic not part of this survey.

8.4. LLMs

Natural language processing with emojis, reinforcement learning with human feedback, and the vulnerabilities of LLMs pose a tremendous challenge in dealing with toxicity. The different individuals perceive emojis differently ([Zhukova & Herring, 2024](#)). [Gupta et al. \(2023\)](#) highlighted that the polarity of a statement's sentiment can be inverted when a single or multiple emojis are included in the sentiment analysis. The use of reinforcement learning with human feedback poses a serious challenge to retaining the models' initial training. New prompts continually enable LLMs to learn and update themselves, thereby improving their performance. In the case of prompt injection and jailbreaking attacks, this may slowly coax the model toward toxicity. This is a research gap suggesting LLMs should be equipped with a mechanism to identify which prompts to learn from or which ones to ignore. Vulnerabilities of LLMs beyond toxicity are not explored in this survey. However, we consider prompt injection and jailbreaking as part of toxicity in LLMs.

8.5. Algorithmic, social, and evaluation bias

While we have showcased research on implicit and explicit toxicity and bias based on target groups, we have not focused on identifying whether toxicity has arisen due to algorithmic bias or social bias. It is necessary to further distinguish which type of algorithmic bias it is afflicted with, whether it is consciously built-in or has crept in subconsciously. Explicit or implicit, algorithmic bias is challenging to evaluate. Evaluation metrics used in research are generic and not self-reflective, suggestive of bias in itself. How to come up with an evaluation metric for toxicity, free of bias, is another huge challenge.

8.6. Explainable AI

Studies have shown that models trained on biased datasets can incorrectly classify neutral identity-related terms such as 'black' or 'gay' as toxic ([Bender et al., 2021](#)). XAI techniques can uncover such patterns by highlighting consistently high attribution scores for these identity terms. This insight allows us to identify and correct biased behavior in models. XAI improves the transparency, fairness, and effectiveness of the toxicity classification models and detoxification systems by generating explanations on their decisions ([Mersha et al., 2024b](#)). XAI techniques, such as LIME ([Ribeiro et al., 2016](#)), SHAP ([Lundberg & Lee, 2017](#)), and Integrated Gradients ([Sundararajan et al., 2017](#)), aid in identifying words that contribute to the classification of toxicity, thereby enhancing interpretability and user trust, providing higher accuracy and fairness in AI systems, by highlighting toxic elements in text, improving evaluation metrics, and quality of the dataset ([Mersha et al., 2024a, 2025c](#)). XAI also enables content moderation and online safety ([Chamola et al., 2023; Mersha et al., 2025a](#)).

9. Conclusion

The toxicity in online platforms and AI Systems continues to remain a challenge, even though there is copious research with various perspectives and contexts being conducted, due to various reasons. One, the effort is usually piecemeal and thus not holistic, and needs consolidation.

Two, the technology is evolving, and so are the facets of toxicity. Adaptability to evolving technology is necessary. Three, the platforms deal with toxicity in a reactive manner and not a proactive manner. Proactive techniques rather than reactive ones should be encouraged. Four, a framework of governance policies is not usually involved in the detection and mitigation of toxicity in digital communication. Five, holistic approaches to develop toxicity evaluation metrics have not been forthcoming. Six, as more and more AI systems are integrated into our lives, more psychological assessments must be considered for effective quality assurance.

Addressing AI-driven toxicities requires improved content moderation strategies, enhanced model transparency, and bias-mitigation techniques, such as adversarial training, XAI, and reinforcement learning with human feedback ([Liu, 2023](#)). Additionally, as discussed in the [Section 2.3](#), a collaborative effort is highly warranted. Adversaries are always several steps ahead in attacking the ecosystem with toxic behavior. All stakeholders need to continuously improve mechanisms for detection and detoxification with evolving technology and omnipresent cyberattacks to mitigate the psychological impacts society is direly facing.

CRediT authorship contribution statement

Smita Khapre: Conceptualization, Writing – original draft, Visualization; **Melkamu Abay Mersha:** Writing – original draft, Visualization; **Hassan Shakil:** Writing – original draft; **Jonali Baruah:** Writing – original draft; **Jugal Kalita:** Supervision, Project administration, Writing – review & editing.

Data availability

No data was used for the research described in the article.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Ahn, H., Kim, Y., Kim, J., & Han, Y.-S. (2024). Sharedcon: Implicit hate speech detection using shared semantics. In *Findings of the association for computational linguistics ACL 2024* (pp. 10444–10455).
- Akar, E. (2025). Exploring the impact of social network structures on toxicity in online mental health communities. *Computers in Human Behavior*, 165, 108542.
- Ali, M., Ullah, R., Khan, U., Rafique, K., Murad, S. (2025). The impact of social media addiction on teenage depression: A correlational study. *Indus Journal of Bioscience Research*, 3(2), 163–173.
- Ali, S. (2024). Youth, social media, and online safety: a holistic approach towards detecting and mitigating risks in online conversations. Ph.D. thesis. Boston University.
- AlShami, A. K., Rabinowitz, R., Lam, K., Shleibik, Y., Mersha, M., Boult, T., & Kalita, J. (2024). Smart-vision: Survey of modern action recognition techniques in vision. *Multimedia Tools and Applications*, 84, 1–72.
- Antypas, D., & Camacho-Collados, J. (2023). Robust hate speech detection in social media: A cross-dataset empirical evaluation. In *The 7th workshop on online abuse and harms (WOAH)* (p. 231).
- Aprin, F. (2024). Supporting young learners in identifying toxic content and misinformation in social media: a virtual learning companion approach. Ph.D. thesis. Dissertation, Duisburg, Essen, Universität Duisburg-Essen, 2023.
- Archana Balkrishna, Y. (2024). An analysis on the use of image design with generative AI technologies. *International Journal of Trend in Scientific Research and Development*, 8(1), 596–599.
- Arora, A., Nakov, P., Hardalov, M., Sarwar, S. M., Nayak, V., Dinkov, Y., Zlatkova, D., Dent, K., Bhatawdekar, A., Bouchard, G., Augenstein, I. (2023). Detecting harmful content on online platforms: What platforms need vs. where research efforts go. *ACM Computing Surveys*, 56(3), 1–17.
- Bączkowska, A., Lewandowska-Tomaszczyk, B., Źitnik, S., Liebeskind, C., Trojszczak, M., & Valunaite Oleskeviciene, G. (2025). Implicit offensive language taxonomy. *Lodz Papers in Pragmatics*.
- Bai, X., Wang, A., Sucholutsky, I., & Griffiths, T. L. (2024). Measuring implicit bias in explicitly unbiased large language models. In *NeurIPS 2024 workshop on behavioral machine learning*.
- Banaji, M. R., Bhaskar, R., & Brownstein, M. (2015). When bias is implicit, how might we think about repairing harm? *Current Opinion in Psychology*, 6, 183–188.

