

From Emotion to Expression: Theoretical Foundations and Resources for Fear Speech

Vigneshwaran Shankaran¹, Gabriella Lapesa^{1,2} and Claudia Wagner^{1,3,4}

¹GESIS - Leibniz Institute for the Social Sciences

²Heinrich-Heine University

³RWTH Aachen University

⁴Complexity Science Hub
first.last@gesis.org

Abstract

Few forces rival fear in their ability to mobilize societies, distort communication, and reshape collective behavior. In computational linguistics, fear is primarily studied as an emotion, but not as a distinct form of speech. Fear speech content is widespread and growing, and often outperforms hate-speech content in reach and engagement because it appears "civilier" and evades moderation. Yet the computational study of fear speech remains fragmented and under-resourced. This can be understood by recognizing that fear speech is a phenomenon shaped by contributions from multiple disciplines. In this paper, we bridge cross-disciplinary perspectives by comparing theories of fear from Psychology, Political science, Communication science, and Linguistics. Building on this, we review existing definitions. We follow up with a survey of datasets from related research areas and propose a taxonomy that consolidates different dimensions of fear for studying fear speech. By reviewing current datasets and defining core concepts, our work offers both theoretical and practical guidance for creating datasets and advancing fear speech research.

1 Introduction

Fear is a fundamental human emotion that shapes decision-making, risk perception, and social interactions. In fact, it is one of the most primal emotions, driving the fundamental "fight or flight" response in most living beings (Cannon, 1939). Fear plays a crucial role in driving precautionary actions, fostering group cohesion, and reinforcing ideological boundaries (Marshall et al., 2007). Unlike emotions such as happiness or surprise, fear has a different kind of profound influence on human behavior, particularly in digital spaces, where it affects engagement dynamics, online discourse, and information diffusion (Brady et al., 2017). Despite its significance, fear as a concept has various inter-

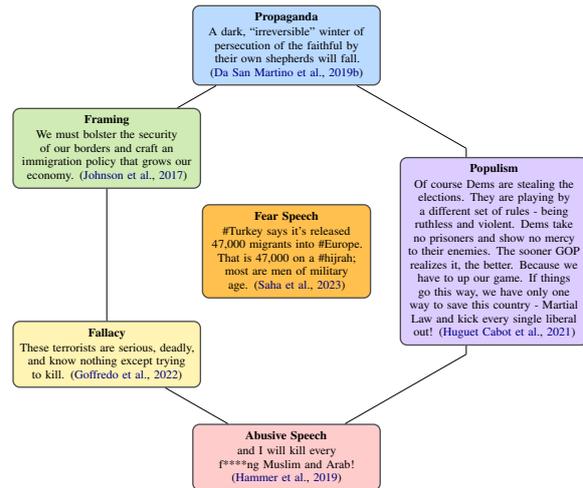


Figure 1: Fear-laden texts are prevalent and have been examined from multiple perspectives. All examples are sourced from datasets surveyed in this paper.

pretations across disciplines that study fear either for its property of being an emotion or its applications as a strategic tool. Regardless, this raises the question of its similarities and divergences and how it could be reconciled to be effectively able to capture the different aspects of it.

Fear lies at the core of the concept of *fear speech*, which has only recently gained traction in computational social science research. In light of the current political climate and the rising prominence of extremist discourse, studying fear speech is crucial, as it can persuade individuals to support or participate in extraordinary measures (Wagner and Morisi, 2019). Recent research found that on social media fear speech content is widespread and growing, often outperforming hate-speech content in reach/engagement (Saha et al., 2023).

Within the domain of NLP, work on computational argumentation has highlighted the persuasive power of emotionality (Benlamine et al., 2017) and its toxic potential (Ziegenbein et al., 2023), focused on emotional language in general and not on spe-

cific emotions. [Greschner and Klinger \(2025\)](#) and [Quensel et al. \(2025\)](#) make a further step annotating discrete emotions (including fear) in arguments and analyzing their impact on persuasion, respectively. Yet, fear has so far been studied as an emotion among others, and it has not yet been put under the interdisciplinary lens we propose in this paper.

We claim that to study fear speech effectively, theories and resources from multiple disciplines should be considered since fear speech is a cross-disciplinary concept. As often is the case in a data-hungry discipline like NLP, the bottleneck is the lack of datasets that encode specifically the *fear speech signal* (while plenty of annotation is available for fear as an emotion). This is why auditing existing datasets can help identify opportunities for part-reuse, i.e., when fear speech compatible notions have been annotated as a subset of a broader phenomenon, such as "*appeal to fear*" in propaganda research ([Da San Martino et al., 2019b](#)) or "*immigrant as a threat*" in framing research ([Mendelsohn et al., 2021](#)). This is precisely the goal of the dataset review presented in this paper, whose scope is illustrated in Figure 1: starting from the core notion of fear speech devised in the social sciences and encoded in small, high quality datasets (center of the figure), we identify a set of related phenomena (i.e., framing, propaganda, populism, etc.) whose annotations can potentially contribute to the investigation of fear speech, and we conduct a systematic review of datasets and related annotation guidelines. This approach not only leads to an interdisciplinary definition of fear speech, it is also sustainable because building new annotated resources is costly, so it is important to carefully evaluate what already exists.

Contributions In this paper, our contributions are threefold. **First**, we provide a comparative overview of how fear is conceptualized across disciplines such as Political science, Communication science, Psychology, and Linguistics. **Second**, we review existing definitions of fear speech and argue that one is particularly suitable and expand it in light of the comparative overview of fear. **Third**, we survey existing datasets that can be used to study fear speech and propose a taxonomy that consolidates and aggregates the different dimensions of fear. Our results show that fear-related constructs are unevenly represented across datasets, with strong concentration in computer science oriented disciplines and significant under representation of social science perspectives.

2 Theoretical Foundations of Fear

In this section we discuss fear through various disciplinary lenses. Since our primary interest concerns datasets (*which can potentially be employed for further modeling experiments*) - Political science, Communication science, Psychology and Linguistics offer the most practical foundation for data acquisition and method development.

2.1 Political Science - Securitization and Elite Rhetoric

In International Relations and Political Science, securitization theory from the Copenhagen School examines how political elites describe certain issues as existential threats that require emergency actions ([Buzan et al., 1998](#)). In this theory, a "securitizing actor," such as a leader, uses a speech act to present a referent object, like the nation or identity, as being in serious danger. When the public accepts this framing, extraordinary measures become justified. Fear, in this context, is not mainly a spontaneous emotional reaction from the public. Instead, it is seen as a discursive tool. Through rhetoric, political actors turn regular political issues into matters of security ([Balzacq, 2010](#)). This approach highlights the role of power and intentional manipulation. For example, after 9/11, U.S. leaders used fear of Islamist terrorism to expand executive powers. Securitization shows how defining something as a threat can create collective fear and bypass normal political discussion. Political science also studies how fear shapes electoral politics. Candidates and media often raise public concerns about crime, the economy, or outsiders to influence voters. Campaign ads with ominous visuals are a common example ([Wagner and Morisi, 2019](#)). In short, **Political science views fear as a tool used by elites**. It is created and spread through rhetoric and policy for strategic purposes.

2.2 Communication Science - Fear Appeals and Persuasion

In communication research, especially in health and risk communication, fear is often studied as a tool of persuasion. Early models like Protection Motivation Theory (PMT) ([Prentice-Dunn and Rogers, 1986](#)) and the Extended Parallel Process Model (EPPM) ([Witte, 1992](#)) explain how people respond to fear-based messages. PMT suggests that people make two types of judgments: how serious and likely a threat is, and whether they feel capa-

ble of handling it. EPPM builds on this idea. If a message makes people feel both threatened and able to act, they are more likely to take protective steps. If they feel threatened but helpless, they may deny the message or avoid it instead (Witte, 1992). Meta-analyses show that strong fear appeals with clear solutions are more likely to change attitudes and behavior than weak or unclear messages. However, if the message causes too much fear without offering ways to cope, it can have the opposite effect (Witte and Allen, 2000). These theories focus on individual reactions. They explain how people notice, interpret, and respond to fear in messages. Fear appeals are common in public health and advertising. For example, anti-smoking campaigns or vaccine messages often use fear to encourage healthy behavior (Witte and Allen, 2000). Overall, **in Communication science, fear is treated as a tool that could be used to drive either beneficial or harmful action.** It is not considered manipulative if it is used with care and includes helpful information.

