Media talks Privacy: Unraveling a Decade of Privacy Discourse around the World

Shujaat Mirza
New York University
New York, USA
shujaat.mirza@nyu.edu

Corban Villa
New York University Abu Dhabi
Abu Dhabi, UAE
corban.villa@nyu.edu

Christina Pöpper
New York University Abu Dhabi
Abu Dhabi, UAE
christina.poepper@nyu.edu

ABSTRACT

Our increasingly digital world has heightened concerns about privacy. Newspaper and media reporting influences and shapes public opinion, which impacts the strategic and operational decisions of a variety of stakeholders, making it crucial to understand how privacy-related issues are portrayed in the media. Leveraging time-series analysis, topic modeling, and sentiment analysis, this paper presents a comprehensive study on the coverage of privacy-related issues in newspapers from 2010 to 2022 across six regions of the world. Temporal trends in privacy coverage reveal a gradual increase in attention to privacy issues globally, with a notable surge observed in newspapers from the Global South, complementing the historically prominent Global North coverage. Topic modeling uncovers dominant themes in privacy reporting, revealing shifts in media focus from government surveillance to data breaches and tech corporations' role. Notably, the majority of privacy reporting carries a negative sentiment, emphasizing the widespread unease that pervades discussions surrounding privacy matters.

1 INTRODUCTION

Media narratives, as reflected in extensive coverage over time, can be used as a proxy for public perception, providing a unique window into the prevailing sentiments and concerns of society. Furthermore, as informed by the agenda-setting theory [24], the influence of news media transcends mere reflection of public opinion. It actively shapes and molds public agendas, steering the societal discourse on privacy. This dual role of the media — as both a mirror and a shaper of public sentiment — underscores the value of analyzing privacy-related reporting over the years. By examining how privacy issues are portrayed in the media, we aim to uncover trends and shifts in the narrative that mirror and potentially influence societal attitudes and policies. This approach is crucial for understanding the evolution of public sentiment in response to technological advancements and legislative developments, providing key insights for stakeholders in shaping future strategies and policies.

Prior research on privacy in media has often been constrained, typically concentrating on single incidents or limited to coverage from a few newspapers or countries, predominantly in the Global North. Our study addresses this by uncovering a diverse array of privacy incidents reported in media from a wide range of countries across different regions. Our dataset comprises 35,655 articles on privacy, collected from 36 newspapers spanning 25 countries across 6 geographic regions, from 2010 to 2022. Notably, our dataset maintains a balance between newspapers from the Global North and Global South, offering a more comprehensive, global perspective.

Our analysis focused on the privacy coverage over this 13-year period, yielding both geographic and temporal insights. This investigation revealed significant variations and spikes in privacy reporting, influenced by major events and stakeholders. Events such as the *PlayStation Network hack* (2011) underscored the importance of security protocols, while the *Snowden Revelations* (2013) [20] shed light on the extent of government surveillance, and high-profile court cases like the EU Court of Justice's *Right to be Forgotten ruling* (2014)[1] spotlighted the judiciary's role. The *Cambridge Analytica scandal* (2018) [21] brought attention to the misuse of data by corporations, and legislation such as the EU's *GDPR* (2016)[2] and India's *DPDP bill* (2022)[3] reinforced the need for legislative oversight. A particularly striking finding from our analysis is the marked increase in privacy coverage in the Global South, complementing the historical predominance of the Global North.

Next, to understand the relationship between coverage and topics, we apply an unsupervised topic model (latent Dirichlet allocation/LDA). This analysis revealed that tracking of users and online abuse remained consistent topics of focus throughout the last decade. Notably, the narrative within the privacy discourse evolved over time: While government surveillance was prominent in the early 2010s, attention shifted toward data breach scandals and subsequent investigations in the latter part of the decade. Additionally, our study uncovered regional disparities in privacy coverage. Newspapers from the Global North tended to focus on data scandals and investigations, whereas those from the Global South centered more on court rulings and user rights. This thematic exploration also highlighted the frequent presence of major stakeholders in privacy reporting, including governments, courts, big tech companies, legislators, corporations, and end-users.

Finally, recognizing that emotionally charged texts can influence readers' perception [13, 49], we complement our understanding of reporting patterns by also analyzing the sentiment and emotional tone expressed in each article. Utilizing IBM's Watson Natural Language Understanding (NLU) API, we analyzed the emotional nuances within each article. Our findings reveal a pronounced negative sentiment in privacy coverage, reflecting escalating public concerns over privacy issues. We observe a surge in emotionality

Proceedings on Privacy Enhancing Technologies 2024(4), 102–122

https://doi.org/10.56553/popets-2024-0109

[1] Google Inc. v Mario Costeja González, 2014, https://curia.europa.eu/jcms/upload/docs/application/pdf/2014-05/cp140070en.pdf
[2] EU's General Data Protection Regulation, 2016, https://eur-lex.europa.eu/eli/reg/2016/679/oj
[3] India's Digital Personal Data Protection Bill, 2022, https://tinyurl.com/yckbc8cn

for major privacy incident investigations, with government surveillance primarily evoking fear and online abuse inciting disgust, highlighting the distinct emotional responses elicited by different privacy-related topics.

In short, the major contributions of this work are:

(1) We present the first global longitudinal and comparative study of privacy-related reporting in online newspapers. We assess and discuss patterns of media reporting that may contribute to increased public awareness or spur legislative proactivity on privacy issues over the last decade.
(2) We addressed the lack of comprehensive privacy datasets by compiling and analyzing a multilingual dataset from 36 newspapers in 25 countries, and conducted a user study to validate translation accuracy. Our study highlights the global evolution of privacy concerns, integrating the Global South's experiences and activism, thereby challenging the Western-centric privacy narrative.
(3) Employing topic modeling, our study tracks a shift in media emphasis from government surveillance to data breaches, and further into the complex terrain of online abuse, highlighting a significant expansion in the scope and depth of privacy concerns.
(4) The negative sentiment dominating privacy media coverage highlights a public trust crisis, necessitating greater transparency and accountability from data custodians to restore and enhance public confidence.

2 PRELIMINARIES

In this section, we present the terminology used in the present study and the research objectives we set.

2.1 Definition of Privacy

The concept of privacy is multifaceted and can be understood differently depending on the context. In this research, we draw upon two comprehensive taxonomies of privacy [2, 42] to shape our understanding and analysis.

Stages of data life cycle. In his influential taxonomy of Privacy [42], Solove categorizes privacy issues into four main groups: information collection, information processing, information dissemination, and invasions. These categories further encompass 16 privacy-related activities that include surveillance, identification, data aggregation, and others. Solove's taxonomy allows us to dissect complex privacy issues and understand how they're portrayed in media. Solove's taxonomy—designed to serve as a guide for the development of privacy legislation, hence broad in its applicability—discusses a vast selection of privacy harms but comes short in addressing types of responses to and preventive measures for such incidents. Hence, we further angle the discussion in terms of attacks and defenses to data privacy, regardless of the life cycle stage.

Data privacy endeavors. As a basic human need or right, privacy can be preserved or exploited. Recent legislative and technical privacy-enhancing developments focus on protective measures that may be covered by the media to raise awareness and empower its readers. In their taxonomy, Antón and Earp distinguish between privacy requirements depending on whether they prevent or contribute to privacy harms [2]. The taxonomy identifies seven main categories of privacy concerns, including notice and awareness, choice and consent, and security, among others. This taxonomy helps us evaluate media coverage of privacy issues in the context of online practices and regulatory compliance.

These taxonomies provide a comprehensive framework to understand and categorize privacy issues, guiding our analysis of newspaper coverage on privacy-related topics.