- Bansal, A., Mahadev, S., Sharma, K. K., Aggarwal, A., Goyal, S., Choudhary, K., Chawla, K., Jain, K., & Bhasin, M. (2012). Classification of flames in computer mediated communications. *International Journal of Computer Applications*, 975, 8887.
- Bansal, A., Mahadev, S., Sharma, K. K., Aggarwal, A., Goyal, S., Choudhary, K., Chawla, K., Jain, K., & Bhasin, M. (2011). Classification of flames in computer mediated communications. *International Journal of Computer Applications*, 975, 8887.
- Barker, S., Lauscher, A., Vulic, I., & Glavaš, G. (2021). Redditbias: A real-world resource for bias evaluation and debiasing of conversational language models. 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)* (pp. 1941–1955).
- Beknazary-Yuzbashev, G., Jiménez-Durán, R., McCrosky, J., & Stalinski, M. (2025). Toxic content and user engagement on social media: Evidence from a field experiment. CESifo Working Papers, CESifo GmbH, 11644.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on fairness, accountability, and transparency* (pp. 610–623).
- Bertoni, A. L. C., & Serafini, M. C. (2023). Ethical content in artificial intelligence systems: A demand explained in three critical points. *Frontiers in Psychology*, 14, 1074787.
- Bhan, M., Vittaut, J.-N., Achache, N., Legrand, V., Chesneau, N., Blangero, A., Murris, J., & Lesot, M.-J. (2024). Mitigating text toxicity with counterfactual generation. CoRR.
- Biju, P. R., & Gayathri, O. (2023). Self-breeding fake news. *The International Journal of Information, Diversity, & Inclusion*, 7(1/2), 1–25.
- Bilewicz, M., & Soral, W. (2020). Hate speech epidemic: the dynamic effects of derogatory language on intergroup relations and political radicalization. *Political Psychology*, 41, 3–33.
- Blackwell, L., Ellison, N., Elliott-Deflo, N., & Schwartz, R. (2019). Harassment in social virtual reality: Challenges for platform governance. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1–25.
- Blodgett, S. L., Barocas, S., Daumé, H., III, & Wallach, H. (2020). Language (technology) is power: A critical survey of “bias” in NLP. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 5454–5476).
- Bouchard, D., Chauhan, M. S., Skarbrevik, D., Bajaj, V., & Ahmad, Z. (2025). Langfair: A python package for assessing bias and fairness in large language model use cases. *Journal of Open Source Software*, 10(105), 7570.
- Bradford, A. (2023). Europe's digital constitution. *Va. J. Int'l L.*, 64, 1.
- Brown, N. I. (2020). Regulatory goldilocks: Finding the just and right fit for content moderation on social platforms. *Tex A&M L Review*, 8, 451.
- Cabrera, J., Loyola, M. S., Magaña, I., & Rojas, R. (2023). Ethical dilemmas, mental health, artificial intelligence, and ILM-based chatbots. In *International work-conference on bioinformatics and biomedical engineering* (pp. 313–326). Springer.
- Cao, Y., Li, S., Liu, Y., Yan, Z., Dai, Y., Yu, P., & Sun, L. (2025). A survey of AI-generated content (aigc). *ACM Computing Surveys*, 57(5), 1–38.
- Chamola, V., Hassija, V., Sulthana Abdul Kareem, R., Ghosh, D., Dhingra, D., & Sikdar, B. (2023). A review of trustworthy and explainable artificial intelligence (XAI). *IEEE Access*, 11, 78994–79015.
- Chandler, D., & Munday, R. (2011). A dictionary of media and communication. Oxford University Press, USA.
- Cheng, L., Varshney, K. R., & Liu, H. (2021). Socially responsible AI algorithms: Issues, purposes, and challenges. *Journal of Artificial Intelligence Research*, 71, 1137–1181.
- Cheng, Z., & Li, Y. (2024). Like, comment, and share on tiktok: Exploring the effect of sentiment and second-person view on the user engagement with tiktok news videos. *Social Science Computer Review*, 42(1), 201–223.
- Chung, Y.-L., & Bright, J. (2024). On the effectiveness of adversarial robustness for abuse mitigation with counterspeech. In *Proceedings of the 2024 conference of the north american chapter of the association for computational linguistics: Human language technologies (volume 1: Long papers)* (pp. 6981–6995).
- Cui, L., & Lee, D. (2020). CoAID: COVID-19 healthcare misinformation dataset. CoRR.
- Dangsawang, B., & Nuchitprasitchai, S. (2024). A machine learning approach for detecting customs fraud through unstructured data analysis in social media. *Decision Analytics Journal*, 10, 100408.
- Delgado-Rodriguez, M., & Llorca, J. (2004). Bias. *Journal of Epidemiology & Community Health*, 58(8), 635–641.
- Dementieva, D., Babakov, N., & Panchenko, A. (2024). MultiParaDetox: Extending text detoxification with parallel data to new languages. In K. Duh, H. Gomez, & S. Bethard (Eds.), *Proceedings of the 2024 conference of the north american chapter of the association for computational linguistics: Human language technologies (volume 2: Short papers)* (pp. 124–140). Mexico City, Mexico: Association for Computational Linguistics. <https://aclanthology.org/2024.nacl-short.12>. <https://doi.org/10.18653/v1/2024.nacl-short.12>
- Derave, T., Prince Sales, T., Gailly, F., & Poels, G. (2021). Comparing digital platform types in the platform economy. In *International conference on advanced information systems engineering* (pp. 417–431). Springer.
- Desai, A. P., Ravi, T., Luqman, M., Sharma, M., Kota, N., & Yadav, P. (2024). Gen-AI for user safety: A survey. In *2024 IEEE International conference on big data (bigdata)* (pp. 5315–5324). IEEE.
- Despot, K. Š., Anić, A. O., & Veale, T. (2023). “somewhere along your pedigree, a bitch got over the wall!” a proposal of implicitly offensive language typology. *Lodz Papers in Pragmatics*, 19(2), 385–414.
- Dhamala, J., Sun, T., Kumar, V., Krishna, S., Praksachatkun, Y., Chang, K.-W., & Gupta, R. (2021). Bold: Dataset and metrics for measuring biases in open-ended language generation. In *Proceedings of the 2021 ACM Conference on fairness, accountability, and transparency* (pp. 862–872).
- Dhingra, M., & Mudgal, R. K. (2019). Historical evolution of social media: An overview. In *International conference on advances in engineering science management & technology (ICAESMT)-2019, Uttarakhand university, Dehradun, India*.