2.3 Psychology - Emotion and Appraisal

Psychologists see fear as a basic emotion tied to survival. It arises when people sense danger and leads to physical changes like increased heart rate or readiness to flee (Cannon, 1939). According to appraisal theories, fear happens when someone judges a situation as harmful to their well-being or self-image. Fear can be a short-term state or a longer-term trait like anxiety (Smith et al., 1990). Research shows that fear and anxiety increase alertness. People become more cautious, seek out information, and pay closer attention to messages (Wagner and Morisi, 2019). Depending on the situation, fear can lead to helpful actions like problem-solving, or harmful responses like panic and avoidance. Communication science's PMT (theory) also draws on these ideas by including how people judge threats and their ability to handle them. Psychologists also study how people learn fear, such as through classical conditioning in phobias or PTSD (Post-Traumatic Stress Disorder) (Rachman, 1991). Terror management theory adds that fear of death can shape a person's beliefs and worldview (Greenberg et al., 1986). Overall, **Psychology focuses on how people experience and manage fear from the inside.** It often uses tools like anxiety scales or physiological tests to measure fear.

2.4 Linguistics - Metaphor, Discourse and Framing

Linguists and communication scholars study how fear is expressed through language and how the way we talk about issues can shape public opinion. In cognitive linguistics, abstract concepts like fear are often explained using metaphors or physical descriptions (Lakoff and Johnson, 1980). For example, in French, fear is sometimes described with bodily images like *"turning pale with fear"* (*pâlir de peur*) or as a growing force, such as *"the fear was growing"* (*la peur grandissait*) (Ströbel, 2015). These examples show how people think and talk about fear either as something internal or as something that acts on them. Discourse analysts look at how public language, such as in media or politics, presents certain topics in fearful terms. Phrases like *"flood of immigrants"* or *"war on drugs"* highlight threats while ignoring other views. Researchers analyze who uses such language and how listeners might understand it. This field shows that fear is not just a feeling. It is also built and shared through words, stories, and cultural references. **Linguists focus on how fear is communicated, what people say, how they say it, and what those choices reveal about social attitudes.**

2.5 Comparative Overview: Theories of Fear

The discipline-specific perspectives outlined in the previous sections can be combined into a multilevel overview for analyzing fear in public discourse.

- At the **Source level**, fear often originates with actors (leaders, media) who select language and frames to present events as threatening. Here, political and linguistic theories intersect: an elite's speech (securitizing act) or a media headline uses specific metaphors/frames to highlight danger.
- At the **Message level**, communication theory applies: the discourse (e.g., speech, article, social media post) is essentially a fear appeal, characterized by severity/susceptibility cues and possibly efficacy information. One can analyze messages for fear-inducing content (e.g., how many threat words, imagery) and assess whether coping information is provided.
- At the **Reception level**, psychological processes kick in: audiences interpret and respond to these fear messages based on personal traits (trait anxiety), perceived control, and group identities. For

Theoretical Approach	Core Constructs	Mechanisms	Contexts/Examples
Securitization (Political Science)	<i>Threat, Security, Referent Object</i>	Framing an issue as existential threat via speech acts; elicits public concern and legitimizes extraordinary measures.	National security, emergency politics, immigration debates (e.g. “war on terror”).
Fear Appeals / EPPM (Communication)	<i>Perceived Severity, Susceptibility, Efficacy</i>	Fear-inducing messages trigger threat appraisal; if high efficacy, leads to danger control (adaptive action); if low efficacy, leads to fear control (denial).	Health campaigns (anti-smoking, vaccines), public safety ads; advertising (product safety claims).
Protection Motivation (Psychology)	<i>Threat Appraisal (vulnerability/severity), Coping Appraisal (self/response efficacy)</i>	Individual cognition: perceiving a threat leads to motivated protective behavior if coping appraisal is favorable; otherwise anxiety or avoidance.	Health behavior change (exercise, diet), phobia treatment, climate change risk communication.
Emotional Appraisal (Psychology)	<i>Emotion (fear/anxiety), Arousal, Cognition</i>	Fear (a high-arousal negative emotion) emerges from appraisal of low control/danger; it increases vigilance and information seeking but can cause withdrawal or risk aversion.	Personal reactions to news (e.g. fear of crime), therapy (anxiety disorders), media effects on audiences.
Framing & Metaphor (Linguistics)	<i>Language, Metaphor, Semantics, Narrative Frame</i>	Linguistic expressions (metaphors like “fear spreads” or “wall of refugees”) shape how audiences conceptualize threats; discourse constructs who/what to fear.	Media/news discourse (framing immigration as a “wave”), political speeches using metaphor (“this virus is an enemy”), public debate narratives.

Table 1: Comparative Overview: Theories of Fear

instance, if a news story frames a disease as uncontrollable, fearful viewers may experience high anxiety and either panic or seek information (as theory suggests). Public opinion surveys or experiments can test how different fear frames affect people’s attitudes (consistent with political psychology research).

- Finally, at the **Outcome level**, fear’s impact can be behavioral or political: individuals may support policies (e.g. security measures) they perceive as protective (political science outcome), or adopt personal health behaviors (communication outcome), or they may retreat in denial if efficacy is missing.

In practical terms, this comparative overview suggests methodologies for empirical analysis: content analysis of speeches and social media can quantify fear-laden language (drawing on linguistics), discourse analysis can reveal securitizing moves (political science) and metaphoric patterns (linguis-

tics), and surveys/experiments can measure individuals’ fear responses and resulting behaviors (psychology/communication). For example, a study of political speeches might code for “security” frames and then survey listeners’ anxiety levels and policy support to test the whole pathway from discourse to cognition to behavior. Similarly, health communication researchers could analyze news articles or social media about a pandemic (fear frames vs efficacy) and track changes in public precautionary actions. Together, this overview treats fear in public discourse as a process: actors use language (conceptualized and measured by linguistics and political science) to generate emotional cues, which audiences process (per psychological and communication models) and then act upon. By mapping core constructs across disciplines (as in Table 1) and tracing them from source to outcome, researchers can capture both the social construction of fear and its psychological effects.

3 Contextualizing Fear Speech

Research on fear speech is fairly in its nascent stages and is yet to attract mainstream attention. As with any growing line of research there exist conflicting definitions of fear speech. We have identified three works that explicitly study fear speech and provide two similar yet distinct definitions.

- [Saha et al. \(2023\)](#) adopts the definition by [Buyse \(2014\)](#) from Human Rights studies where fear speech is defined as “an expression aimed at instilling (existential) fear of a target (ethnic or religious) group”.¹
- [Greipl et al. \(2024a\)](#) defines fear speech as “any deliberate communicative act that explicitly or implicitly portrays a particular entity, e.g., a group or an institution, as inherently and/or imminently harmful on a cultural, societal, or existential level.”

We integrate these two definitions with the broader theoretical perspectives on how fear is generated and utilized in discourse. This expanded view, drawing from disciplines like Political Science, Communication Science, Psychology and Linguistics, sees fear speech as a more complex strategic communicative process. Actors, such as political elites or media outlets, often employ specific linguistic tools like securitizing frames (“*war on terror*”) or evocative metaphors (“*flood of immigrants*”) to construct certain issues, groups, or events as existential threats. These messages frequently function as “fear appeals”, designed to influence public perception, legitimize extraordinary measures, or motivate particular behaviors. The ultimate impact of such speech is not uniform, often depending on factors like the audience’s psychological processing and whether the message includes information on how to cope with the presented threat (efficacy).

Building upon the definition from [Greipl et al. \(2024a\)](#) and [Saha et al. \(2023\)](#), we define **fear speech** as “a deliberate communicative act that, through the strategic use of language, rhetoric, and framing, explicitly or implicitly portrays a targeted entity (such as a group, institution, or referent object like the nation or an identity) as an inherent or imminent threat on a cultural, societal, or existential level. The primary aim of the communicative

¹We cite the latest work of Fear speech by Saha but we acknowledge that the definition was first introduced to the computational research community in ([Saha et al., 2021](#)).

act is to evoke fear or widespread concern in an audience, often to influence their attitudes, justify specific actions, or shape public policy.”

We break the definition into three components, which we will use in our dataset analysis:

- **C1:** Deliberate communicative acts that use strategic language, rhetoric, and framing.
- **C2:** Portrayal of a targeted entity as an inherent or imminent threat at cultural, societal, or existential levels.
- **C3:** Primary aim of evoking fear or concern to influence attitudes, justify actions, or shape policy.

Fear speech and linguistic theory We return to linguistic theory, this time with a specific focus on fear speech, motivated by our interest in advancing its study within computational linguistics and exploring how existing linguistic frameworks can inform future modeling approaches for fear speech. Fear speech can also be examined through the lens of linguistic theory, particularly Speech Act Theory ([Austin, 1975](#); [Searle, 1969](#)). Speech act theory offers a useful framework for distinguishing between the literal content of an utterance (locutionary force), the communicative intention of it (illocutionary force), and its potential impact on recipients (perlocutionary force). From this perspective, fear speech is best understood in terms of its illocutionary force, that is, the communicative intention to evoke fear rather than merely to describe a threatening situation (the literal meaning). From a measurement perspective, the literal meaning and the communicative intention of a fear speech instance are definitely less challenging than the perlocutionary effects. Whether fear is actually induced in the audience, are considerably more difficult to observe and can typically only be approximated through annotation or by analyzing reactions in subsequent discourse.