2.2 Research Objectives

We aim to explore the evolving landscape of privacy-related news coverage, encompassing four key dimensions:

i) *Temporal trends*: We aim to identify how the coverage of privacy-related issues has changed over time across newspapers from varying regions. The intention is to understand the potential impact of key events and legislative changes.
ii) *Dominant themes in privacy reporting*: We aim to uncover dominant themes and observe their shift over the past decade, allowing us to recognize which topics have gained or lost prominence over the years.
iii) *Sentiment injected in privacy coverage*: We aim to investigate the tone of privacy-related articles, and how it varies across different themes and regions, revealing much about the framing of privacy issues in the public discourse.
iv) *Main stakeholders featured in privacy-related news*: By examining entities like governments, corporations, and courts in articles, we aim to understand their portrayed roles—whether as enforcers, violators, or victims of privacy practices.

These objectives guide our subsequent analysis and discussions, establishing a structured framework for this study.

3 RELATED WORK

Privacy remains a critical concern as our world becomes increasingly digitized [17, 25, 26, 30, 37]. Studies on media coverage of privacy-related endeavors have analyzed media reporting of major events such as the Snowden revelations [4, 11, 18, 45, 48] or general national or cross-national coverage of issues concerning digital privacy [9, 33, 36, 40, 44]. Works investigating the privacy-related news landscape employed frame and sentiment analyses to understand reporting patterns. Frame analysis seeks to uncover how news sources, most commonly, construct their discourse on issues of wide interest. Teutsch and Niemann [44] explore how German newspapers portray privacy in social network sites. The authors find that the amount of coverage over a period of seven years varies for the different identified frames and across local and national outlets. Kuehn [18] analyzes New Zealand's news reporting of the Snowden revelations from both a frame and sentiment point of view, and reveals that the majority (51%) of articles express a negative tone towards surveillance. Sheshadri et al. [40] compared privacy reporting in The New York Times and The Guardian with coverage of human suffering events, noting a more negative tone in privacy news. Our study broadens the scope in terms of timeline and geographical distribution.

Research leveraging priming theory reveals that privacy news consumption can heighten privacy concerns and literacy while diminishing trust in data institutions [27]. A study categorizing security and privacy (S&P) news into types such as financial and corporate data breaches, and politicized cybersecurity, finds distinct patterns in public sharing and awareness, influenced by factors including age and gender [8]. In examining privacy perceptions, it is notable that news about government surveillance can increase concerns about intrusion while lowering self-efficacy in privacy, sometimes even leading to weaker passwords [22]. Through a combination of topic modeling, temporal analysis, and regional distribution, our study aims to uncover how privacy reporting has evolved vis-a-vis different stakeholders over the last decade.

Researchers such as Druckman and Parkin [13] have analyzed news sentiments, showing how media's linguistic choices, such as the tone in political coverage, can influence reader attitudes. Similarly, Whitley et al. [49] found a shift toward a more positive tone in Canadian newspapers' mental health coverage, underscoring the subtle influence of media portrayal on public perception. Motivated by these insights, our study aims to unravel the tone conveyed in media reports on privacy events, exploring its influence on public understanding and attitudes.

4 METHODS

We detail the methods applied in our research. We first present the strategies employed for data collection and cleaning (Sec. 4.1). Then we describe our text-classification process and human validation (Sec. 4.2) and explain the temporal analysis we performed to track privacy coverage trends over time (Sec. 4.3). Finally, we outline our approach to topic modeling for the identification of prevailing themes in privacy reporting (Sec. 4.4) and discuss the sentiment analysis used to decipher the tone of the reporting (Sec. 4.5).

4.1 Data Collection & Cleaning

Our primary analytical lens prioritizes the Global North-South divide—in terms of economic development, digital access, and cultural factors—to ensure that our study reflects the complex, real-world landscape of global privacy issues. This dichotomy is essential to understanding the diversity in privacy issues. To systematically categorize the countries within this framework, we employ the United Nations' M49 standard, which delineates six global regions.[4] Overall, our study surveyed articles from 36 newspapers within 25 countries, ensuring balance by selecting 18 newspapers each from the Global North and the Global South, across six world regions.

To construct a representative dataset, we commenced with a preliminary selection of widely circulated and popular newspapers from each region. We refined our choices by considering several factors: The newspapers' rankings, the availability of their archives, their publication frequency—prioritizing those with daily issues—and their reporting style, specifically excluding tabloids. This refinement process involved iterative adjustments based on regional and international rankings, particularly utilizing the International Media and Newspapers (4IMN) ranking[5] to identify leading publications. Additionally, we sourced articles exclusively through the Lexis/Nexis archival service,[6] which afforded us a uniform data collection method across all regions. This approach ensured consistency and reliability in the data gathered, allowing for a more standardized comparative analysis.

Our analysis primarily spans the decade from 2013 to 2022, with articles from 2010 to 2012 included as available to broaden the historical context of our study. With the exception of two financial newspapers (*El Economista* and *Business & Financial Times*), the rest have coverage for at least one decade (2013-2022). Where available, we favored regional language newspapers to capture an authentic representation of the local privacy discourse. When faced with archival constraints, we turned to leading English-language newspapers such as the *Times of India* and *China Daily*. These publications have a wide local readership and can effectively cover diverse regional viewpoints. Our dataset comprises six languages: English, Spanish, French, German, Arabic, and Portuguese.

Based on the 4IMN ranking, the newspapers we selected fall in the top 10 or top 100 of their country or region, respectively, except for *Times of India* (141^{th} in regional rankings), and *The Moscow Times* (11^{th} in country rankings). Our selection criteria were designed to encompass both nationally influential and regionally significant newspapers. We also included financial-centric newspapers to cover economic impact and markets' reactions to breaches and regulations. To allow for better coverage of such topics and regions, we chose to include newspapers for which no ranking was assigned in the ranking list: *Financial Post*, *Nikkei Asia*, *Manawatu Standard*, *Caribbean News Agency*, and *The Dominion Post*.

For each newspaper source, we queried the term "privacy" or its local language equivalent terms against the L/N database. Due to Lexis/Nexis's download limits, collecting all news articles was infeasible. We thus focused on retrieving articles that specifically mentioned the keyword "privacy", ensuring our dataset was both manageable and relevant to our research objectives. For each article, we collected its title, content, an extract highlighting query matches, date of publishing, and word count. For newspapers published in languages other than English, we employ the Google Cloud Translation AI API [7] to translate them into English. These translated versions are stored alongside the original content in our database. Overall, we collected a total of 112, 572 articles.

Validation of Machine Translation Quality. To validate the quality of the automated translations, we designed a user study that required participants to post-edit machine-translated texts. Post-editing involves human processing of the text after machine translation [47]. Participants were provided with the original text and its machine translation, presented as sentence-by-sentence pairs. They were tasked with making minimal yet precise adjustments to ensure that machine translations closely mirrored the original texts in meaning, tone, and sentiment. We conducted the study through Prolific [28], a platform renowned for its engaged and attentive respondents [12]. We recruited 50 bilingual participants, evenly distributed across language pairs, each bringing an average of 26 years of linguistic experience. Each task involved post-editing a single article, followed by a questionnaire designed to assess the quality of the machine translations across dimensions

[4]"Standard Country or Area Codes for Statistical Use" – The M49 coding classification divides the world into six regional groups: Africa, Americas - Northern, Americas - Latin & the Caribbean, Asia, Europe, and Oceania (https://unstats.un.org/unsd/methodology/m49/)

[5]4 International Media & Newspapers is an international directory for newspapers, accessible at https://www.4imn.com/about/

[6]https://www.lexisnexis.com/en-us/professional/data-as-a-service/daas.page

[7]https://cloud.google.com/translate

such as accuracy, tone, and sentiment. Each task took approximately 33 minutes, with participants receiving $15 per task as compensation. Prolific also charged a $5 service fee per task. We employed stringent attention checks using deliberately misaligned translation pairs that required significant corrections. The ten participants who failed these checks were excluded from the analysis but were compensated, maintaining integrity and ensuring 50 valid responses through an additional $200 budget for replacement participants. Further details on the approach are provided in Appendix B.