- Diao, M., Li, R., Liu, S., Liao, G., Wang, J., Cai, X., & Xu, W. (2024). Seas: Self-evolving adversarial safety optimization for large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)-2025* 39(22), 385–414.
- Dubber, M. D., Pasquale, F., & Das, S. (2020). The oxford handbook of ethics of AI. Oxford University Press.
- Duchêne, C., Jamet, H., Guillaume, P., & Dehak, R. (2023). A benchmark for toxic comment classification on civil comments dataset. CoRR.
- Duncombe, C. (2020). Social media and the visibility of horrific violence. *International Affairs*, 96(3), 609–629.
- Eberl, J.-M., Tolochko, P., Jost, P., Heidenreich, T., & Boomgaarden, H. G. (2020). What's in a post? How sentiment and issue salience affect users' emotional reactions on facebook. *Journal of Information Technology & Politics*, 17(1), 48–65.
- ElSherief, M., Ziems, C., Muchlinski, D., Anupindi, V., Seybolt, J., De Choudhury, M., & Yang, D. (2021). Latent hatred: A benchmark for understanding implicit hate speech. In *Proceedings of the 2021 conference on empirical methods in natural language processing* (pp. 345–363).
- Ermis, B., Pozzobon, L., Hooker, S., & Lewis, P. (2024). From one to many: Expanding the scope of toxicity mitigation in language models. In *Findings of the association for computational linguistics ACL 2024* (pp. 15041–15058).
- Evangeline, S. I. (2025). The double edged sword of AI: Legal and regulatory implications of AI in cyberbullying. In *Combating cyberbullying with generative AI* (pp. 299–328). IGI Global Scientific Publishing.
- Fan, L., Li, L., & Hemphill, L. (2024a). Toxicity on social media during the 2022 mpox public health emergency: Quantitative study of topical and network dynamics. *Journal of Medical Internet Research*, 26, e52997.
- Fan, Z., Chen, R., Xu, R., & Liu, Z. (2024b). Biasalert: A plug-and-play tool for social bias detection in LLMs. In *Proceedings of the 2024 conference on empirical methods in natural language processing* (pp. 14778–14790).
- Fang, X., Che, S., Mao, M., Zhang, H., Zhao, M., & Zhao, X. (2024). Bias of AI-generated content: An examination of news produced by large language models. *Scientific Reports*, 14(1), 5224.
- Fanton, M., Bonaldi, H., Tekiroglu, S. S., & Guerini, M. (2021). Human-in-the-loop for data collection: A multi-target counter narrative dataset to fight online hate speech. 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)* (pp. 3226–3240).
- Ferrara, E. (2023). Should chatGPT be biased? Challenges and risks of bias in large language models. *Challenges and Risks of Bias in Large Language Models (October 26, 2023)*.
- Fitzgerald, J. M., Kinney, K. L., Phan, K. L., & Klumpp, H. (2020). Distinct neural engagement during implicit and explicit regulation of negative stimuli. *Neuropsychologia*, 145, 106675.
- Flew, T. (2021). Regulating platforms. John Wiley & Sons.
- Florindi, E., Franzoni, V., & Milani, A. (2024). Ethical boundaries in AI-generated imagery: Analysis of controversial content constraints. In *International conference on computational science and its applications* (pp. 292–302). Springer.
- Fu, Y., Xiao, W., Chen, J., Li, J., Papalexakis, E., Chien, A., & Dong, Y. (2024). Cross-task defense: Instruction-tuning LLMs for content safety. In *Proceedings of the 4th workshop on trustworthy natural language processing (trustNLP 2024)* (pp. 85–93).
- Gallegos, I. O., Rossi, R. A., Barrow, J., Tanjim, M. M., Kim, S., Dernoncourt, F., Yu, T., Zhang, R., & Ahmed, N. K. (2024). Bias and fairness in large language models: A survey. *Computational Linguistics*, 50(3), 1097–1179.
- Gambin, Á. F., Yazidi, A., Vasilakos, A., Haugerud, H., & Djenouri, Y. (2024). Deepfakes: Current and future trends. *Artificial Intelligence Review*, 57(3), 64.
- Gamson, W. A., Croteau, D., Hoynes, W., & Sassen, T. (1992). Media images and the social construction of reality. *Annual Review of Sociology*, 18(1), 373–393.
- Gao, S., He, L., Chen, Y., Li, D., & Lai, K. (2020). Public perception of artificial intelligence in medical care: Content analysis of social media. *Journal of Medical Internet Research*, 22(7), e16649.
- Gao, Y., Qin, W., Murali, A., Eckart, C., Zhou, X., Beel, J. D., Wang, Y.-C., & Yang, D. (2024). A crisis of civility? Modeling incivility and its effects in political discourse online. In *Proceedings of the international AAAI conference on web and social media* (pp. 408–421). (vol. 18).
- Garg, S., & Sharma, D. K. (2020). New politifac: A dataset for counterfeit news. In *2020 9th International conference system modeling and advancement in research trends (SMART)* (pp. 17–22). IEEE.
- Garg, T., Masud, S., Suresh, T., & Chakraborty, T. (2023). Handling bias in toxic speech detection: A survey. *ACM Computing Surveys*, 55(13s), 1–32.
- Gehman, S., Gururangan, S., Sap, M., Choi, Y., & Smith, N. A. (2020). Realtoxicityprompts: Evaluating neural toxic degeneration in language models. In *Findings of the association for computational linguistics: EMNLP 2020* (pp. 3356–3369).
- Gencoglu, O. (2020). Cyberbullying detection with fairness constraints. *IEEE Internet Computing*, 25(1), 20–29.
- George, A. S. (2023). Toxicity in the workplace: The silent killer of careers and lives. *Partners Universal International Innovation Journal*, 1(2), 1–21.
- Ghosh, S., Varshney, P., Galinkin, E., & Parisien, C. (2024). Aegis: Online adaptive AI content safety moderation with ensemble of ILM experts. CoRR.
- Gjerazi, B., & Skana, P. (2023). Impact of politically motivated fake news on public opinion: A case study of deliberate dissemination of disinformation. *Balkan Social Science Review*, 22(22), 365–383.
- Gossipcop 2025. <https://huggingface.co/datasets/osusume/Gossipcop>.
- Gongane, V. U., Munot, M. V., & Anuse, A. D. (2022). Detection and moderation of detrimental content on social media platforms: Current status and future directions. *Social Network Analysis and Mining*, 12(1), 129.
- Greenfield, S. (2015). Mind change: How digital technologies are leaving their mark on our brains. Random House.