A further relevant linguistic dimension concerns presupposition. The presupposition of a statement is the knowledge that is assumed to be shared by the speaker and the receiver of the message ([Levinson, 1983](#); [Beaver, 1997](#)). Such assumption of shared beliefs can be a very powerful communication trick. Consider the use of *again* in the utterance “*Refugees are going to destroy our country again*”: it presupposes that refugees have already destroyed our country in the past without explicitly

arguing for it. Indeed, fear speech often relies on presupposition triggers like *again* that present contested assumptions as given, thereby reinforcing threat narratives implicitly. Prior work on political discourse has shown that such presuppositions play an important role in shaping audience interpretations, casting the use of presupposition triggers like *again* as a form of pragmatic framing (Yu, 2022).

Relation to Hate speech Fear speech and Hate speech may share certain surface characteristics but we conceptualize fear speech as a broader and analytically distinct phenomenon. In particular, fear speech is not inherently tied to hostility toward a **protected group**. For instance, public health campaigns such as anti-smoking or anti-drunk-driving messages frequently rely on fear appeals to influence behavior, yet they do not constitute hate speech. While some instances of hate speech may draw on fear-inducing narratives, this overlap alone is insufficient to justify treating the two phenomena within the same analytical framework. We refer to hate speech primarily because its research trajectory illustrates how sustained research attention can emerge in response to strong societal relevance. We argue that fear speech warrants similar attention, especially given the increasingly fear-driven nature of contemporary political and media discourse.

4 Resource Landscape for Fear Speech

To foster research on the topic of fear speech, we audit existing resources about it. We conduct a survey on the resource landscape of fear speech from a multi-disciplinary perspective. We use a novel AI-supported prompt based paper discovery method rather than traditional keyword-based literature search methods; this helps us to address the vocabulary mismatch that may hinder cross-disciplinary literature reviews (Schneider and Matthes, 2024).

4.1 Search Methodology

Paper Discovery To systematically discover candidate papers, we employ Allen AI's Ai2 Paper Finder (*Now part of Ai2 Asta*)². Since prompts are the important element of the process, our overall strategy for designing prompts expands the core concept of “*datasets for fear speech*” and breaking

²This tool decomposes prompts into relevant components, searches for papers, follows citation trails, evaluates results for relevance, conducts follow-up queries as needed, and then presents a curated list of papers (Allen Institute for Artificial Intelligence, 2025).

it into component parts, related concepts, and relevant academic contexts. We developed prompts starting from the concept of fear speech and broadened the discovery range by creating prompts for framing, propaganda, emotion appeal, Political science and Communication science. Using this strategy, we ended up with a total of 27 prompts, which together resulted in 810 candidate papers for further screening (Refer to Appendix A). To check if the tool was functioning well, we looked at the candidate paper pool to confirm that the results included 3 "core" fear speech papers and 9 prominent papers from propaganda and framing studies that we already identified as relevant using keyword search on Google Scholar.

Categorization Taxonomy Since we collect datasets from related phenomena, we come across datasets with different labels with similar meanings and definitions. To standardize them, we develop a taxonomy to consolidate and aggregate the different dimensions of fear (Refer to Appendix C). The taxonomy is developed to encompass resources that cover at least one dimension of fear from one of these phenomena: Fear Speech, Propaganda, Framing, Populism, Fallacy and Abusive Speech. However, given the heterogeneity of sources and contexts, we cannot assume that each label fully reflects all components described in the taxonomy. We therefore indicate, in parentheses, the specific components that might apply on a case-by-case basis. We omit datasets that are created for the purpose of emotion detection since that is not of interest to studying fear speech.

Paper Screening Once the candidate papers are identified, the next step is to identify the final set of papers with datasets that could be directly or indirectly used for fear speech research. To support this, we define two buckets of papers to structure our selection process:

- **B1-Fear Speech:** Resources that directly focus on identifying, labeling, or analyzing fear speech.
- **B2-Adjacent Phenomena:** Resources that do not explicitly address fear speech but are instrumental for studying fear through related phenomena like framing, propaganda, populism, fallacies and abusive speech which are the top contenders for encoding fear.

To systematically review papers and assign them to a suitable bucket, we develop guidelines that

align with the fear speech definition presented in Section 3 (Refer to Appendix D). To make the bucketing process robust, we use the components of fear speech. Bucket **B1** ideally aims to capture all three components directly, whereas bucket **B2** is intended to capture at least one of the components by focusing on adjacent phenomena that share similar emotional strategies. To streamline the paper selection process, we utilize ASReview, an active learning tool that supports researchers in narrowing down relevant papers.³ The paper selection was carried out by the first author. To ensure the robustness and quality of the process, we implemented a validation step. After identifying 30 resource-related papers, a stratified sample of 30 papers (including these and others reviewed up to that point) were independently evaluated by the second author and achieved 100 percent consensus. The stopping criterion for ASReview was set at **50 consecutive non-relevant papers**. In total, **317 papers** were reviewed, resulting in **40 papers** that support our research objectives on fear speech. However, during the time of information extraction, two papers were removed because we identified that they did not conform to our research objectives. A complete list of all the **37 resources** relevant to studying fear speech is provided in Appendix E.

4.2 Current Landscape

To map the current landscape of resources on fear-related discourse, we conducted a structured analysis of the existing datasets collected for this study. Our aim is to characterize how fear and related constructs are currently represented and operationalized across research communities. We frame our investigation around four research questions (RQs) that address disciplinary orientations, data practices, conceptualizations, and dataset reusability.

RQ1: What are the disciplinary orientations of existing studies on fear-related discourse? We first examine the disciplinary orientation of each paper. Each dataset was assigned to a research community based on the venue of publication, while for preprints we used the primary subject area indicated by the authors. Our corpus spans *11 distinct research communities*. Of the *37 datasets*

³ASReview leverages the titles and abstracts of candidate papers, incorporating real-time human input to train machine learning models continuously in the background. This iterative process, referred to by the tool's creators as a "researcher-in-the-loop" approach, enhances both efficiency and accuracy in identifying relevant literature (Van De Schoot et al., 2021).

we found, 27 of them are contributed by *Computational Linguistics (CL)* venues. The remaining communities - *Computer Science (CS)*, *Machine Learning (ML)*, *Artificial Intelligence (AI)*, *Information Retrieval (IR)*, *Web Mining (WM)*, *Social & Information Networks (SI)*, *Data Processing (DP)*, *Data Accessibility (DA)*, *Social Science (SocSci)* and *Political Science (PolSci)* each contribute one dataset. When aggregated, 32 out of 37 datasets can be traced to **computer science-related venues**, reflecting a strong disciplinary concentration and limited engagement from the social sciences.

RQ2: What are the primary data sources used to study fear-related discourse? We next coded the data sources underlying each dataset. As shown in Figure 2, *news media* and *Twitter* dominate the landscape, together accounting for 26 datasets. Interestingly, datasets explicitly designed to capture fear speech were instead sourced from personal or less-moderated communication channels such as *WhatsApp*, *Telegram*, and *Gab* – platforms often associated with private or ideologically charged discourse. All datasets contained textual data; two included images, two were audiovisual, and one was multi-modal. In terms of language coverage, *English* overwhelmingly dominates (23 datasets), followed by *German* (6), *Arabic* (5), and *Urdu* (2), with the remainder spanning *Hindi*, *Czech*, *French*, *Italian*, *Polish*, *Russian*, *North Macedonian*, *Portuguese*, *Slovak*, and *Norwegian* each represented once. Figure 3 illustrates the topical distribution across datasets. Among those with identifiable topics,⁴ immigration appears most frequently (10 datasets), followed by politics including US, Indian, and UK politics in 8 datasets. Other topics include *climate change*, *wars (Israel-Palestine and Russo-Ukrainian)*, *sports*, and *finance*, suggesting a moderate topical diversity despite strong thematic clustering around political and social conflict.

RQ3: How are fear-related phenomena conceptualized and labeled across datasets? We identified 19 distinct labels across all datasets and aligned them with our proposed taxonomy of fear dimensions (Figure 4). The most frequent conceptualizations correspond to propaganda and framing, underscoring the dominance of these perspectives in studying fear discourse. Notably, labels "*Appeal to fear*" and "*Appeal to fear/prejudice*" even

⁴12 datasets had no explicit topic information, either because it was not provided or because the data was source oriented (e.g., Reddit communities; curated news media lists).