To quantify the quality of our translations, we computed BLEURT (a BERT-based evaluation metric) [38], BLEU (Bilingual Evaluation Understudy) [31] and TER (Translation Edit Rate) [41] scores, providing objective measures of the translations' fidelity and fluency. The results, as depicted in Table 1, indicate high fidelity in translations across the languages we processed [3], with BLEURT and BLEU scores consistently reflecting a high degree of accuracy, and TER scores demonstrating minimal edits were required. Additionally, we confirmed that post-editing preserved the original machine translation's dominant tone and sentiment.

Removal of Duplicates. We consider articles published by the same newspaper to be duplicates if they have highly similar titles and were published within the same calendar week. Duplicates may appear because of editorial reasons, e. g., typographical corrections or narrative development. Besides editorial adjustments, articles may be republished at later times, with or without significant changes in the content, to bring fresh attention to past issues. We expect that our one-week timeframe is long enough to capture most of the duplicates caused by editorial updates and short enough as to not remove many intentional reprints. Appendix E provides further details about the similarity threshold used for duplicate removal. We only consider for further analysis the record with the latest date or, if the dates coincide, the one with the higher word count assuming the article was updated following a narrative development. Removing duplicates narrowed our set to 96, 275 articles (85.5% of the original collection).

Our final newspaper selection is shown in Table 2.

Language	BLEURT	BLEU	TER	Tone	Sentiment
Arabic	86.0	79.9	14.2	9	7
French	93.6	96.0	3.7	10	9
German	90.7	88.0	9.3	10	10
Portuguese	92.6	90.1	7.7	10	10
Spanish	92.0	90.6	5.7	10	10

Table 1: Translation Metrics. BLEURT & BLEU Scores: higher are better (max. is 100). TER Score: lower is better (min. is 0). Tone and Sentiment columns show the number of articles (out of 10) where post-editing maintained the original machine translation's dominant tone and sentiment, respectively.

4.2 Privacy Text Classification

When extracting articles from the L/N database, we anticipated collecting articles that only collaterally mention our query term and do not, in fact, focus on digital privacy as defined in our study. To ensure the validity of our data set, we proceeded to remove such false positives. We needed a binary classifier to distinguish between *privacy* and *non-privacy* articles. Text classification is a fundamental problem in Natural Language Processing (NLP). In recent years, pre-trained language models have proven exceptionally effective at learning universal language representations by leveraging extensive corpora of unlabeled text. For our privacy filter, we utilized two prominent models: OpenAI's GPT [5] and Google's BERT [10].

Ground Truth. Our study leverages the comprehensive privacy frameworks established by Solove [42] and Antón & Earp [2] to construct a nuanced operational definition of digital privacy, detailed in Appendix D (cf. Listing 14). Solove's framework categorizes privacy issues into four groups—information collection, processing, dissemination, and invasion—each with specific privacy concerns. Antón & Earp's taxonomy, developed by applying grounded theory to online privacy policies, reveals twelve categories of privacy elements spread across two broad classifications: privacy protection goals and potential vulnerabilities. Combining these insights, the study's definition addresses the handling of personal information, the importance of protective measures, and ethical considerations. Central to this definition is informed consent, highlighting the individual's right to control their personal data in the digital space. These frameworks categorize privacy issues and delineate protection goals and vulnerabilities, respectively, guiding our methodology for classifying newspaper content by privacy relevance.

We operationalized these definitions into explicit inclusion and exclusion criteria for our classification task. For instance, discussions on surveillance (reflecting Solove's "Information Collection") and articles examining online services' data management for personalized experiences (aligned with Antón & Earp's "Information Personalization") were flagged as privacy-centric. We paid particular attention to "Secondary Use" and "Information Transfer" practices, emphasizing transparency and individual consent, critical elements derived from our foundational frameworks.

In refining our exclusion criteria, we focused on articles that, despite mentioning personal data, lacked depth in privacy analysis—such as cursory technological reports devoid of privacy implications. Additionally, we filtered out articles that, though employing privacy-related terms, diverged from our study's emphasis on digital privacy. For example, narratives centered on individuals seeking seclusion from public exposure—such as defendants desiring privacy in legal contexts—and discussions praising the privacy advantages of specific real estate, were deemed peripheral. To maintain a sharp focus, such articles were excluded, aligning our analysis closely with the digital privacy issues our theoretical frameworks aim to highlight.

Our meticulous annotation process involved two expert privacy researchers, who individually annotated 600 randomly selected articles for privacy focus (*privacy*, *non-privacy*). The high Cohen's Kappa score of approx. 0.936 not only attests to the reliability of our annotations but also underscores the effectiveness of our operational definitions in facilitating a shared understanding of privacy-focused content. Our manual annotation yielded 44.31% *privacy* and 55.69% *non-privacy* articles. We divided the jointly agreed upon annotated subset (571 articles) into a training (456 articles) and test (115 articles) set. Our training set was split into .9 training (410 articles) and .1 validation (46 articles) sets.

Table 2: Newspapers Included in the Study: Newspaper ID, Country of Origin/Publishing, Language of Publishing, Focus (G: General; F: Financial), Ranking (Newspaper Rank for both Region & Country, if not available N.A. is used), Coverage Duration (Start Year – End Year), and Article Count (# of Articles Focused on Digital Privacy).

Region Newspaper	**ID**	**Country**	**Language**	**Focus**	**Ranking** Region / Country	**Start Year**	**End Year**	**Article Count**
				Global North				
Americas - Northern								
The Toronto Star	TS	Canada	English	G	13 / 2	2010	2022	2837
La Presse Canadienne	LPC	Canada	French	G	58 / 5	2010	2022	256
Financial Post	FPC	Canada	English	F	N.A.	2010	2022	1055
The New York Times	NYT	United States	English	G	1 / 1	2010	2022	2100
The Hill	THU	United States	English	G	8 / N.A.	2010	2022	360
USA Today	USA	United States	English	G	3 / N.A.	2010	2022	1343
Europe								
The Daily Telegraph	DT	England	English	G	3 / 3	2010	2022	2494
Financial Times	FTL	England	English	F	7 / 5	2010	2022	2894
Le Figaro	LFF	France	French	G	16 / 2	2010	2022	252
Sueddeutsche Zeitung	SZG	Germany	German	G	24 / 4	2010	2022	1421
The Moscow Times	TMT	Russia	English	G	82 / 11	2010	2022	159
El Pais	EPS	Spain	Spanish	G	5 / 1	2010	2022	1187
Oceania								
Australian Financial Review	AFR	Australia	English	F	7 / 6	2010	2022	1267
Herald Sun (Melbourne)	HSM	Australia	English	G	5 / 4	2010	2022	1160
Sydney Morning Herald	SMH	Australia	English	G	1 / 1	2010	2022	1683
Manawatu Standard	MSN	New Zealand	English	G	N.A.	2010	2022	550
The New Zealand Herald	NZH	New Zealand	English	G	4 / 1	2010	2022	2901
The Dominion Post	TDP	New Zealand	English	G	N.A.	2010	2022	1227
				Global South				
Africa								
Daily News Egypt	DNE	Egypt	English	G	35 / 5	2010	2022	161
Business and Financial Times	BFT	Ghana	English	F	N.A. / 3	2016	2022	89
Daily Nation	DNK	Kenya	English	G	1 / 1	2013	2022	348
The Sun	TSN	Nigeria	English	G	27 / 6	2013	2022	156
This Day (Lagos)	TDL	Nigeria	English	G	17 / 5	2010	2022	115
The Daily Monitor	TDM	Uganda	English	G	16 / 1	2013	2022	219
Americas - Latin & The Caribbean								
La Nacion	LNA	Argentina	Spanish	G	4 / 3	2010	2022	540
O Estado de S. Paulo	ESP	Brazil	Portuguese	G	7 / 1	2010	2022	912
Caribbean News Agency	CAN	Caribbean	English	G	N.A.	2012	2022	96
El Economista	EEM	Mexico	Spanish	F	29 / 4	2018	2022	493
El Universal	EUM	Mexico	Spanish	G	6 / 1	2010	2022	895
El Comercio	ECP	Peru	Spanish	G	8 / 1	2010	2022	191
Asia								
China Daily	CD	China	English	G	3 / 1	2013	2022	749
The Times of India	TOI	India	English	G	N.A. / 140	2010	2022	4185
Nikkei Asia	NA	Japan	English	F	N.A.	2010	2022	65
Dawn	DN	Pakistan	English	G	19 / 1	2013	2022	355
Asharq Alawsat	AAA	Saudi/Pan-Arab	Arabic	G	89 / 2	2012	2022	339
Khaleej Times	KT	UAE	English	G	50 / 3	2010	2022	601
Total								35,655