- Guerreiro, N. M., Alves, D. M., Waldendorf, J., Haddow, B., Birch, A., Colombo, P., & Martins, A. F. T. (2023). Hallucinations in large multilingual translation models. *Transactions of the Association for Computational Linguistics*, 11, 1500–1517.
- Guess, A. M., & Munger, K. (2023). Digital literacy and online political behavior. *Political science research and methods*, 11(1), 110–128.
- Guest, E., Vidgen, B., Mittos, A., Sastry, N., Tyson, G., & Margetts, H. (2021). An expert annotated dataset for the detection of online misogyny. In *Proceedings of the 16th conference of the european chapter of the association for computational linguistics: main volume* (pp. 1336–1350).
- Gunturi, U. S., Ding, X., & Rho, E. H. (2023). Toxvis: Enabling interpretability of implicit vs. explicit toxicity detection models with interactive visualization. *CoRR*.
- Gupta, S., Singh, A., & Kumar, V. (2023). Emoji, text, and sentiment polarity detection using natural language processing. *Information*, 14(4), 222.
- Hall, J. A. (2025). Loneliness and social media. *Annals of the New York Academy of Sciences*, 1543(1), 5–16.
- Hancock, J. T., & Bailenson, J. N. (2021). The social impact of deepfakes. *Cyberpsychology, behavior, and social networking*, 34(3), 149–152.
- Hartvigsen, T., Gabriel, S., Palangi, H., Sap, M., Ray, D., & Kamar, E. (2022). Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 3309–3326).
- Hassan, S., & Alikhani, M. (2023). DisCGen: A framework for discourse-informed counterspeech generation. In *Proceedings of the 13th international joint conference on natural language processing and the 3rd conference of the asia-pacific chapter of the association for computational linguistics (volume 1: Long papers)* (pp. 420–429).
- Helwe, C., Calamai, T., Paris, P.-H., Clavel, C., & Suchanek, F. M. (2024). Mafalda: A benchmark and comprehensive study of fallacy detection and classification. In *Naacl 2024-North American chapter of the association for computational linguistics*.
- Hengle, A., Padhi, A., Singh, S., Bandhakavi, A., Akhtar, M. S., & Chakraborty, T. (2024). Intent-conditioned and non-toxic counterspeech generation using multi-task instruction tuning with RLAIF. In *Proceedings of the 2024 conference of the north american chapter of the association for computational linguistics: Human language technologies (volume 1: Long papers)* (pp. 6716–6733).
- Hoang, N., Do, X. L., Do, D. A., Vu, D. A., & Luu, A. T. (2024). ToXCL: A unified framework for toxic speech detection and explanation. In K. Duh, H. Gomez, & S. Bethard (Eds.), *Proceedings of the 2024 conference of the north american chapter of the association for computational linguistics: Human language technologies (volume 1: Long papers)* (pp. 6460–6472). Mexico City, Mexico: Association for Computational Linguistics. <https://aclanthology.org/2024.naacl-long.359>. <https://doi.org/10.18653/v1/2024.naacl-long.359>
- Hu, B., Mao, Z., & Zhang, Y. (2025). An overview of fake news detection: From a new perspective. *Fundamental Research*, 5(1), 332–346.
- Hu, B., Sheng, Q., Cao, J., Shi, Y., Li, Y., Wang, D., & Qi, P. (2024). Bad actor, good advisor: Exploring the role of large language models in fake news detection. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 22105–22113). (vol. 38).
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., Liu, T. (2025). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2), 1–55.
- Huang, Y., Liu, Y., Thirukovalluru, R., Cohan, A., & Dhingra, B. (2024). Calibrating long-form generations from large language models. In *Findings of the association for computational linguistics: EMNLP 2024* (pp. 13441–13460).
- Imran, M., & Almusharraf, N. (2024). Google gemini as a next generation AI educational tool: A review of emerging educational technology. *Smart Learning Environments*, 11(1), 22.
- Ismail, H., Khalil, A., & Jasmy, A. (2025). Enhancing online toxicity detection on gaming networks: A novel embeddings-based valence lexicon approach. *International Journal of Data Science and Analytics*, 20, 1–12.
- Jawad, M., Talreja, K., Bhutto, S. A., & Faizan, K. (2024). Investigating how AI personalization algorithms influence self-perception, group identity, and social interactions online. *Review of Applied Management and Social Sciences*, 7(4), 533–550.
- Ji, J., Liu, M., Dai, J., Pan, X., Zhang, C., Bian, C., Chen, B., Sun, R., Wang, Y., & Yang, Y. (2024). Beavertails: Towards improved safety alignment of llm via a human-preference dataset. *Advances in Neural Information Processing Systems*, 36, 24678–24704.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM computing surveys*, 55(12), 1–38.
- Jiang, H., Zhang, X., Cao, X., Breazeal, C., Roy, D., & Kabbara, J. (2024). PersonalLLM: Investigating the ability of large language models to express personality traits. In *Findings of the association for computational linguistics: NAACL 2024* (pp. 3605–3627).
- Johnson, G. M. (2021). Algorithmic bias: On the implicit biases of social technology. *Synthese*, 198(10), 9941–9961.
- Keen, A. (2011). The cult of the amateur: How blogs, myspace, youtube and the rest of today's user generated media are killing our culture and. John Murray Business.
- Keller, D. (2018). Internet platforms: Observations on speech, danger, and money. *Hoover Institution's Aegis paper series*, (1807).
- Khapre, S. (2025). Mitigation of social media platforms impact on the users. *WSCG 2025 Proceedings, CSRN 3501, 33. International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision WSCG 2025*, 33, 139–148.
- Kiela, D., Firooz, H., Mohan, A., Goswami, V., Singh, A., Ringshia, P., & Testuggine, D. (2020). The hateful memes challenge: Detecting hate speech in multimodal memes. *Advances in neural information processing systems*, 33, 2611–2624.
- Kim, J., Derakhshan, A., & Harris, I. (2024). Robust safety classifier against jailbreaking attacks: Adversarial prompt shield. In *Proceedings of the 8th workshop on online abuse and harms (WOAH 2024)* (pp. 159–170).
- Kim, M. (2024). Unveiling the e-servicescape of chatGPT: Exploring user psychology and engagement in AI-powered chatbot experiences. *Behavioral Sciences*, 14(7), 558.
- Kim, S., & Lee, G. (2024). Adversarial DPO: Harnessing harmful data for reducing toxicity with minimal impact on coherence and evasiveness in dialogue agents. In *Findings of the association for computational linguistics: NAACL 2024* (pp. 1821–1835).
- Kiritchenko, S., Nejadgholi, I., & Fraser, K. C. (2021). Confronting abusive language online: A survey from the ethical and human rights perspective. *Journal of Artificial Intelligence Research*, 71, 431–478.
- Kirk, H., Vidgen, B., Röttger, P., Thrush, T., & Hale, S. (2022). Hatemoji: A test suite and adversarially-generated dataset for benchmarking and detecting emoji-based hate. In *Proceedings of the 2022 conference of the north american chapter of the association for computational linguistics: Human language technologies* (pp. 1352–1368).
- Koh, H., Kim, D., Lee, M., & Jung, K. (2024). Can LLMs recognize toxicity? A structured investigation framework and toxicity metric. Miami, Florida, USA: Association for Computational Linguistics.
- Kordzadeh, N., & Ghasemaghaei, M. (2022). Algorithmic bias: Review, synthesis, and future research directions. *European Journal of Information Systems*, 31(3), 388–409.
- Korre, K., Pavlopoulos, J., Gajo, P., & Barrón-Cedeño, A. (2024). Hate speech according to the law: An analysis for effective detection. *arXiv preprint arXiv:2412.06144*.
- Kumar, A., & Sachdeva, N. (2022). Multimodal cyberbullying detection using capsule network with dynamic routing and deep convolutional neural network. *Multimedia Systems*, 28(6), 2043–2052.
- Kumar, D., Jain, U., Agarwal, S., & Harshangi, P. (2024). Investigating implicit bias in large language models: A large-scale study of over 50 LLMs. In *Neurips safe generative AI workshop 2024*.
- Kumar, M., & Mondal, A. (2018). A study on internet addiction and its relation to psychopathology and self-esteem among college students. *Industrial psychiatry journal*, 27(1), 61–66.
- Latif, Y. A. (2025). Hallucinations in large language models and their influence on legal reasoning: Examining the risks of AI-generated factual inaccuracies in judicial processes. *Journal of Computational Intelligence, Machine Reasoning, and Decision-Making*, 10(2), 10–20.
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. D., Sunstein, C. R., Thorson, E. A., Watts, D. J., Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096.
- Lee, H. (2005). Behavioral strategies for dealing with flaming in an online forum. *The Sociological Quarterly*, 46(2), 385–403.
- Lehtimäki, M.-L. (2024). Navigating ethical and practical challenges in AI-driven visual content creation. Master's thesis. Haaga-Helia University of Applied Sciences.
- Lemaire, M., Ye, S., Le Stanc, L., Borst, G., & Cassotti, M. (2025). The development of media truth discernment and fake news detection is related to the development of reasoning during adolescence. *Scientific Reports*, 15(1), 6854.
- Lewandowska-Tomaszczyk, B., Baczkowska, A., Liebeskind, C., Valunaite Oleskeviciene, G., & Žitnik, S. (2023). An integrated explicit and implicit offensive language taxonomy. *Lodz Papers in Pragmatics*, 19(1), 7–48.
- Li, D., Gupta, K., Bhaduri, M., Sathiadoss, P., Bhatnagar, S., & Chong, J. (2025a). Comparative diagnostic accuracy of GPT-4o and LLaMA 3.70b: Proprietary vs. open-source large language models in radiology. *Clinical Imaging*, 118, 110382.
- Li, J., Cao, P., Wang, C., Jin, Z., Chen, Y., Zeng, D., Liu, K., & Zhao, J. (2024). Focus on your question! interpreting and mitigating toxic cot problems in commonsense reasoning. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 9206–9230).
- Li, L., Shi, Z., Hu, X., Dong, B., Qin, Y., Liu, X., Sheng, L., & Shao, J. (2025b). T2ISafety: Benchmark for assessing fairness, toxicity, and privacy in image generation. *Proceedings of the Computer Vision and Pattern Recognition Conference* (pp. 13381–13392).
- Lin, B., Bounouffou, D., Cecchi, G., & Varshney, K. (2024a). Towards healthy AI: Large language models need therapists too. In A. Ovalle, K.-W. Chang, Y. T. Cao, N. Mehrabi, J. Zhao, A. Galstyan, J. Dhamala, A. Kumar, & R. Gupta (Eds.), *Proceedings of the 4th workshop on trustworthy natural language processing (trustNLP 2024)* (pp. 61–70). Mexico City, Mexico: Association for Computational Linguistics. <https://aclanthology.org/2024.trustnlp-1.6>. <https://doi.org/10.18653/v1/2024.trustnlp-1.6>
- Lin, R., & Utz, S. (2015). The emotional responses of browsing facebook: happiness, envy, and the role of tie strength. *Computers in human behavior*, 52, 29–38.
- Lin, S., Hilton, J., & Evans, O. (2022). TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 3214–3252).
- Lin, S.-Y., Chien, S.-Y., Chen, Y.-Z., & Chien, Y.-H. (2024b). Combating online malicious behavior: Integrating machine learning and deep learning methods for harmful news and toxic comments. *Information Systems Frontiers*, (pp. 1–16).
- Liu, G. K.-M. (2023). Transforming human interactions with AI via reinforcement learning with human feedback (RLHF). *Massachusetts Institute of Technology*.
- Liu, L. (2024). The algorithmic bias in recommendation systems and its social impact on user behavior: Algorithmic bias in recommendation systems. *International Theory and Practice in Humanities and Social Sciences*, 1(1), 290–303.
- Liu, Y., Cai, C., Zhang, X., Yuan, X., & Wang, C. (2024). Arondight: Red teaming large vision language models with auto-generated multi-modal jailbreak prompts. In *Proceedings of the 32nd ACM international conference on multimedia* (pp. 3578–3586).
- Logacheva, V., Dementieva, D., Ustyantsev, S., Moskovskiy, D., Dale, D., Krotova, I., Semenov, N., & Panchenko, A. (2022). Paradetox: Detoxification with parallel data. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 6804–6818).
- Lopez, H., & Kübler, S. (2025). Context in abusive language detection: On the interdependence of context and annotation of user comments. *Discourse, Context & Media*, 63, 100848.

- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.
- Luong, T., Le, T.-T., Ngo, L., & Nguyen, T. (2024). Realistic evaluation of toxicity in large language models. In *Findings of the association for computational linguistics ACL 2024* (pp. 1038–1047).
- Mademlis, I., Mancuso, M., Paternoster, C., Evangelatos, S., Finlay, E., Hughes, J., Radoglou-Grammatikis, P., Sarigiannidis, P., Stavropoulos, G., Votis, K., Papadopoulos, G. Th. (2024). The invisible arms race: Digital trends in illicit goods trafficking and AI-enabled responses. *IEEE Transactions on Technology and Society*, 6(2), 181–199.
- Maity, K., Poornash, A. S., Saha, S., & Pasupu, K. (2024a). ToxVI: A multimodal LLM-based framework for generating intervention in toxic code-mixed videos. In *Proceedings of the 33rd ACM international conference on information and knowledge management* (pp. 3937–3942).
- Maity, K., Sangeetha, P., Saha, S., & Bhattacharyya, P. (2024b). ToxvidLM: A multimodal framework for toxicity detection in code-mixed videos. In *Findings of the association for computational linguistics ACL 2024* (pp. 11130–11142).
- Maleki, M., Arami, M., Mead, E., Kready, J., & Agarwal, N. (2022). Applying an epidemiological model to evaluate the propagation of toxicity related to COVID-19 on twitter. *Hawaii International Conference on System Sciences, Proceedings of the 55th Hawaii International Conference on System Sciences*, 55, 3275–3284.
- Manheim, K., & Kaplan, L. (2019). Artificial intelligence: Risks to privacy and democracy. *Yale JL & Tech.*, 21, 106.
- Markov, T., Zhang, C., Agarwal, S., Nekoul, F. E., Lee, T., Adler, S., Jiang, A., & Weng, L. (2023). A holistic approach to undesired content detection in the real world. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 15009–15018). (vol. 37).
- Mathew, B., Saha, P., Yimam, S. M., Biemann, C., Goyal, P., & Mukherjee, A. (2021). Hatexplain: A benchmark dataset for explainable hate speech detection. In *Proceedings of the AAAI conference on artificial intelligence*. (vol. 35), (pp. 14867–14875).
- Mersha, M., Bitewa, M., Abay, T., & Kalita, J. (2024a). Explainability in neural networks for natural language processing tasks. arXiv preprint arXiv:2412.18036.
- Mersha, M., Lam, K., Wood, J., AlShami, A., & Kalita, J. (2024b). Explainable artificial intelligence: A survey of needs, techniques, applications, and future direction. *Neurocomputing*, 599, 128111.
- Mersha, M. A., Yigezu, M. G., Bade, G. Y., Kalita, J., Kolesnikova, O., Gelbukh, A. (2024c). Ethio-fake: Cutting-edge approaches to combat fake news in under-resourced languages using explainable ai. *Procedia Computer Science*, 244, 133–142.
- Mersha, M. A., Yigezu, M. G., & Kalita, J. (2025a). Evaluating the effectiveness of XAI techniques for encoder-based language models. *Knowledge-Based Systems*, 310, 113042.
- Mersha, M. A., Yigezu, M. G., Shakil, H., AlShami, A. K., Byun, S., & Kalita, J. (2025b). A unified framework with novel metrics for evaluating the effectiveness of xai techniques in LLMs. arXiv:2503.05050.
- Mersha, M. A., Yigezu, M. G., Tonja, A. L., Shakil, H., Iskandar, S., Kolesnikova, O., & Kalita, J. (2025c). Explainable AI: XAI-guided context-aware data augmentation. *Expert Systems with Applications*, (p. 289, 128364).
- Mitra, R., & Zulkernain, I. (2025). Music generation using deep learning and generative AI: A systematic review. *IEEE Access*, 13, 18079–18106.
- Moher, D., Liberati, A., Tetzlaff, J., & Altman, D. G. (2009). Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. *BMJ (Clinical research ed.)*, 96(5), 1695–1722.
- Nasri, L., Bagherzad, M., Gruss, R., Marucchi, N. S. W., Abrahams, A. S., & Ehsani, J. P. (2018). An investigation into online videos as a source of safety hazard reports. *Journal of safety research*, 65, 89–99.
- Nazir, A., Chakravarthy, T. K., Cecchini, D. A., Khajuria, R., Sharma, P., Mirik, A. T., Kocaman, V., & Talby, D. (2024). Langtest: A comprehensive evaluation library for custom LLM and NLP models. *Software Impacts*, 19, 100619.
- Nedungadi, P., Veena, G., Tang, K.-Y., Menon, R. R. K., & Raman, R. (2025). Ai techniques and applications for online social networks and media: Insights from bertopic modeling. *IEEE Access*, 13, 37389–37407.
- Niu, T., Xiong, C., Zhou, Y., & Yavuz, S. (2024). Parameter-efficient detoxification with contrastive decoding. In N. Soni, L. Flek, A. Sharma, D. Yang, S. Hooker, & H. A. Schwartz (Eds.), *Proceedings of the 1st human-centered large language modeling workshop* (pp. 30–40). TBD: ACL. <https://aclanthology.org/2024.hucllm-1.3>. <https://doi.org/10.18653/v1/2024.hucllm-1.3>
- Ognibene, D., Wilkens, R., Taibi, D., Hernández-Leo, D., Kruschwitz, U., Donabauer, G., Theophilou, E., Lomonaco, F., Bursic, S., Lobo, R. A., Sánchez-Reina, J. R., Scifo, L., Schwarze, V., Börsting, J., Hoppe, U., Aprin, F., Malzahn, N., Eimler, S. (2023). Challenging social media threats using collective well-being-aware recommendation algorithms and an educational virtual companion. *Frontiers in Artificial Intelligence*, 5, 654930.
- O'Neill, J., Subramanian, S., Lin, E., Satish, A., & Mugunthan, V. (2024). Guardformer: Guardrail instruction pretraining for efficient safeguarding. In *Neurips safe generative AI workshop 2024*.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A. et al. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35, 27730–27744.
- Pachinger, P., Hanbury, A., Neidhardt, J., & Planitzer, A. (2023). Toward disambiguating the definitions of abusive, offensive, toxic, and uncivil comments. In *Proceedings of the first workshop on cross-cultural considerations in NLP (c3NLP)* (pp. 107–113).
- Pan, K., & Zeng, Y. (2023). Do llms possess a personality? Making the mbti test an amazing evaluation for large language models. arXiv preprint arXiv:2307.16180.
- Pandiani, D. S. M., Sang, E. T. K., & Ceolin, D. (2025). 'Toxic'memes: A survey of computational perspectives on the detection and explanation of meme toxicities. *Online Social Networks and Media*, 47, 100317.