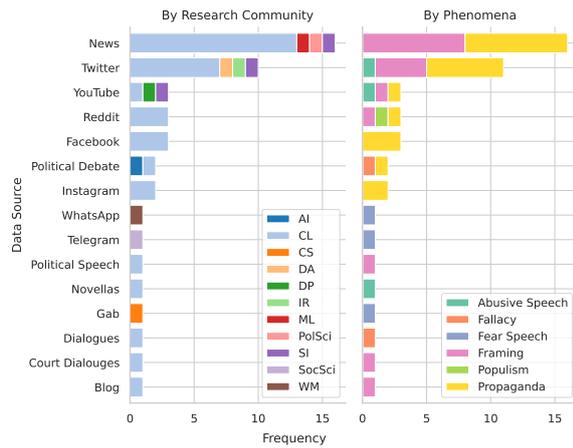


Figure 2: Distribution of research communities and phenomena by data source.

though a common propaganda technique has minor difference with the latter using emotional appeal to capitalize on people’s existing bias and stereotypes. Several fallacy focused datasets included the generic label “*Appeal to emotion*”, but were excluded in absence of fear-specific sub labels. As a result, *fallacy* and *populism* emerge as the most underrepresented phenomena in the current dataset landscape.

RQ4: To what extent are existing datasets accessible and reusable? Finally, to assess the potential for resource part reuse, we evaluated each dataset’s accessibility, annotation transparency, and licensing status. Of the 37 datasets, 26 provided some form of access mechanism most commonly via *direct link* (23), followed by *registration portal* (2) or *contact form* (1). However, only 21 of these links remained active at the time of our auditing, raising concerns about long term resource persistence. Moreover, only 13 datasets included annotation guidelines, reflecting limited documentation of labeling practices. Among the 33 manually annotated datasets, only 22 reported inter-annotator agreement scores and just 7 mention information of the annotators. Reusability is further constrained by incomplete reporting of corpus size. For 7 datasets, the number of fear related instances was not specified in the paper; for 4 others, it was determined manually by direct inspection of the data. These findings highlight significant variability in transparency and documentation practices, which impede reproducibility and reuse across the fear discourse research landscape.

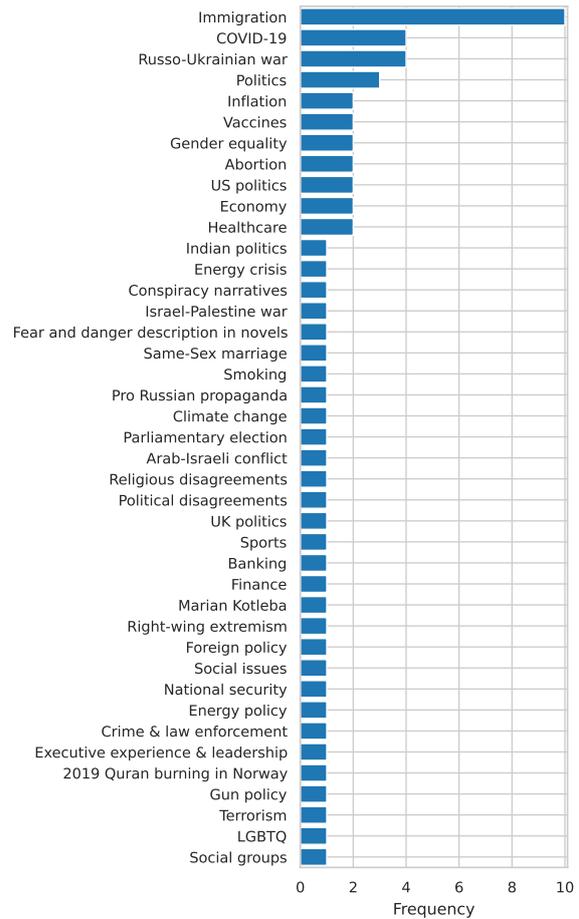


Figure 3: Distribution of topics across datasets

5 Discussion

The motivation behind this study lies in a simple observation that despite fear’s profound role in shaping public discourse, computational linguistics lacks a coherent direction or resource base for studying it. While hate speech and toxicity have inspired large-scale datasets and shared tasks, fear speech remains theoretically diffuse and empirically under-sourced. Our cross-disciplinary approach and resource mapping aim to establish fear as a valid construct for computational modeling grounded in theory yet operationally feasible.

From a computational perspective, this study provides a strategic starting point for modularizing fear. The comparative overview in Section 2 shows that fear manifests through identifiable linguistic and rhetorical cues such as metaphors, threat framing and emotional appeals which can be detected with existing NLP techniques. Future work can decompose fear into modular signals aligned with our definition components: strategic language, portrayal of threat and intent to evoke fear. This modu-

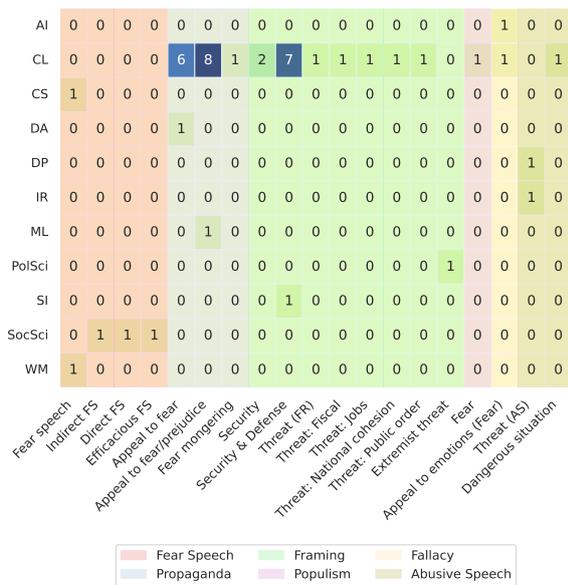


Figure 4: Distribution of fear related labels across research communities grouped by phenomena.

lar view enables researchers to reuse and combine methods like propaganda detection and framing while keeping theoretical coherence.

For the broader computational linguistics community, the value of our resource analysis lies in guidance. Understanding where fear-related signals already exist enables more targeted corpus design, whether through extending current resources, refining annotation schemes, or addressing gaps in underrepresented contexts such as everyday online discourse, local news, and crisis communication. Expanding linguistic coverage beyond English and German and improving annotation transparency through open guidelines and agreement reporting directly address the accessibility gaps identified in our audit. Our actively maintained repository⁵ continuously updated with new relevant datasets by the community, will serve as shared infrastructure for the emerging fear speech research.

Looking ahead, connecting theoretical insight with scalable NLP practice can advance fear speech as a distinct computational task. Future fear speech research can integrate modular signal detection with dynamic, community driven datasets that evolve alongside social discourse. Such an agenda not only advances fear speech as a computational task but also strengthens our capacity to study how emotions structure public communication at scale. Ultimately, our findings invite the computational

⁵<https://fearspeechdata.com/> (A successful pull request to this repository would expand the catalog of datasets)

linguistics community to view fear not only as a marginal emotional category among established phenomena, but also as a core analytic dimension.

6 Conclusion

This paper bridges theoretical perspectives and empirical resources to position fear speech as a distinct yet integrative phenomenon. By mapping how fear is conceptualized across disciplines and auditing datasets from adjacent research areas, we identify both the conceptual gaps and practical opportunities: fear is rarely modeled as a communicative construct, annotation practices are inconsistent, and existing resources remain concentrated within computational rather than interdisciplinary domains.

Our comparative overview shows that fear can be operationalized through modular linguistic and rhetorical signals. The resource landscape we present offers a foundation for informed, continuous, collaborative data development, encouraging the computational linguistics community to extend fear-relevant datasets as shared infrastructure. In doing so, we move towards a research timeline where fear is not treated as a marginal emotional category but as a core analytic dimension that connects emotion, social influence and linguistic form.

Limitations

Dual use While the goal of this work is to build a solid, theory-based empirical foundation for the understanding (and modeling) of fear speech, we cannot exclude that the our dataset survey may be used to develop systems which use fear to manipulate users.

Scope of the investigated phenomena Our overview of interdisciplinary phenomena and constructs is centered on aspects that have to do with the use of fear speech as a *"a deliberate communicative act"*. That is, we focus on the side of the *"producer"* of the fear speech and on the communicative strategies/motives that lead to its use. We acknowledge that a whole set of other constructs from social-psychology can in future work be integrated into the study of fear speech: these are phenomena that define which type of *"audience"* is more susceptible to fear speech. This could be the case for example for individuals: with a preference for values such as "security" (Human Value Theory, [Schwartz and Boehnke \(2004\)](#)); with a strong loading on moral foundations such as "care/harm" (Moral Foundation Theory, [Haidt and](#)

Joseph (2004)); cultural dimensions such as "uncertainty avoidance" (Hofstede's theory of cultural dimensions, Hofstede (2001)); or with a regulatory focus oriented to "prevention" (Regulatory Focus Theory, Tory Higgins (2015)).