BERT Baseline. BERT (Bidirectional Encoder Representations from Transformers) has achieved notable results in many language comprehension tasks [10]. Trained on plain text for masked word prediction and next-sentence prediction tasks, BERT can be fine-tuned to enhance its performance on text classification tasks. Since BERT is trained in the general domain with a data distribution different from our target domain of privacy filter, we further pre-trained BERT with our human-annotated article set.

We fine-tuned the BERT model for sequence classification on the jointly agreed upon annotated training set. Our training iterated over 10 epochs in batches of 16 articles (Figure 24). One limitation of BERT is its encoding sequence maximum size of 512 tokens. Existing works have employed truncation (e. g., first 512 tokens) and hierarchical strategies (iteratively obtaining BERT representations for each fraction of a long article, then combining the outputs). Sun et al. [43] compared the performance of different fine-tuning approaches for long texts from IMDb and Sogou News. The authors found that truncating the head and tail of the documents returned the best performance. Since the median number of words per article in our dataset is 696, we trained and tested the model over the first 512 tokens (words) of each article only, on the assumption that this will be enough to reveal the intended focus of an article.

Out of the ten models, we picked the best-performing one in terms of accuracy over the validation test (0.809) and average training loss (0.038). Upon running the trained model on the test set, we obtained a Matthew correlation coefficient (MCC) of 0.836. The approach yielded a 91.3% accuracy, with detailed performance metrics provided in Table 3.

GPT Classifier. In our privacy filtering process, we harnessed the capabilities of the gpt-3.5-turbo-0301 model accessible via OpenAI API. We selected the GPT-3.5 Turbo model due to its scalability and cost-efficiency, aligning with our budget and API rate limits for processing a vast dataset of 96,275 articles, and its proven track record in similar text classification and annotation tasks [14, 15, 19]. Moreover, as we will detail in the section later on, the performance of GPT-3.5-turbo already exceeded the BERT baseline.

We refined our prompt query through multiple iterations and finalized it as detailed in Listing 1 in Appendix A, which asks, *"Has the article discussed aspects of digital privacy? Answer 1 if True, 0 if False or unknown."* To assist the model in accurately interpreting this task, we provided a comprehensive definition of digital privacy, referenced in Listing 14 (Appendix A), drawing from established privacy frameworks by Solove [42] and Antón & Earp [2].

We present the evaluation of the approach on the test set in a zero-shot setting in Table 3. The numbers demonstrate the superior performance of the GPT-based approach over BERT for our classification task. The approach achieves precision, recall, and F-1 score of 0.902, 0.958, and 0.929 respectively. We recognize that the 93.9% accuracy rate of our GPT-based filtering, while high, is not perfect and may introduce some systematic errors in identifying privacy-related articles. Nevertheless, alternatives such as employing crowdworkers for such nuanced tasks bring challenges in ensuring consistent interpretations of 'privacy' and could demand substantial time and resources. Given these trade-offs, we opted for the automated approach, acknowledging its limitations while providing a practical balance for our study's scale.

Table 3: Performance comparison between the BERT baseline and the GPT filter on the test set.

Cohen's kappa coeff.		0.936	
Set size *(# articles)*	Training	410	
	Validation	46	
	Test	115	
		BERT baseline	GPT filter
Training set	Avg. train. loss	0.038	-
Validation set	Accuracy	0.809	-
Test set	Matthew corr. coeff.	0.836	-
	Accuracy	0.913	0.939
	Precision	0.906	0.902
	Recall	0.931	0.958
	F-1 Score	0.911	0.929

Privacy Filter. Our GPT-based filter was applied to the duplicate-free dataset, which resulted in the identification of 35,655 (37.03%) *privacy* and 60,620 (62.97%) *non-privacy* articles, the former of which are analyzed further (see Table 2 for their distribution by newspaper). Table 5 (Appendix H) provides a breakdown by year and newspaper of the number of articles published on privacy.

4.3 Temporal Analysis

The articles from a 13-year period were processed into time-series data and grouped by month, offering a balance between spotting short-term trends and maintaining a manageable data volume for analysis. We also evaluated and plotted a quarterly moving average, where necessary. This was to smooth out short-term fluctuations and highlight longer-term trends or cycles, providing a clearer view of the data's overall direction, especially when monthly data appeared too volatile.

We first used the Augmented Dickey-Fuller (ADF) [6] test to validate the stationarity of our time-series data, a prerequisite for reliable trend analysis. With confirmed stationarity, we used the Mann-Kendall test [16, 23] to detect any monotonic trends in the privacy-related articles' frequency. We used Sen's Slope Estimator [39] for the rate of change, giving us a specific slope value to better comprehend the evolution of privacy coverage over time.

4.4 Topic Modeling

We then delve into our topic modeling process, which reveals dominant themes and their shifts in privacy reporting over the past decade. This exercise offers insights into the substantive content of privacy coverage, unveiling which facets of privacy have been spotlighted in media discourse.

GPT-inferred Focus Topics. Our initial step was to use the GPT-3.5-turbo large language model as an automated tool for generating 3 to 5 keywords that encapsulate the focus of each article.

To validate the effectiveness of these LLM-generated topics, we conducted a user study with 50 participants (the same evaluators of translation quality in Table 1) who rated the relevance and comprehensiveness of these topics for a series of articles on a 5-point Likert scale. Where participants found gaps, they were encouraged to suggest additional terms that would better encapsulate the article's content, thereby offering insights into any missing perspectives.

The relevancy ratings, illustrated in Figure 1 (left), depict a clear tendency towards high scores (4 or 5) for 224 topics under evaluation, indicating that participants generally found the LLM-generated keywords to be aligned with the content of the articles. In terms of comprehensiveness, as shown in Figure 1 (right), the majority of our participants rated the collection of keywords as covering the key points of the articles effectively.

For our focus topic frequency analysis, we culled the top 1000 recurrent keywords. To enhance relevance, we carried out further preprocessing using an automated script that excluded common terms such as tech company and country names, focusing the dataset on privacy issues. To encapsulate broader themes, we consolidated related terms into more expansive categories using a combination of manual review and automated scripts. For instance, terms like "privacy invasion," "invasion of privacy," and "privacy violation" were consolidated under the broad banner of "privacy invasion." We implemented a similar strategy for other vital themes such as data breaches, legislation, surveillance, and social media.