- Papageorgiou, E., Chronis, C., Varlamis, I., & Himeur, Y. (2024). A survey on the use of large language models (llms) in fake news. *Future Internet*, 16(8), 298.
- Pennycook, G., & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, 25(5), 388–402.
- Pöldvere, N., Uddin, Z., & Thomas, A. (2023). The politifact-oslo corpus: A new dataset for fake news analysis and detection. *Information*, 14(12), 627.
- Polemi, N., Praça, I., Kioski, K., & Bécue, A. (2024). Challenges and efforts in managing AI trustworthiness risks: A state of knowledge. *Frontiers in Big Data*, 7, 1381163.
- Przybyla, P., McGill, E., & Saggion, H. (2024). Know thine enemy: Adaptive attacks on misinformation detection using reinforcement learning. In *Proceedings of the 14th workshop on computational approaches to subjectivity, sentiment, & social media analysis* (pp. 125–140).
- Qiu, Y. (2024). Social comparison on social media platforms: A media and communication perspective. In *Shs web of conferences* (p. 03008). EDP Sciences (vol. 185).
- Raj, C., Agarwal, A., Bharathy, G., Narayan, B., & Prasad, M. (2021). Cyberbullying detection: Hybrid models based on machine learning and natural language processing techniques. *Electronics*, 10(22), 2810.
- Rao, H., Leung, C., & Miao, C. (2023). Can chatGPT assess human personalities? A general evaluation framework. In *Findings of the association for computational linguistics: EMNLP 2023* (pp. 1184–1194).
- Ray, G., McDermott, C. D., & Nicho, M. (2024). Cyberbullying on social media: Definitions, prevalence, and impact challenges. *Journal of Cybersecurity*, 10(1), tyae026.
- Rebedea, T., Dinu, R., Sreedhar, M. N., Parisien, C., & Cohen, J. (2023). Nemo guardrails: A toolkit for controllable and safe LLM applications with programmable rails. In *Proceedings of the 2023 conference on empirical methods in natural language processing: System demonstrations* (pp. 431–445).
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). " Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135–1144).
- Röttger, P., Vidgen, B., Nguyen, D., Waseem, Z., Margetts, H., & Pierrehumbert, J. (2021). Hatecheck: Functional tests for hate speech detection models. 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)* (pp. 41–58).
- Roy, D., & Kumar, V. (2024). An analysis of illegal wildlife trade with the aid of social media and prevention strategies. *Journal of Wildlife and Biodiversity*, 8(1), 386–401.
- Rozado, D. (2023). The political biases of chatGPT. *Social Sciences*, 12(3), 148.
- Saarinen, T. (2017). Toxic behavior in online games. Master's thesis. T. Saarinen.
- Saha, K., Chandrasekharan, E., & De Choudhury, M. (2019). Prevalence and psychological effects of hate speech in online college communities. In *Proceedings of the 10th ACM conference on web science* (pp. 255–264).
- Sakketou, F., Plepi, J., Cervero, R., Geiss, H. J., Rosso, P., & Flek, L. (2022). Factoid: A new dataset for identifying misinformation spreaders and political bias. In *Proceedings of the thirteenth language resources and evaluation conference* (pp. 3231–3241).
- Salminen, J., Sengün, S., Corporan, J., Jung, S.-g., & Jansen, B. J. (2020). Topic-driven toxicity: Exploring the relationship between online toxicity and news topics. *PLoS One*, 15(2), e0228723.
- Sandua, D. (2024). The double sides of technology: Internet addiction and its impact on today's society. David Sandua.
- Sap, M., Gabriel, S., Qin, L., Jurafsky, D., Smith, N. A., & Choi, Y. (2020a). Social bias frames: Reasoning about social and power implications of language. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 5477–5490).
- Sap, M., Gabriel, S., Qin, L., Jurafsky, D., Smith, N. A., & Choi, Y. (2020b). Social bias frames: Reasoning about social and power implications of language. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 5477–5490).
- Schäfer, J., Heid, U., & Klinger, R. (2024). Hierarchical adversarial correction to mitigate identity term bias in toxicity detection. In *Proceedings of the 14th workshop on computational approaches to subjectivity, sentiment, & social media analysis* (pp. 35–51).
- Scheuerlein, M. K., Jiang, J. A., Fiesler, C., & Brubaker, J. R. (2021). A framework of severity for harmful content online. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2), 1–33.
- Scheufele, D. A., & Krause, N. M. (2019). Science audiences, misinformation, and fake news. *Proceedings of the National Academy of Sciences*, 116(16), 7662–7669.
- Selvakumar, P., Sudheer, P., & Kannan, N. (2025). Balancing innovation and privacy: Understanding AI in the digital world. In *Digital citizenship and the future of AI engagement, ethics, and privacy* (pp. 279–304). IGI Global Scientific Publishing.
- Serapio-Garcia, G., Saifdari, M., Crepy, C., Sun, L., Fitz, S., Abdulhai, M., Faust, A., & Mataric, M. (2023). Personality traits in large language models. *arXiv e-prints*, arXiv-2307.
- Shahbaznezhad, H., Dolan, R., & Rashidirad, M. (2021). The role of social media content format and platform in users' engagement behavior. *Journal of Interactive Marketing*, 53(1), 47–65.
- Shaikh, O., Zhang, H., Held, W., Bernstein, M., & Yang, D. (2023). On second thought, let's not think step by step! bias and toxicity in zero-shot reasoning. In *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 4454–4470).
- Shakil, H., Farooq, A., & Kalita, J. (2024a). Abstractive text summarization: State of the art, challenges, and improvements. *Neurocomputing*, 603, 128255.
- Shakil, H., Mahi, A. M., Nguyen, P., Ortiz, Z., Kalita, J., & Mardini, M. T. (2024b). Evaluating text summaries generated by large language models using openAI's GPT. In *2024 international conference on machine learning and applications (ICMLA)* (pp. 951–956). IEEE.
- Shakil, H., Ortiz, Z., Forbes, G. C., & Kalita, J. (2024c). Utilizing GPT to enhance text summarization: A strategy to minimize hallucinations. *Procedia Computer Science*, 244, 238–247.