Theoretical Scope and Exclusion of Methodological Analysis

This paper is about theories and datasets, and is not concerned with modeling choices or specific methods. For example, most political science, psychology and communication science works rely on qualitative screening, dictionary-based approaches and off-the-shelf tools like LIWC, while NLP research adopts models that are more powerful but increasingly less interpretable. While we acknowledge this to be a very important aspect, an overview of methods and approaches falls out of the scope of this paper and we reserve it for future work.

Scope of the investigated papers While we highlight a prevalence of computer-science contributions as dataset sources, we cannot exclude that this finding is driven by the fact that papers are labeled by venue and authors from social, political or communication science could be involved as authors.

AI-Assisted Literature Survey The literature review process was conducted across multiple disciplinary domains that lie outside our primary areas of expertise. As a result, conventional keyword-based search strategies were at times insufficient, as they rely heavily on prior familiarity with domain-specific terminology. This limitation became evident, for example, in our delayed identification of the concept of securitization, which only emerged after consultation with a political scientist. In such contexts, semantic and AI-assisted search methods can offer substantial advantages by enabling the discovery of relevant work that may not be captured through standard keyword queries (Schneider and Matthes, 2024). While a systematic evaluation of these methods against established literature review standards would constitute a valuable contribution to research on review methodologies, such an analysis lies beyond the scope of the present study. Accordingly, the use of AI-assisted search should be interpreted as a pragmatic support tool rather than a formally validated alternative to conventional systematic review procedures.

Acknowledgments

We would like to thank Luisa Kotthoff and Annika Herbertz for their support in the collection of meta data of the surveyed datasets.

References

- Firoj Alam, Hamdy Mubarak, Wajdi Zaghrouani, Giovanni Da San Martino, and Preslav Nakov. 2022. [Overview of the WANLP 2022 shared task on propaganda detection in Arabic](#). In *Proceedings of the Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 108–118, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Allen Institute for Artificial Intelligence. 2025. AI2 Paper Finder. <https://allenai.org/blog/paper-finder>. Accessed: 04.07.2025.
- Bodor Moheel Almotairy, Manal Abdullah, and Dimah Hussein Alahmadi. 2024. Dataset for detecting and characterizing arab computation propaganda on x. *Data in Brief*, 53:110089.
- Maaz Amjad, Alisa Zhila, Grigori Sidorov, Andrey Labunets, Sabur Butta, Hamza Imam Amjad, Oxana Vitman, and Alexander Gelbukh. 2022. Overview of abusive and threatening language detection in urdu at fire 2021. *arXiv preprint arXiv:2207.06710*.
- Arnav Arora, Srishti Yadav, Maria Antoniak, Serge Belongie, and Isabelle Augenstein. 2025. Multi-modal framing analysis of news. *arXiv preprint arXiv:2503.20960*.
- John Langshaw Austin. 1975. *How to do things with words*. Harvard university press.
- Vít Baisa, Ondřej Herman, and Ales Horak. 2019. [Benchmark dataset for propaganda detection in Czech newspaper texts](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 77–83, Varna, Bulgaria. INCOMA Ltd.
- Oana Balalau and Roxana Horincar. 2021. [From the stage to the audience: Propaganda on Reddit](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3540–3550, Online. Association for Computational Linguistics.
- Thierry Balzacq, editor. 2010. *Securitization Theory: How Security Problems Emerge and Dissolve*, 0 edition. Routledge.
- David Ian Beaver. 1997. Presupposition. In *Handbook of logic and language*, pages 939–1008. Elsevier.
- Mohamed S. Benlamine, Serena Villata, Ramla Ghali, Claude Frasson, Fabien Gandon, and Elena Cabrio. 2017. Persuasive argumentation and emotions:

- An empirical evaluation with users. In *Human-Computer Interaction. User Interface Design, Development and Multimodality*, pages 659–671, Cham. Springer International Publishing.
- William J Brady, Julian A Wills, John T Jost, Joshua A Tucker, and Jay J Van Bavel. 2017. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318.
- Antoine Buyse. 2014. Words of violence: "fear speech," or how violent conflict escalation relates to the freedom of expression. *Human Rights Quarterly*, 36(4):779–797.
- Barry Buzan, Ole Wæver, and Jaap De Wilde. 1998. *Security: A new framework for analysis*. Lynne Rienner Publishers.
- Walter Bradford Cannon. 1939. The wisdom of the body. *Psychological Bulletin*.
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. [The media frames corpus: Annotations of frames across issues](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.
- Giovanni Da San Martino, Alberto Barrón-Cedeño, and Preslav Nakov. 2019a. [Findings of the NLP4IF-2019 shared task on fine-grained propaganda detection](#). In *Proceedings of the Second Workshop on Natural Language Processing for Internet Freedom: Censorship, Disinformation, and Propaganda*, pages 162–170, Hong Kong, China. Association for Computational Linguistics.
- Giovanni Da San Martino, Alberto Barrón-Cedeño, Henning Wachsmuth, Rostislav Petrov, and Preslav Nakov. 2020. [SemEval-2020 task 11: Detection of propaganda techniques in news articles](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1377–1414, Barcelona (online). International Committee for Computational Linguistics.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019b. [Fine-grained analysis of propaganda in news articles](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5636–5646, Hong Kong, China. Association for Computational Linguistics.
- Agnese Daffara, Sourabh Dattawad, Sebastian Padó, and Tanise Ceron. 2025. Generalizability of media frames: Corpus creation and analysis across countries. *arXiv preprint arXiv:2506.16337*.
- Dimitar Dimitrov, Firoj Alam, Maram Hasanain, Abul Hasnat, Fabrizio Silvestri, Preslav Nakov, and Giovanni Da San Martino. 2024. [SemEval-2024 task 4: Multilingual detection of persuasion techniques in memes](#). In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 2009–2026, Mexico City, Mexico. Association for Computational Linguistics.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. [Detecting propaganda techniques in memes](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6603–6617, Online. Association for Computational Linguistics.
- Pierpaolo Goffredo, Shohreh Haddadan, Vorakit Vorakitphan, Elena Cabrio, and Serena Villata. 2022. Fallacious argument classification in political debates. In *Thirty-First International Joint Conference on Artificial Intelligence {IJCAI-22}*, pages 4143–4149. International Joint Conferences on Artificial Intelligence Organization.
- Jennifer Golbeck, Zahra Ashktorab, Rashad O Banjo, Alexandra Berlinger, Siddharth Bhagwan, Cody Buntain, Paul Cheakalos, Alicia A Geller, Rajesh Kumar Gnanasekaran, Raja Rajan Gunasekaran, and 1 others. 2017. A large labeled corpus for online harassment research. In *Proceedings of the 2017 ACM on web science conference*, pages 229–233.
- Jeff Greenberg, Tom Pyszczynski, and Sheldon Solomon. 1986. The causes and consequences of a need for self-esteem: A terror management theory. In *Public self and private self*, pages 189–212. Springer.
- Simon Greipl, Julian Hohner, Heidi Schulze, Patrick Schwabl, and Diana Rieger. 2024a. ["You are doomed!" Crisis-specific and Dynamic Use of Fear Speech in Protest and Extremist Radical Social Movements](#). *Journal of Quantitative Description: Digital Media*, 4.
- Simon Greipl, Julian Hohner, Heidi Schulze, Patrick Schwabl, and Diana Rieger. 2024b. ["you are doomed!" crisis-specific and dynamic use of fear speech in protest and extremist radical social movements](#). *Journal of Quantitative Description: Digital Media*, pages 1–46.
- Lynn Greschner and Roman Klinger. 2025. [Fearful falcons and angry llamas: Emotion category annotations of arguments by humans and LLMs](#). In *Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities*, pages 628–646, Albuquerque, USA. Association for Computational Linguistics.
- Jonathan Haidt and Craig Joseph. 2004. [Intuitive ethics: how innately prepared intuitions generate culturally variable virtues](#). *Daedalus*, 133(4):55–66.