LDA Topic Modeling. To gain insights into the topics covered in the privacy dataset, we performed an exploratory analysis using Latent Dirichlet Allocation (LDA) topic modeling. LDA is a widely used approach for discovering hidden thematic structures within text data without the need for labeled training data. By assigning topics to articles and words to topics, LDA can distill a large set of articles down to a few representative topics. We chose to apply LDA to article focus keywords, summaries, and titles, as sourced from L/N database, for its better precision (perplexity) and coherence in topic generation. Perplexity assesses how well the model predicts samples, with lower scores being better. The topic coherence (c_v) measure [34] evaluates topic quality by assessing the semantic similarity of high-scoring words, with higher scores indicating more meaningful topics. Upon manual inspection, the topics generated were more focused, relevant, and insightful than those from full-text LDA. The number of topics was chosen to be 30 based on coherence score metrics, supplemented by manual examination for meaningful interpretability. For ease of interpretation, we inspected the individual topics and aggregated them into twelve larger categories. The trained model was then applied to all newspaper articles to retrieve topic probabilities.

4.5 Tone Analysis

To examine the reporting style of privacy-related topics, we analyze the content of each article using the IBM Watson Natural Language Understanding (NLU) Standard Plan, version 4.7.1. This service conducts linguistic analysis of written text and provides a scorecard for each detected sentiment, and emotional tone(s). While sentiment analysis categorizes attitudes as *positive*, *negative*, or *neutral*, emotion analysis delves deeper to identify specific underlying emotions contributing to sentiments. For this study, we focused on detecting emotional tones like *anger*, *fear*, *joy*, *disgust*, and *sadness*. The emotional tone predictive algorithm considers features such as n-grams, punctuation, and sentiment polarity. We conducted a document-level analysis to capture a holistic view of the sentiment

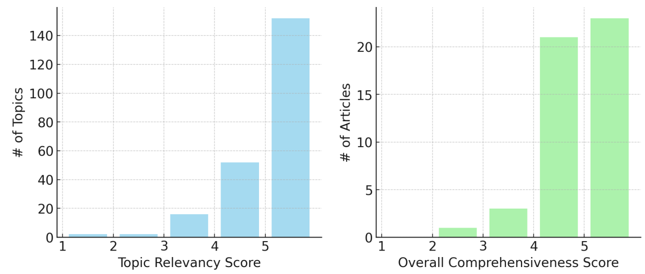

Figure 1: Topic relevance (left) and comprehensiveness (right) validation. Based on a 5-point Likert-scale rating. Higher is better.

conveyed in privacy-related articles. Sentiment scores, ranging from -1 to 1, indicated negative, positive, and neutral sentiments for scores less than 0, greater than 0, and equal to 0, respectively. For emotion analysis, each of the five tones—anger, fear, joy, disgust, and sadness—received a score between 0 and 1, with higher scores representing a stronger emotional indication. While IBM's NLU API supports sentiment analysis across all languages in our dataset, tone analysis is confined to English and French content. For tone analysis, our study focused on English-translated versions of the articles, acquired via the Google Translations API.

5 RESULTS

In this section, we report the results of the temporal analysis, topic modeling, and tone analysis on the refined dataset of 35,655 data privacy articles. Table 2 presents a comprehensive breakdown of the total count of articles.

5.1 Temporal Analysis

5.1.1 Overall Coverage Trend. Figure 2 depicts a consistent and steady increase in media coverage of privacy-related issues over the past decade. We investigated the temporal trends in article publications during this period, employing the ADF, Mann-Kendall, and Sen's Slope Estimator statistical tests. The ADF test confirmed the dataset's stationarity without differencing, enabling direct trend interpretation. The Mann-Kendall test revealed a statistically significant positive correlation ($p < 0.001$) between time and article count, indicating a weak but evident increasing trend. Sen's Slope Estimator further supported this finding, estimating a slight upward trend in the article count.

5.1.2 Regional Coverage Trend. Figure 3 presents the coverage trend of privacy articles as a three-month moving average for different regions. As demonstrated in the figure and verified by the Mann-Kendall test, there is an upward trend in the coverage of privacy issues in the Global South ($p < 0.001$). In contrast, coverage in the Global North increases around major developments but is roughly consistent over time.

The global trends observed earlier in the overall media coverage of privacy are also visible within each of the six regions, with spikes occurring around significant global developments, irrespective of the absolute number of articles published. This finding suggests that privacy-related matters garner increased attention during critical

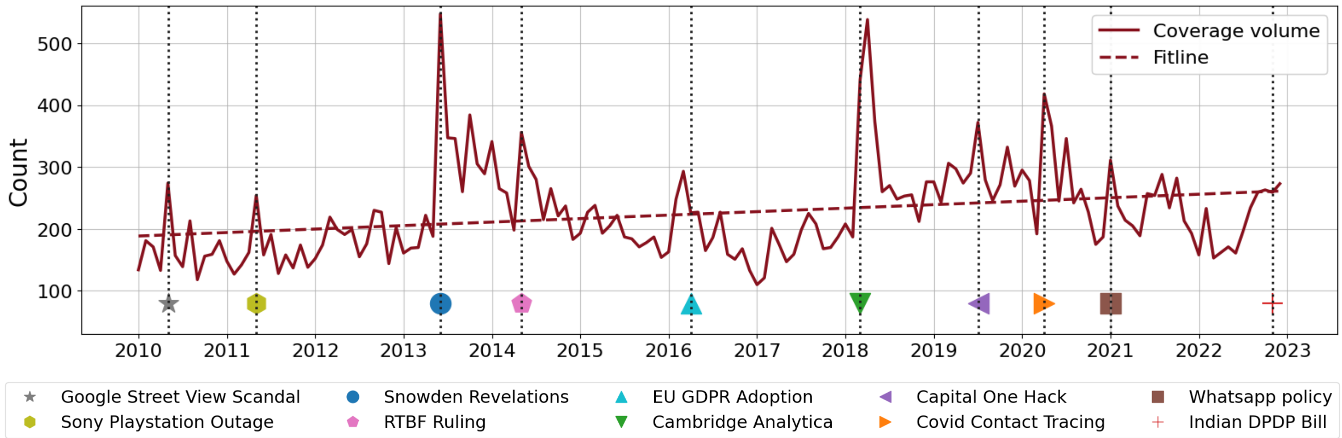

Figure 2: Time series of monthly newspaper coverage of digital privacy across 25 countries and 6 regions. for the duration of 13 years (2010 - 2022). To better understand the evolution, we limit this analysis to those 34 newspapers with collection start year in 2013 or before.