- Sharma, K. R., Chiu, W.-Y., & Meng, W. (2023). Security analysis on social media networks via STRIDE model. *ICNS 2023*, (p. 36).
- Shen, M., Tan, Z., Niyato, D., Liu, Y., Kang, J., Xiong, Z., Zhu, L., Wang, W., & Shen, X. (2024). Artificial intelligence for web 3.0: A comprehensive survey. *ACM Computing Surveys*, 56(10), 1–39.
- Shen, X., Wu, Y., Qu, Y., Backes, M., Zannettou, S., & Zhang, Y. (2025). Hatebench: Benchmarking hate speech detectors on LLM-generated content and hate campaigns. *Proceedings of 34th USENIX Security Symposium*, 221–240.
- Shu, K., Mahudeswaran, D., Wang, S., Lee, D., & Liu, H. (2020). Fakenewsnet: A data repository with news content, social context and spatiotemporal information for studying fake news on social media. *Journal on big data*, 8(3), 171–188.
- Si, W. M., Backes, M., Blackburn, J., De Cristofaro, E., Stringhini, G., Zannettou, S., & Zhang, Y. (2022). Why so toxic? Measuring and triggering toxic behavior in open-domain chatbots. In *Proceedings of the 2022 ACM SIGSAC conference on computer and communications security* (pp. 2659–2673).
- Siapera, E. (2022). Ai content moderation, racism and (de) coloniality. *International Journal of Bullying Prevention*, 4(1), 55–65.
- Siegelmann, R., Mehrabi, N., Goyal, P., Goyal, P., Bauer, L., Dhamala, J., Galstyan, A., Gupta, R., & Ghanadan, R. (2024). Mico: Preventative detoxification of large language models through inhibition control. In *Findings of the association for computational linguistics: NAACL 2024* (pp. 1696–1703).
- Singh, S., Gupta, A. K., Singh, S. K., Rastogi, N., & Singh, S. (2022). Technical challenges to automated detection of toxic language. In *International joint conference on advances in computational intelligence* (pp. 707–722). Springer.
- Singh, S., & Nambiar, V. (2024). Role of artificial intelligence in the prevention of online child sexual abuse: A systematic review of literature. *Journal of Applied Security Research*, 19(4), 586–627.
- Skiba, R. (2024). Shadows of catastrophe: Navigating modern suffering risks in a vulnerable society. After Midnight Publishing.
- Smith, E. M., Hall, M., Kambadur, M., Presani, E., & Williams, A. (2022). “I’m sorry to hear that”: Finding new biases in language models with a holistic descriptor dataset. In *Proceedings of the 2022 conference on empirical methods in natural language processing* (pp. 9180–9211).
- Soni, D., & Singh, V. K. (2018). See no evil, hear no evil: Audio-visual-textual cyberbullying detection. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), 1–26.
- Stahl, B. C. (2021). Artificial intelligence for a better future: An ecosystem perspective on the ethics of AI and emerging digital technologies. Springer Nature.
- Štrkalj Despot, K., Ostroški Anić, A., & Veale, T. (2023). Somewhere along your pedigree, a bitch got over the wall!—a data-driven approach to a typology of implicitly offensive language. *A Data-Driven Approach to a Typology of Implicitly Offensive Language*.
- Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. In *International conference on machine learning* (pp. 3319–3328). PMLR.
- Thomas, L., & Bhat, S. (2022). An overview of facebook’s journey to meta-a case study. *International Journal of Case Studies in Business, IT, and Education (IJCSE)*, 6(1), 268–287.
- Throuvala, M. A., Griffiths, M. D., Rennoldson, M., & Kuss, D. J. (2021). Perceived challenges and online harms from social media use on a severity continuum: A qualitative psychological stakeholder perspective. *International Journal of Environmental Research and Public Health*, 18(6), 3227.
- Tommoy, S. M. T. I., Zaman, S. M. M., Jain, V., Rani, A., Rawte, V., Chadha, A., & Das, A. (2024). A comprehensive survey of hallucination mitigation techniques in large language models. *CoRR*.
- Toraman, C., Şahinuç, F., & Yilmaz, E. (2022). Large-scale hate speech detection with cross-domain transfer. In *Proceedings of the thirteenth language resources and evaluation conference* (pp. 2215–2225).
- Tuong, W., Larsen, E. R., & Armstrong, A. W. (2014). Videos to influence: A systematic review of effectiveness of video-based education in modifying health behaviors. *Journal of Behavioral Medicine*, 37, 218–233.
- Twenge, J. M. (2020). Increases in depression, self-harm, and suicide among US adolescents after 2012 and links to technology use: Possible mechanisms. *Psychiatric Research and Clinical Practice*, 2(1), 19–25.
- Vaidya, A., Mai, F., & Ning, Y. (2020). Empirical analysis of multi-task learning for reducing identity bias in toxic comment detection. In *Proceedings of the international AAAI conference on web and social media* (pp. 683–693). (vol. 14).
- van Aken, B., Risch, J., Krestel, R., & Löser, A. (2018). Challenges for toxic comment classification: An in-depth error analysis. *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)* (pp. 33–42).
- Vayadande, K., Baviskar, A., Avhad, J., Bahadkar, S., Bhalerao, P., & Chimkar, A. (2024). A comprehensive review on navigating the web 3.0 landscape. In *2024 second international conference on inventive computing and informatics (ICICI)* (pp. 456–463). <https://doi.org/10.1109/ICICI62254.2024.000080>
- Verhoeven, I., Mishra, P., Beloch, R., Yannakoudakis, H., & Shutova, E. (2024). A (more) realistic evaluation setup for generalisation of community models on malicious content detection. In *Findings of the association for computational linguistics: NAACL 2024* (pp. 437–463).
- Vidgen, B., Thrush, T., Waseem, Z., & Kiela, D. (2021a). Learning from the worst: Dynamically generated datasets to improve online hate detection. 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)* (pp. 1667–1682).
- Vidgen, B., Thrush, T., Waseem, Z., & Kiela, D. (2021b). Learning from the worst: Dynamically generated datasets to improve online hate detection. 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)* (pp. 1667–1682).
- Vidgen, B., Thrush, T., Waseem, Z., & Kiela, D. (2021c). Learning from the worst: Dynamically generated datasets to improve online hate detection. 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)* (pp. 1667–1682).
- Villate-Castillo, G., Lorente, J. D. S., & Urquijo, B. S. (2024). A systematic review of toxicity in large language models: Definitions, datasets, detectors, detoxification methods and challenges.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.