- Hugo L Hammer, Michael A Riegler, Lilja Øvrelid, and Erik Velldal. 2019. Threat: A large annotated corpus for detection of violent threats. In *2019 International Conference on Content-Based Multimedia Indexing (CBMI)*, pages 1–5. IEEE.
- Maram Hasanain, Fatema Ahmad, and Firoj Alam. 2024. Can GPT-4 identify propaganda? annotation and detection of propaganda spans in news articles. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2724–2744, Torino, Italia. ELRA and ICCL.
- Maram Hasanain, Firoj Alam, Hamdy Mubarak, Samir Abdaljalil, Wajdi Zaghouni, Preslav Nakov, Giovanni Da San Martino, and Abed Freihat. 2023. ArAIEval shared task: Persuasion techniques and disinformation detection in Arabic text. In *Proceedings of ArabicNLP 2023*, pages 483–493, Singapore (Hybrid). Association for Computational Linguistics.
- Geert Hofstede. 2001. *Culture’s consequences: Comparing values, behaviors, institutions and organizations across nations*. Sage publications.
- Aleš Horák, Radoslav Sabol, Ondřej Herman, and Vít Baisa. 2024. Recognition of propaganda techniques in newspaper texts: Fusion of content and style analysis. *Expert Systems with Applications*, 251:124085.
- Pere-Lluís Hugué Cabot, David Abadi, Agneta Fischer, and Ekaterina Shutova. 2021. Us vs. them: A dataset of populist attitudes, news bias and emotions. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1921–1945, Online. Association for Computational Linguistics.
- Rahul Radhakrishnan Iyer and Katia Sycara. 2019. An unsupervised domain-independent framework for automated detection of persuasion tactics in text. *arXiv preprint arXiv:1912.06745*.
- Zhijing Jin, Abhinav Lalwani, Tejas Vaidhya, Xiaoyu Shen, Yiwen Ding, Zhiheng Lyu, Mrinmaya Sachan, Rada Mihalcea, and Bernhard Schoelkopf. 2022. Logical fallacy detection. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7180–7198, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Kristen Johnson, Di Jin, and Dan Goldwasser. 2017. Leveraging behavioral and social information for weakly supervised collective classification of political discourse on Twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 741–752, Vancouver, Canada. Association for Computational Linguistics.
- Anna-Katharina Jung, Gautam Kishore Shahi, Jennifer Fromm, Kari Anne Røysland, and Kim Henrik Gronert. 2024. Multi-platform framing analysis: A case study of kristiansand quran burning. In *Multi-disciplinary International Symposium on Disinformation in Open Online Media*, pages 101–130. Springer.
- Avneet Kaur and Arnav Arora. 2025. Beyond the battlefield: Framing analysis of media coverage in conflict reporting. *arXiv preprint arXiv:2506.10421*.
- Alena Kluknavská. 2015. A right-wing extremist or people’s protector? media coverage of extreme right leader marian kotleba in 2013 regional elections in slovakia. *Intersections. East European Journal of Society and Politics*, 1(1).
- George Lakoff and Mark Johnson. 1980. The metaphorical structure of the human conceptual system. *Cognitive science*, 4(2):195–208.
- Stephen C Levinson. 1983. *Pragmatics*. Cambridge university press.
- Abdurahman Maarouf, Dominik Bär, Dominique Geissler, and Stefan Feuerriegel. 2024. HQP: A human-annotated dataset for detecting online propaganda. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6064–6089, Bangkok, Thailand. Association for Computational Linguistics.
- Randall D Marshall, Richard A Bryant, Lawrence Amstel, Eun Jung Suh, Joan M Cook, and Yuval Neria. 2007. The psychology of ongoing threat: relative risk appraisal, the september 11 attacks, and terrorism-related fears. *American Psychologist*, 62(4):304.
- Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. Modeling framing in immigration discourse on social media. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263, Online. Association for Computational Linguistics.
- Wiktoria Mieszczewicz-Kowszewska, Beata Bajcar, Maciej Markiewicz, Jolanta Babiak, Berenika Dyczek, Przemysław Kazienko, and 1 others. 2025. Unraveling sitt: Social influence technique taxonomy and detection with llms. *arXiv preprint arXiv:2506.00061*.
- Benjamin Moffitt. 2015. How to perform crisis: A model for understanding the key role of crisis in contemporary populism. *Government and Opposition*, 50(2):189–217.
- Jakub Piskorski, Nicolas Stefanovitch, Nikolaos Nikolaidis, Giovanni Da San Martino, and Preslav Nakov. 2023. Multilingual multifaceted understanding of online news in terms of genre, framing, and persuasion techniques. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3001–3022, Toronto, Canada. Association for Computational Linguistics.
- Lavanya Prahallad and Radhika Mamidi. 2025. Analyzing biases in political dialogue: Tagging us presidential debates with an extended dams1 framework. *arXiv e-prints*, pages arXiv–2505.

- Steven Prentice-Dunn and Ronald W Rogers. 1986. Protection motivation theory and preventive health: Beyond the health belief model. *Health education research*, 1(3):153–161.
- Carlotta Quensel, Neele Falk, and Gabriella Lapesa. 2025. Investigating subjective factors of argument strength: Storytelling, emotions, and hedging. In *Proceedings of the 12th Argument mining Workshop*, pages 126–139, Vienna, Austria. Association for Computational Linguistics.
- S Rachman. 1991. Neo-conditioning and the classical theory of fear acquisition. *Clinical Psychology Review*, 11(2):155–173.
- Julia Romberg, Maximilian Maurer, Henning Wachsmuth, and Gabriella Lapesa. 2025. Towards a perspectivist turn in argument quality assessment. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7458–7485, Albuquerque, New Mexico. Association for Computational Linguistics.
- Punyajoy Saha, Kiran Garimella, Narla Komal Kalyan, Saurabh Kumar Pandey, Pauras Mangesh Meher, Binny Mathew, and Animesh Mukherjee. 2023. On the rise of fear speech in online social media. *Proceedings of the National Academy of Sciences*, 120(11):e2212270120.
- Punyajoy Saha, Binny Mathew, Kiran Garimella, and Animesh Mukherjee. 2021. “short is the road that leads from fear to hate”: Fear speech in indian whatsapp groups. In *Proceedings of the Web conference 2021*, pages 1110–1121.
- Muhammad Salman, Asif Hanif, Shady Shehata, and Preslav Nakov. 2023. Detecting propaganda techniques in code-switched social media text. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16794–16812, Singapore. Association for Computational Linguistics.
- Phillip Schneider and Florian Matthes. 2024. Conversational exploratory search of scholarly publications using knowledge graphs. In *Proceedings of the 7th International Conference on Natural Language and Speech Processing (ICNLSP 2024)*, pages 384–396, Trento. Association for Computational Linguistics.
- Shalom H Schwartz and Klaus Boehnke. 2004. Evaluating the structure of human values with confirmatory factor analysis. *Journal of Research in Personality*, 38(3):230–255.
- John R Searle. 1969. *Speech acts: An essay in the philosophy of language*. Cambridge university press.
- Craig A Smith, Richard S Lazarus, and 1 others. 1990. Emotion and adaptation. *Handbook of personality: Theory and research*, 21:609–637.
- Liane Ströbel. 2015. Linguistic realizations of the concept of fear.
- E. Tory Higgins. 2015. *Regulatory Focus Theory*, pages 1–18. John Wiley & Sons, Ltd.
- Rens Van De Schoot, Jonathan De Bruin, Raoul Schram, Parisa Zahedi, Jan De Boer, Felix Weijdema, Bianca Kramer, Martijn Huijts, Maarten Hoogerwerf, Gerbrich Ferdinands, and 1 others. 2021. An open source machine learning framework for efficient and transparent systematic reviews. *Nature machine intelligence*, 3(2):125–133.
- Prashanth Vijayaraghavan and Soroush Vosoughi. 2022. TWEETSPIN: Fine-grained propaganda detection in social media using multi-view representations. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3433–3448, Seattle, United States. Association for Computational Linguistics.
- Markus Wagner and Davide Morisi. 2019. Anxiety, fear, and political decision making. In *Oxford research encyclopedia of politics*.
- Kim Witte. 1992. Putting the fear back into fear appeals: The extended parallel process model. *Communications Monographs*, 59(4):329–349.
- Kim Witte and Mike Allen. 2000. A meta-analysis of fear appeals: Implications for effective public health campaigns. *Health education & behavior*, 27(5):591–615.
- Qi Yu. 2022. “again, dozens of refugees drowned”: A computational study of political framing evoked by presuppositions. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*, pages 31–43, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Qi Yu. 2023. Towards a more in-depth detection of political framing. In *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 162–174, Dubrovnik, Croatia. Association for Computational Linguistics.
- Qi Yu and Anselm Fliethmann. 2022. Frame detection in German political discourses: How far can we go without large-scale manual corpus annotation? *Journal for Language Technology and Computational Linguistics*, 35(2):15–31.
- Seunghak Yu, Giovanni Da San Martino, and Preslav Nakov. 2019. Experiments in detecting persuasion techniques in the news. *arXiv preprint arXiv:1911.06815*.
- Albin Zehe, Julian Schröter, and Andreas Hotho. 2023. Towards a computational analysis of suspense: Detecting dangerous situations. *arXiv preprint arXiv:2305.06818*.