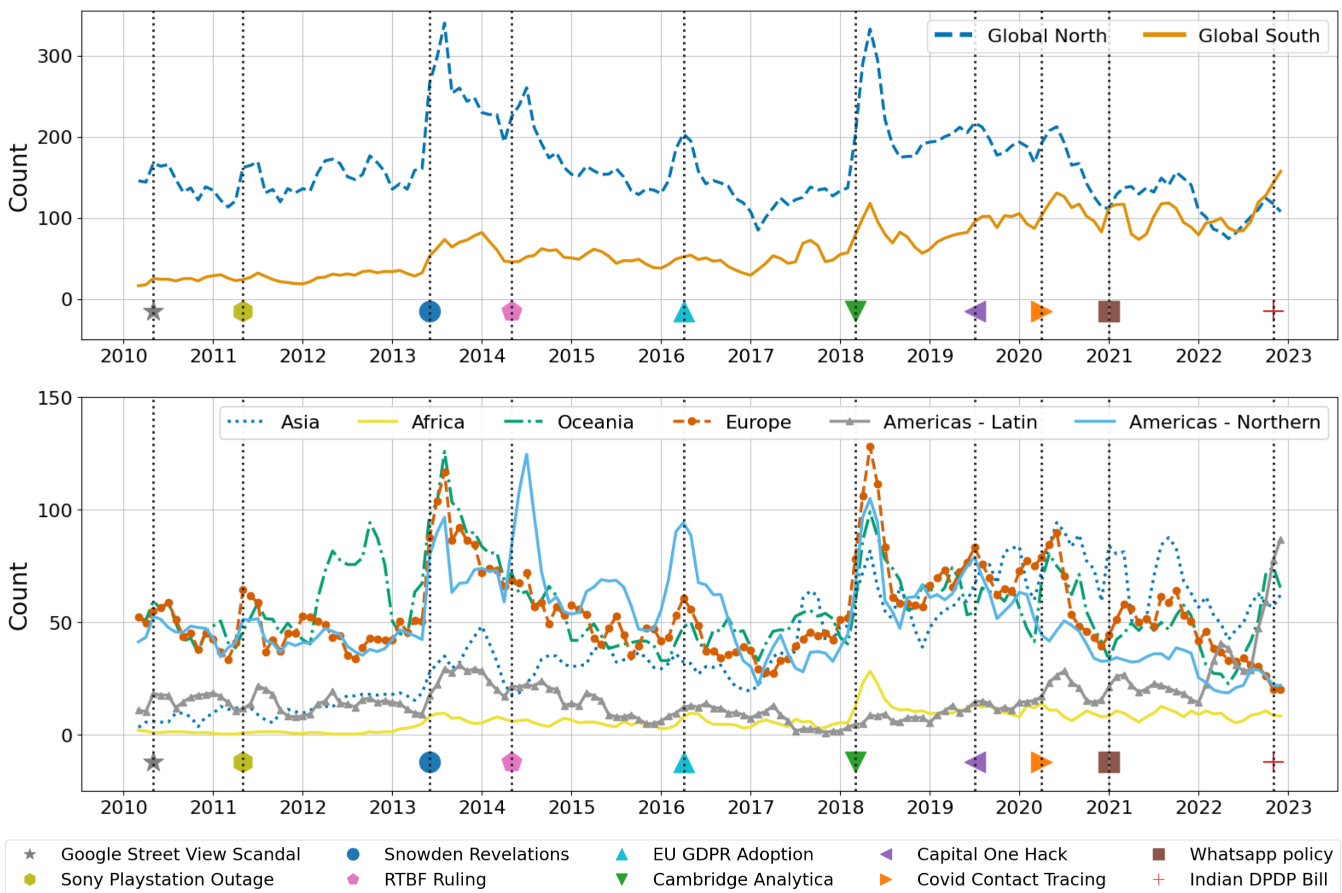

Figure 3: Time series of the quarterly moving average of privacy coverage across different regions. The top chart shows a comparison between the Global North and the Global South whereas the bottom chart shows trends for finer-grained regions.

global events, highlighting the interconnectedness of privacy concerns on a global scale. To group the regions based on the similarity of their time series data, we performed time series clustering using a k-means clustering algorithm, which resulted in three clusters (cf. Figure 16 in Appendix F).

5.1.3 High Profile Incidents. Several events have dramatically influenced the public discourse around digital privacy over the last decade. Attention in online newspapers has spiked during pivotal years marked by significant incidents. Earlier, in 2010, Google's Street View scandal had drawn attention to the vulnerabilities and potential misuse of geolocation data. In 2013, the global surveillance disclosures by former NSA contractor, Edward Snowden, revealed extensive surveillance programs, awakening a heightened global consciousness of privacy rights and governmental oversight. The following year, the "Right to be Forgotten" ruling by the Court of Justice of the European Union set a crucial precedent for personal data control and reshaped the discourse on data privacy rights. In 2018, the Facebook-Cambridge Analytica scandal underscored the pervasive risks of personal data misuse on social media platforms, prompting a clamor for stringent regulations and transparency. The advent of COVID-19 contact tracing apps in 2020 introduced novel privacy concerns, balancing public health initiatives against individual privacy rights, thus underscoring the complexities inherent in policy-making for an increasingly interconnected world.

5.2 Topic Modeling

5.2.1 GPT-inferred Focus Topics. Figure 4 visually represents the top 30 focus areas in the digital privacy landscape discourse. The range of issues is vast, covering areas such as the digital economy, health data management, the policing system, social networks, and online advertising. The potential for abuse of emerging technologies such as artificial intelligence and facial recognition received significant attention. The digital privacy issues faced by susceptible groups, including children, victims of online sexual harassment, and the elderly, have also been highlighted in the news media.

To ascertain which tech companies have been at the forefront of the digital privacy discourse over the last decade, we plotted the ten most frequently featured companies in Figure 13 in Appendix H. Due to the Cambridge Analytica scandal, Facebook emerged as the focus of approximately 10% of all articles in our dataset. Apple, due to its standoff with the NSA over iPhone unlocking, and Google, due to various legal battles over the right to be forgotten and Google's Street View scandal, also remained significant points of focus. Updates to WhatsApp's policies in late 2022 incited considerable uproar in the Global South, particularly in India. Conversely, in North America and Europe, concerns over TikTok's use of personal data have consistently been a point of contention.

5.2.2 LDA Topic Modeling. To gain insights into the topics generated by LDA (Section 4.4), we manually inspected and categorized them into twelve broad themes. The results are presented in Table 4, showcasing cohesive and recognizable topics.

Using the trained LDA model, we assigned each article in our dataset to the topic with the highest likelihood based on its content. Figure 5 provides a detailed analysis of the evolving themes within the privacy discourse from 2010 to 2022. Each distinct color in the

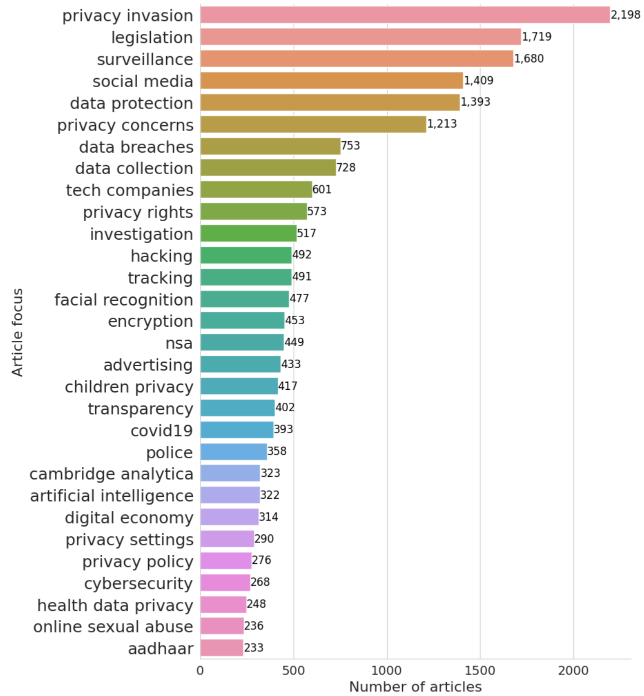

Figure 4: 30 most frequent focus topics of newspaper articles as annotated by GPT-3.5-turbo.

plot corresponds to a specific topic category, and the size of the colored areas represents their relative contribution to the cumulative privacy discourse over the specified timeline.