- Wang, M., Zhang, N., Xu, Z., Xi, Z., Deng, S., Yao, Y., Zhang, Q., Yang, L., Wang, J., & Chen, H. (2024a). Detoxifying large language models via knowledge editing. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 3093–3118).
- Wang, M., Zhang, N., Xu, Z., Xi, Z., Deng, S., Yao, Y., Zhang, Q., Yang, L., Wang, J., & Chen, H. (2024b). Detoxifying large language models via knowledge editing. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 3093–3118).
- Wang, S., Li, R., Chen, X., Yuan, Y., Wong, D. F., & Yang, M. (2025). Exploring the impact of personality traits on LLM bias and toxicity. *arXiv preprint arXiv:2502.12566*.
- Wang, W., Gutt, D., & Li, T. (2024c). Balancing openness and control: The impact of entry barriers on user-generated content in DAO. *ICIS 2024 Proceedings, Sharing Economy, Platforms, and Crowds*, 20.
- Wang, Y. (2023). Synthetic realities in the digital age: Navigating the opportunities and challenges of ai-generated content. *Authorea Preprints*.
- Wang, Y., Li, H., Han, X., Nakov, P., & Baldwin, T. (2024d). Do-not-answer: Evaluating safeguards in LLMs. In *Findings of the association for computational linguistics: EACL 2024* (pp. 896–911).
- Wang, Y., Pan, Y., Yan, M., Su, Z., & Luan, T. H. (2023). A survey on chatGPT: AI-generated contents, challenges, and solutions. *IEEE Open Journal of the Computer Society*, 4, 280–302.
- Wang, Z., Yang, F., Wang, L., Zhao, P., Wang, H., Chen, L., Lin, Q., & Wong, K.-F. (2024e). Self-guard: Empower the llm to safeguard itself. In *Proceedings of the 2024 conference of the north american chapter of the association for computational linguistics: Human language technologies (volume 1: Long papers)* (pp. 1648–1668).
- Wardle, C., & Derakhshan, H. (2017). Information disorder: Toward an interdisciplinary framework for research and policymaking (vol. 27). Council of Europe Strasbourg.
- Webster, C. S., Taylor, S., Thomas, C., & Weller, J. M. (2022). Social bias, discrimination and inequity in healthcare: Mechanisms, implications and recommendations. *BJA Education*, 22(4), 131.
- Wei, A., Haghtalab, N., & Steinhardt, J. (2024). Jailbroken: How does llm safety training fail? *Advances in Neural Information Processing Systems*, 36, 80079–80110.
- Wei, X., Kumar, N., & Zhang, H. (2025). Addressing bias in generative AI: Challenges and research opportunities in information management. *Information & Management*, 62 (2), 104103.
- Wen, J., Ke, P., Sun, H., Zhang, Z., Li, C., Bai, J., & Huang, M. (2023). Unveiling the implicit toxicity in large language models. In *Proceedings of the 2023 conference on empirical methods in natural language processing* (pp. 1322–1338).
- Wiegand, M., Ruppenhofer, J., & Kleinbauer, T. (2019). Detection of abusive language: The problem of biased datasets. In *Proceedings of the 2019 conference of the north american chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)* (pp. 602–608).
- Wyrwoll, C. (2014). User-generated content. In *Social media: Fundamentals, models, and ranking of user-generated content* (pp. 11–45). Springer.
- Xu, J., Ju, D., Li, M., Bourreau, Y.-L., Weston, J., & Dinan, E. (2021). Bot-adversarial dialogue for safe conversational agents. In *Proceedings of the 2021 conference of the north american chapter of the association for computational linguistics: Human language technologies* (pp. 2950–2968).
- Xu, Z., Liu, Y., Deng, G., Li, Y., & Picek, S. (2024). A comprehensive study of jailbreak attack versus defense for large language models. In *Findings of the association for computational linguistics ACL 2024* (pp. 7432–7449).
- Yaprak, B. (2024). Generative artificial intelligence in marketing: The invisible danger of AI hallucinations. *Ekonimi İşletme ve Yönetim Dergisi*, 8(2), 133–158.
- Yousefi, M., & Emmanouilidou, D. (2021). Audio-based toxic language classification using self-attentive convolutional neural network. In *2021 29th European signal processing conference (EUSIPCO)* (pp. 11–15). IEEE.
- Yu, J., Zhang, X., Xu, Y., Lei, X., Yao, Z., Zhang, J., Hou, L., & Li, J. (2024). A cause-effect look at alleviating hallucination of knowledge-grounded dialogue generation. In *Proceedings of the 2024 joint international conference on computational linguistics, language resources and evaluation (LREC-COLING 2024)* (pp. 102–112).
- Yu, Z., Zhai, S., & Zhang, N. (2023). Antifake: Using adversarial audio to prevent unauthorized speech synthesis. In *Proceedings of the 2023 ACM SIGSAC conference on computer and communications security* (pp. 460–474).
- Zellers, R., Holtzman, A., Rashkin, H., Bisk, Y., Farhadi, A., Roesner, F., & Choi, Y. (2019). Defending against neural fake news. *Advances in Neural Information Processing Systems*, 32.
- Zhang, C., Sedoc, J., D’Haro, L. F., Banchs, R., & Rudnicky, A. (2022). Automatic evaluation and moderation of open-domain dialogue systems. *arXiv e-prints*, arXiv-2111.
- Zhang, R. (2024). The impact of the “regulations on the governance of online violence information” on user behavior in chinese social media platforms. *Law and Economy*, 3(11), 23–27.
- Zhang, Y. (2025). Adaptive detoxification: Safeguarding general capabilities of LLMs through toxicity-aware knowledge editing. In *Computer science undergraduate conference 2025@ XJTU*.

- Zheng, A., Rana, M., & Stolcke, A. (2025). Lightweight safety guardrails using fine-tuned BERT embeddings. In *Proceedings of the 31st international conference on computational linguistics: Industry track* (pp. 689–696).
- Zheng, J., Liu, X., Haque, M., Qian, X., Yang, G., & Yang, W. (2024). Hatemoderate: Testing hate speech detectors against content moderation policies. In *Findings of the association for computational linguistics: NAACL 2024* (pp. 2691–2710).
- Zhou, P., Wang, L., Liu, Z., Hao, Y., Hui, P., Tarkoma, S., & Kangasharju, J. (2024). A survey on generative AI and LLM for video generation, understanding, and streaming. *CoRR*.
- Zhukova, M., & Herring, S. C. (2024). Benign or toxic? differences in emoji interpretation by gender, generation, and emoji type. *Language@ Internet*, 22(Special Issue), 74–108.
- Zollo, F., & Quattrociocchi, W. (2018). Misinformation spreading on facebook. *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks*, (pp. 177–196).