Timon Ziegenbein, Shahbaz Syed, Felix Lange, Martin Potthast, and Henning Wachsmuth. 2023. [Modeling appropriate language in argumentation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4344–4363, Toronto, Canada. Association for Computational Linguistics.

A Prompts

Table 3 lists the prompts used to discover candidate papers.

B Information Extraction

We trained and employed two student assistants to extract metadata from the datasets used in the selected papers. Our procedure adapts the systematic literature database developed by (Romberg et al., 2025) to guide future research on fear speech. We began by collecting fundamental information about each dataset: the **dataset name**, and its **availability** (*online, upon request, upon registration, or -*). For datasets available online, we additionally recorded the **resource link** and its **status** (*accessible, not accessible or - as of September 2025*). We also noted the **license, annotation type** (*in-house, crowd-sourced, external, automatic*⁶), and, in the case of manual annotation, the **availability of annotation guidelines** (*Yes or -*). For the corresponding papers that introduced the datasets, we recorded the **title, authors, year of publication, and URL**, along with the targeted **research community** (based on the publication venue). At the dataset level, we documented the **size** in terms of **annotation units**, the **topic**, the **modality** (*text, image, audio-visual or multi-modal*), and the **language**. To align with our focus on fear speech, we included **label** as presented in the original papers and mapped these to the **dimension** from our *fear speech taxonomy codes* (Refer to Appendix C). For manually annotated datasets, we also tracked inter-annotator agreement through the reported **IAA score** and **IAA measure**.

C Taxonomy

Table 2 describes the taxonomy developed for categorizing fear dimensions.

⁶Some papers implement novel methods to label data and validate with a gold corpus. An example is in line with the recent advances and trends in CL research, prompting LLMs.

D Paper Screening Guidelines

Objective: Your task is to review a list of academic papers and classify each one into one of three categories:

- **B1:** Fear Speech
- **B2:** Adjacent Phenomena
- Discard

This classification is for a research project on identifying relevant datasets and resource for studying fear speech.

Definition: We define fear speech as a *deliberate communicative act that, through the strategic use of language, rhetoric, and framing, explicitly or implicitly portrays a targeted entity* (such as a group, institution, or referent object like the nation or an identity) *as an inherent or imminent threat on a cultural, societal, or existential level*. Its *primary aim is to evoke fear or widespread concern in an audience, often to influence their attitudes, justify specific actions, or shape public policy*.

Breaking it down into core components:

- **C1:** Deliberate communicative acts that use strategic language, rhetoric, and framing.
- **C2:** Portrayal of a targeted entity as an inherent or imminent threat at cultural, societal, or existential levels.
- **C3:** Primary aim of evoking fear or concern to influence attitudes, justify actions, or shape policy.

B1: Fear Speech

- **Focus:** Papers that directly study the concept of "fear speech" even if their definition of fear speech is not the one defined above.
- **Criteria:** The paper must explicitly mention "fear speech".

B2 - Adjacent Phenomena:

- **Focus:** Papers that don't use the term "fear speech" but focus on related phenomena like propaganda, framing, populism and fallacy.
- **Criteria:** The paper focuses on at least one of the components: **Component 1 (strategic language)**, **Component 2 (portrayal as threat)** and **3 (evoking fear)**. It studies the mechanisms that are used in fear speech, even if it labels them differently.

- **Think of it this way:** Could this paper’s dataset or method be used to analyze fear speech, even if the authors didn’t call it fear speech? If yes, it’s likely B2.

The Classification Process: A Step-by-Step Workflow

Step 1: Initial Scan : Read the Title, Keywords, and Abstract. This is often enough to infer if the paper is introducing a dataset. If you are not able to find the mention of “dataset” in the abstract, CTRL+F “dataset” and skim the highlighted marks area if it communicates the creation of a new dataset.

Step 2: B1 Test : Ctrl+F and search for “*fear speech*” and “*fearspeech*”. Now examine if the paper introduces a dataset and is explicitly labeled for fear speech. If this is not the case, its not in B1 bucket, proceed further.

Step 3: B2 Test : Ask these questions to determine if paper belongs to B2. Note: Emotion detection resources do not belong in this bucket and must be discarded.

1. Does the paper have one of the key words: “*fear*”, “*anxiety*”, “*threat*”, “*security*”, “*risk*”, “*crisis*”? Note: In case you find similar keywords that are not listed here and could be used to extend it, please note it down and we will resolve it during discussion session.
2. Does it come under one of the following phenomena: “*framing*”, “*propaganda*”, “*populism*”, “*fallacy*”, “*crisis communication*”, “*abusive speech*” ? Note: In case you find a similar phenomenon that is not listed here, please note it down and we will resolve it during discussion session.
3. Does the dataset contain one of the following labels or similar labels: “*appeal to fear*”, “*fear mongering*”, “*security*”. “*threat*”, “*crisis*”, “*appeal to emotion*”? Note: When there is a mention of Appeal to emotion as in the case of paper from Fallacy studies, choose them only if have “fear” has sub-category else discard.

Step 4: Discard : If the paper does not fit B1 or B2, classify it as Discard.

1. It could be **too broad** like “*hate speech*” or “*offensive language*” but **lacks any dimension of fear, anxiety or threat.**

2. It could be a **non-useful threat** (ex: virus, malware in cybersecurity research).
3. **Method-only studies** that reuse existing datasets without expanding on existing dataset or introducing new datasets or resources.
4. If it’s an **emotion detection resource.**

E List of Datasets

Table 4 lists the 37 datasets that we identified in our systematic literature search.

Dimension	Description	Components
Fear Speech	Fear speech is a deliberate communicative act that, through the strategic use of language, rhetoric, and framing, explicitly or implicitly portrays a targeted entity as an inherent or imminent threat on a cultural, societal, or existential level. Its primary aim is to evoke fear or widespread concern in an audience, often to influence their attitudes, justify specific actions, or shape public policy.	C1: Deliberate communicative acts that use strategic language, rhetoric, and framing. C2: Portrayal of a targeted entity as an inherent or imminent threat at cultural, societal, or existential levels. C3: Primary aim of evoking fear or concern to influence attitudes, justify actions, or shape policy.
Propaganda		
Appeal to fear	Seeking to build support for an idea by instilling anxiety and/or panic in the population. towards an alternative, possibly based on preconceived judgments. (Da San Martino et al., 2020)	C1: A deliberate rhetorical strategy to persuade an audience. (C2): It typically portrays a group, event, or idea as an imminent or existential danger.
Fear mongering	Is the text trying to appeal to fear, uncertainty or other threat? (Horák et al., 2024)	C1: A deliberate rhetorical strategy to persuade an audience. (C2): It typically portrays a group, event, or idea as an imminent or existential danger.
Framing		
Security	Used as a framing label to present issues (e.g., immigration, terrorism) as threats to safety, emphasizing protection or defense	C1: An intentional framing device. C2: By design it casts issues (for example migration) in terms of threats to safety or the social order. (C3): Primarily used to prioritize protective measures and shape policy preferences, they commonly aim to induce concern or support for defensive action.
Threat	Frames an entity, group, or event as posing danger and overlapping with security framing or as a broader construct of fear or crisis.	C1: An intentional framing device. C2: It explicitly depicts an actor or event as dangerous at cultural, societal, or existential levels. (C3): Primarily used to prioritize protective measures and shape policy preferences, they commonly aim to induce concern or support for defensive action.
Populism		
Crisis rhetoric	A communication strategy that identifies a "failure" or threatening event and elevates it to the level of an existential crisis. (Moffitt, 2015)	C1: It is an intentional rhetorical strategy aimed at mobilizing "the people". C2: It casts the situation as a systemic or existential breakdown, often locating blame in elites or out-groups and portraying those actors as threats to the people's way of life. (C3): It seeks to evoke fear, anger, or urgency among the audience to legitimize the populist leader's proposed solutions or policies.
Fallacy		
Appeal to emotion	Appeal to Emotion Manipulation of the recipient's emotions in order to win an argument. (Jin et al., 2022)	C1: A deliberate rhetorical strategy to persuade an audience. (C2): It typically portrays a group, event, or idea as an imminent or existential danger. This applies only if the emotion appealed to is fear.
Abusive Speech		
Threat	These have language intended to make the target or a broader group fearful or to feel unsafe. The threats may be personal or general. Maybe explicit or they may generally make the target feel unsafe without a specific threat of direct action. (Golbeck et al., 2017)	C3: Its immediate objective is to make the target or targets fearful and unsafe.