The analysis of digital privacy reporting over time has consistently shown a notable focus on the *tracking and tracing* of individuals. Tracking individuals online through their browsing patterns, location tracking through apps and the increased adoption of smart home devices remained popular subtopics within this category. Notably, the Sidewalk smart city project in Toronto has raised concerns about potential privacy invasions. During the COVID-19 pandemic, the topic gained significant attention due to privacy concerns related to contact-tracing apps.

Another noteworthy and consistently high-reporting topic in the digital privacy discourse pertains to various forms of *online abuse*, particularly concerning vulnerable populations. With the mainstreaming of digital platforms over the last decade, concerns have arisen regarding children's usage of online platforms without adequate parental supervision. Instances of unsolicited explicit content, revenge porn, cyberstalking, and harassment have also been widely reported, highlighting the pressing need to address these issues and safeguard vulnerable individuals in the digital realm. The sustained attention to these topics underscores their relevance and calls for sustained efforts to combat digital abuse.

There was a noteworthy surge in reporting on *government surveillance*, particularly following the Snowden Leaks in 2013. Edward Snowden's revelations about extensive surveillance activities conducted by government agencies, such as the NSA, served as a catalyst for heightened public awareness. It sparked intense discussions

Table 4: Broad topic categories derived from the LDA model alongside the top words for each topic.

Topic	Top Words
Online Abuse	child, student, woman, sexual, abus, victim, parent, video, pay, photo, publish, million, violat, lawsuit, protect, block, consent, law, websit, breach, safeti, regul, lose, famili
Social Media	social, medium, network, site, share, account, platform, profil, post, friend, peopl, content, like, concern, delet, privat, photo, experi, protect, allow, control, access
Corporate Responsibility & FinTech	card, ident, credit, employe, web, work, employ, servic, manag, system, secur, free, govern, financi, plan, bank, compani, corpor, access, safe, number, public, issu, monitor
Surveillance Technologies	camera, recognit, facial, polic, surveil, instal, softwar, watch, control, home, system, crime, citi, imag, devic, identifi, public, video, state, civil, offic, record, hide, spi
Privacy Incident Investigations	breach, investig, email, commiss, polic, probe, complaint, journalist, watchdog, offic, report, govern, document, illeg, reveal, minist, bank, hack, alleg, agenc, law, releas
Government Surveillance	surveil, spi, snowden, nsa, agenc, govern, nation, intellig, program, state, presid, terror, collect, phone, snoop, foreign, secret, call, citizen, record, servic, law, monitor
Court Rulings & User Rights	court, right, rule, case, justic, search, order, judg, law, union, protect, violat, govern, legal, human, request, decis, remov, state, public, lawyer, feder, act, europ, battl
Consumer Tracking & Tracing	app, mobil, encrypt, messag, smartphon, trace, android, contact, user, use, iphon, ban, hack, applic, track, concern, access, call, devic, allow, store, launch, health, collect
Regulation & Governance	regul, govern, protect, discuss, articl, need, transpar, right, concern, highlight, risk, global, challeng, technolog, surveil, individu, law, intellig, public, trust, futur, market
Legislation & Policy	law, bill, freedom, legisl, protect, govern, tax, right, propos, enforc, commun, pass, press, act, civil, power, express, access, feder, minist, critic, regul, agenc, surveil, reform
Data Breach Scandals	hack, charg, hacker, secur, attack, breach, stole, million, data, steal, crime, victim, target, charg, state, compani, nation, militari, cybersecur, protect, defenc, war, foreign
Big Tech & Public Perception	tech, big, compani, new, polici, regul, announc, updat, improv, search, engin, web, servic, user, featur, control, concern, busi, protect, servic, custom, deal, trust, competit

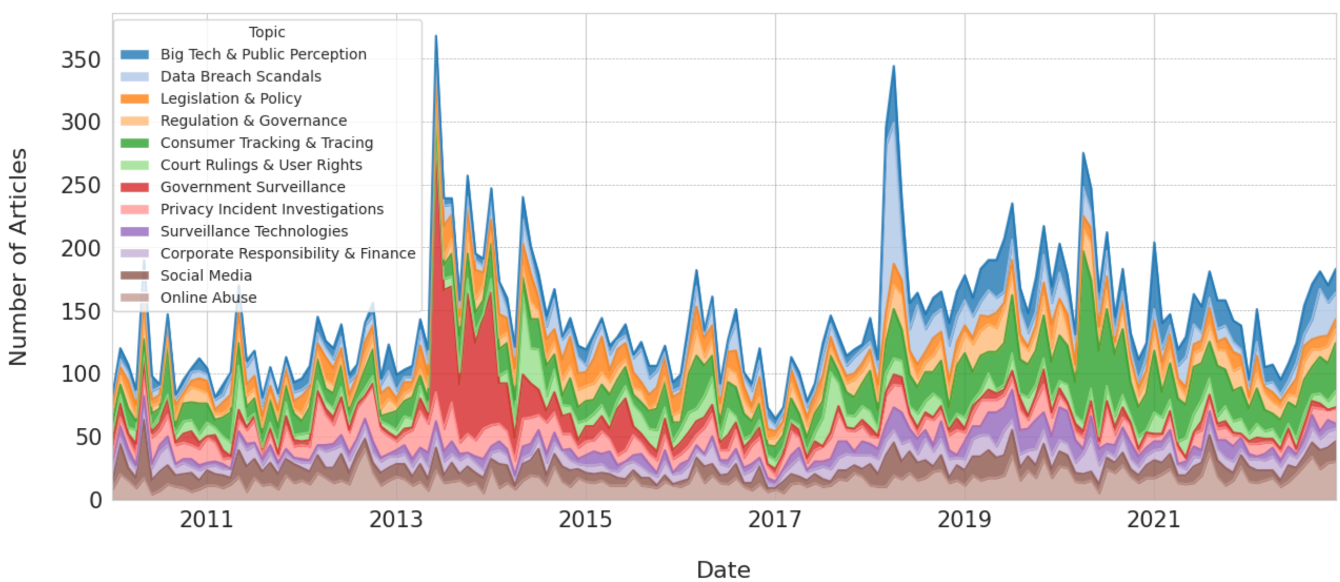

Figure 5: Prevalence of Topics in Privacy Discourse from 2010 to 2022. The colored regions in the stacked area plot correspond to twelve different topic categories, each showing their contribution to the overall discourse over time.

in the digital privacy discourse, highlighting the need for greater scrutiny of government surveillance practices and advocating for transparency and accountability in the digital age.

Over time, the focus in the digital privacy discourse shifted from primarily centering on government surveillance to encompassing

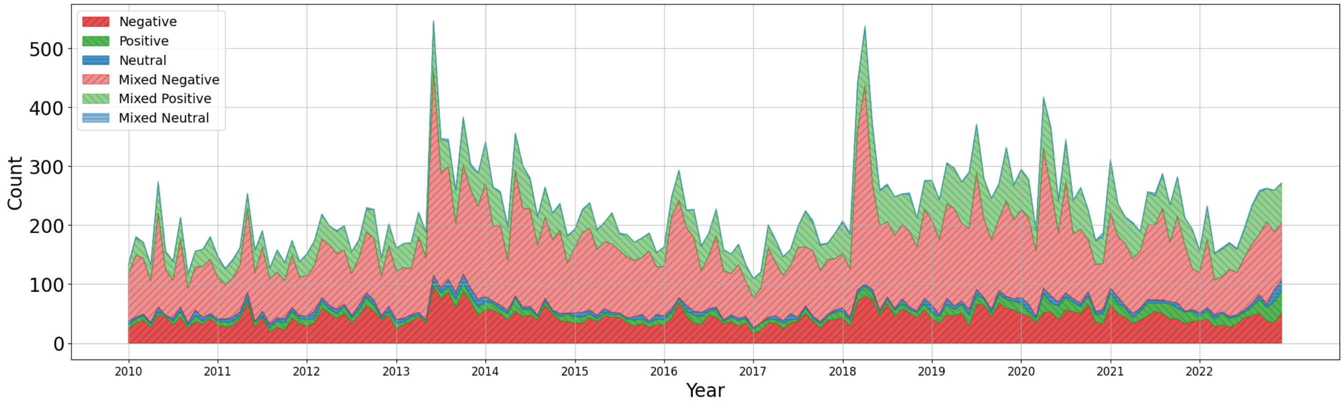

Figure 6: Evolution of sentiment present in privacy-related coverage in newspapers. Each color's expansion and contraction over time provide a visual representation of the sentiment's prominence within the overarching privacy discourse during the given period.