Table 2: Taxonomy for dimensions of fear.

Fear Speech	
1	Papers that present datasets for detecting fear speech.
2	Papers that present datasets for detecting deliberate communicative act that, through the strategic use of language, rhetoric, and framing, explicitly or implicitly portrays a targeted entity (such as a group, institution, or referent object like the nation or an identity) as an inherent or imminent threat on a cultural, societal, or existential level and its primary aim is to evoke fear or widespread concern in an audience, often to influence their attitudes, justify specific actions, or shape public policy.
3	Papers that present datasets for detecting deliberate communicative acts that use strategic language, rhetoric, and framing.
4	Papers that present datasets for detecting portrayal of a targeted entity as an inherent or imminent threat at cultural, societal, or existential levels.
5	Papers that present datasets for detecting text with the primary aim of evoking fear or concern to influence attitudes, justify actions, or shape policy.
Framing	
6-9	Papers that present datasets for detecting threat-based framing strategies (<i>in [political speech, media coverage, online discourse]</i>), where framing is defined as the process of selecting certain aspects of a perceived reality and making them more salient in a communicating text, in order to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation.
10-13	Papers that present datasets for detecting strategic framing used to depict existential or cultural threats (<i>in [political text, media coverage, online discourse]</i>), where framing is defined as the process of selecting certain aspects of a perceived reality and making them more salient in a communicating text, in order to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation.
14-17	Papers that present datasets for detecting framing devices that promote fear (<i>in [political speech, media coverage, online discourse]</i>), where framing is defined as the process of selecting certain aspects of a perceived reality and making them more salient in a communicating text, in order to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation.
Propaganda	
18	Papers that present datasets for detecting fear-based propaganda.
19	Papers that present datasets for detecting propaganda techniques, specifically "appeal to fear," "threat invocation," or "whataboutism."
20	Papers that present datasets for identifying propaganda techniques used to portray an out-group as a danger.
Emotion Appeal	
21	Papers that present datasets for detecting manipulative language strategies that evoke fear or anxiety.
22	Papers that present datasets for detecting logical fallacies driven by emotion like appeal to fear.
Political Science	
23	Papers on securitization theory that use computational methods to analyze speech acts that frame issues as existential threats.
24	Papers with datasets of parliamentary debates or policy documents annotated for securitizing moves and referent object construction.
Communication Science	
25	Papers that present datasets for detecting fear-inducing narratives in mass media reporting.
26	Papers that present datasets for detecting fear-based rhetoric during crises.
27	Papers that present datasets for analyzing emergency communication strategies that evoke fear or urgency.

Table 3: Prompts used to identify candidate papers through Allen AI's AI2 Paper Finder. For the Framing category, each row includes three prompt options, with the values in brackets indicating the chosen option.

Paper	Community	Label	Dimension	Topic	Data Source
Saha et al. (2021)	WM	Fear speech	FS:FS	Indian politics	WhatsApp
Saha et al. (2023)	CS	Fear speech	FS:FS	-	Gab
Greipl et al. (2024b)	SocSci	Indirect FS, Direct FS, Efficacious FS	FS:FS	COVID-19, Energy Crisis, Inflation, Immigration, Russo-Ukrainian war, Conspiracy Narratives	Telegram
Mendelsohn et al. (2021)	CL	Threat: Jobs, Threat: Public order, Threat: National cohesion, Threat: Fiscal, Security & Defense	FR:SE, FR:TH	Immigration	Twitter
Arora et al. (2025)	CL	Security & Defense	FR:SE	Immigration	News
Huguet Cabot et al. (2021)	CL	Fear	PO:CR	Social groups (Immigrants, Refugees, Muslims, Jews, Liberals, Conservatives)	Reddit
Kaur and Arora (2025)	CL	Security & Defense	FR:SE	Israel-Palestine war	News
Yu and Fliethmann (2022)	CL	Security	FR:SE	Immigration	News
Yu (2023)	CL	Security	FR:SE	Immigration	News
Zehe et al. (2023)	CL	Dangerous situation	AS:TH	Fear and danger description in novels	Novellas
Card et al. (2015)	CL	Security & Defense	FR:SE	Immigration, Same-Sex Marriage, Smoking	News
Salman et al. (2023)	CL	Appeal to fear/prejudice	PR:AF	-	Twitter, Facebook, Instagram, YouTube
Dimitrov et al. (2021)	CL	Appeal to fear/prejudice	PR:AF	Vaccines, Politics, COVID-19, Gender Equality	Facebook
Hasanain et al. (2024)	CL	Appeal to fear/prejudice	PR:AF	-	News

Continued on next page

Table 4 – continued from previous page

Paper	Community	Label	Dimension	Topic	Data Source
Da San Martino et al. (2019b)	CL	Appeal to fear/prejudice	PR:AF	-	News
Vijayaraghavan and Vosoughi (2022)	CL	Appeal to fear/prejudice	PR:AF	-	Twitter
Yu et al. (2019)	ML	Appeal to fear/prejudice	PR:AF	-	News
Alam et al. (2022)	CL	Appeal to fear/prejudice	PR:AF	-	Twitter
Baisa et al. (2019)	CL	Fear mongering	PR:FM	Pro Russian Propaganda	News
Piskorski et al. (2023)	CL	Security & Defense, Appeal to fear/prejudice	PR:AF	COVID-19, Abortion, Immigration, Russo-Ukrainian war, Climate Change, Parliamentary Election	News
Hammer et al. (2019)	DP	Threat	AS:TH	Arab-Israeli conflict, Religious Disagreements, Political Disagreements	YouTube
Hasanain et al. (2023)	CL	Appeal to fear	PR:AF	-	Twitter, News
Dimitrov et al. (2024)	CL	Appeal to fear/prejudice	PR:AF	Politics, Vaccines, COVID-19, Gender Equality, Russo-Ukrainian war	Facebook, Instagram
Amjad et al. (2022)	IR	Threat	AS:TH	-	Twitter
Da San Martino et al. (2019a)	CL	Appeal to fear	PR:AF	-	News
Da San Martino et al. (2020)	CL	Appeal to fear	PR:AF	-	News
Iyer and Sycara (2019)	CL	Threat	FR:TH	US politics	Reddit, Court Dialogues, Political Speech, Blog
Balalau and Horincar (2021)	CL	Appeal to fear	PR:AF	UK politics, US politics	Reddit

Continued on next page

Table 4 – continued from previous page

Paper	Community	Label	Dimension	Topic	Data Source
Maarouf et al. (2024)	CL	Appeal to fear	PR:AF	Russo-Ukrainian war	Twitter
Almotairy et al. (2024)	DA	Appeal to fear	PR:AF	Sports, Banking	Twitter
Daffara et al. (2025)	CL	Security & Defense	FR:SE	Politics, Economy, Finance	News
Mieleszczenko-Kowszewicz et al. (2025)	CL	Appeal to emotions (fear)	FA:AE	-	Dialogues
Kluknavská (2015)	PolSci	Extremist threat	FR:TH	Marian Kotleba, Right-wing Extremism	News
Goffredo et al. (2022)	AI	Appeal to emotions (fear)	FA:AE	Foreign policy, Economy, Immigration, Healthcare, Social issues	Political Debate
Prahallad and Mamidi (2025)	CL	Appeal to fear	PR:AF	Immigration, Inflation, National Security, Energy policy, Crime & law enforcement, Executive Experience & Leadership	Political Debate
Jung et al. (2024)	SI	Security & Defense	FR:SE	2019 Quran burning in Norway	Twitter, YouTube, News
Johnson et al. (2017)	CL	Security & Defense	FR:SE	Abortion, Healthcare, Gun policy, Immigration, Terrorism, LGBTQ	Twitter

Table 4: Overview of datasets, including the research community targeted (*CL* - Computational Linguistics, *CS* - Computer Science, *ML* - Machine Learning, *AI* - Artificial Intelligence, *IR* - Information Retrieval, *WM* - Web Mining, *SI* - Social & Information Networks, *DP* - Data Processing, *DA* - Data Accessibility, *SocSci* - Social Science, *PolSci* - Political Science), the label(s) in the dataset, the assigned dimension of fear according to our taxonomy (*FS:FS* - Fear speech, *PR:AP* - Appeal to fear in Propaganda, *PR:FM* - Fear mongering in Propaganda, *FR:SE* - Security framing, *FR:TH* - Threat framing, *PO:CR* - Crisis rhetoric in populism, *FA:AE* - Appeal to emotion in Fallacy, *AS:TH* - Threat in Abusive speech), topic of the dataset and the data source