the practices of *big tech* companies. A series of *data breaches and privacy scandals* brought these companies' data practices into question, raising concerns about the appropriate use and protection of personal information. As a result, the privacy conversation expanded to include *corporate responsibility* and the necessity of robust privacy regulations to safeguard individuals' sensitive information.

Figure 18 (Appendix I) and Figure 19 illustrates a comparative examination of the temporal trends in topic popularity between the Global North and the Global South. Intriguingly, the analysis reveals noteworthy privacy developments in both regions. For instance, the increased presence of reporting on *Court Rulings and User Rights* in the Global South during 2017 can be predominantly attributed to the Indian Supreme Court's decision to declare privacy a fundamental right. This landmark ruling significantly impacted privacy discourse in the region and received substantial media attention.

Furthermore, the attention given to *Big Tech and Public Perception* is significantly increasing in the Global South, signifying a growing interest in discussions about the considerable power and influence of major technology companies. On the other hand, *Data Breaches* and *Privacy Incident Investigations* attract considerably more attention in the Global North than in the Global South. Oceania, Europe, and North America tend to report the most on these topics, highlighting the heightened concerns and media scrutiny surrounding data breaches and privacy violations in these regions.

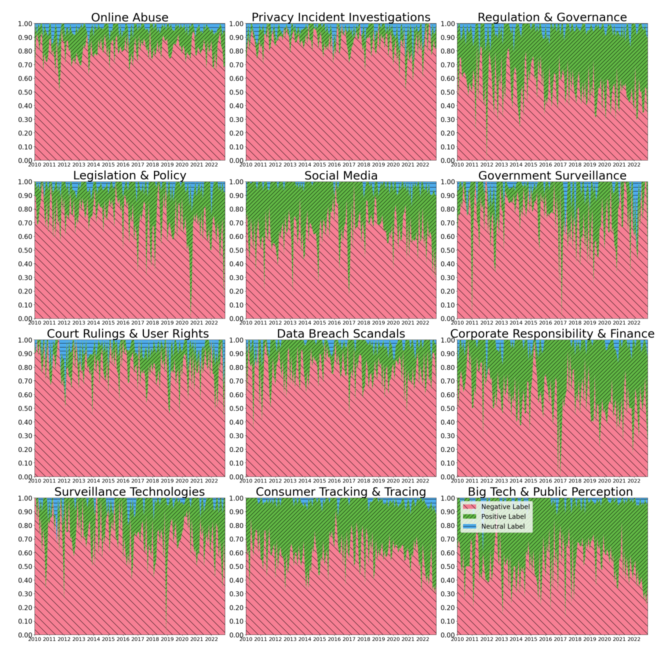

Figure 7: A time-series view of sentiment prevalence across 12 distinct LDA-derived privacy-related topics.

5.3 Tone Analysis

5.3.1 Sentiment Analysis. We conducted sentiment analysis to assess the reporting trends of privacy-related articles over the years in terms of positive, negative, and neutral sentiments. During this process, we also identified articles conveying mixed sentiments, although most of them exhibited a dominant tone of either positive or negative sentiment. Figure 6 presents the trend of six sentiment categories for the entire corpus of articles. We found that the majority of the articles exhibited a negative sentiment, underscoring the prevailing apprehension and concern surrounding privacy matters. That said, the majority of articles categorized as predominantly negative also exhibited mixed sentiments, containing elements of positivity alongside negativity.

Figure 7 presents a comprehensive area chart depicting the sentiment split over time for each of the twelve identified privacy topics. Through visual exploration, we can discern the changing emotional landscape surrounding various privacy concerns. For instance, topics like *Government Surveillance* and *Corporate Responsibility* exhibit fluctuations in sentiment as public perceptions respond to major developments or incidents. On the other hand, *Online Abuse* and *Data Breach Scandals* reveal consistent sentiments over time, reflecting enduring public sentiments and concerns in those areas.

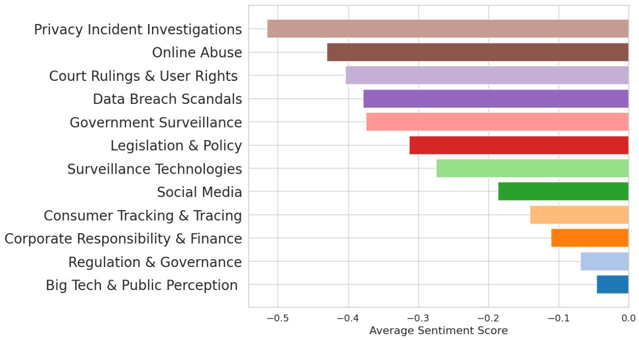

Figure 8: Average sentiment score for 12 LDA-derived privacy-related topics.

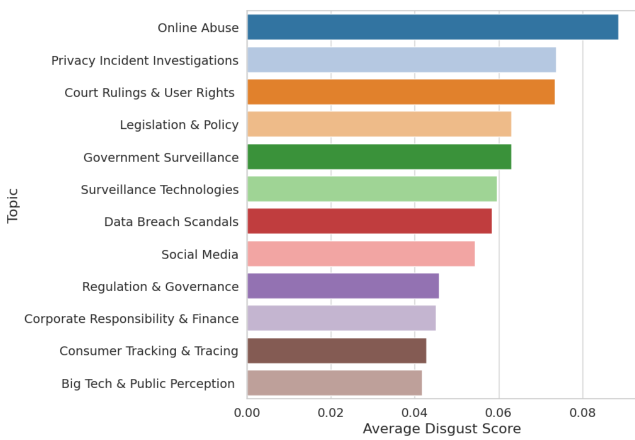

Figure 9: Average score for Disgust emotion across 12 LDA-derived topics.

Our analysis of the average sentiment scores across 12 topics reveals a distinct pattern, as depicted in Figure 8. The topic of *Privacy Incident Investigations* records the highest negative sentiment, suggesting that such investigations often reveal the extent of non-compliance with privacy regulations, thereby intensifying public distrust and negative sentiment. The less negative sentiment towards *Regulation & Governance* may reflect public recognition of the importance of regulations and control measures in safeguarding privacy. The relatively balanced sentiment towards *Big Tech & Public Perception* is shaped by media narratives that highlight both the efforts and shortcomings of tech giants in privacy matters. Notable instances like WhatsApp's policy reversal amid public backlash underscore the potential of big tech companies to adapt to be on the favorable side of public perception. Figure 20 depicts the overall split of sentiments across the regions and languages present in our dataset. Notably, there were distinct differences in sentiment between Latin America and the other regions. In Latin America, the sentiment tended to be more positive in contrast to predominantly negative sentiments observed elsewhere. This disparity in sentiment may be attributed to varying cultural perspectives and public attitudes towards privacy in different regions.

5.3.2 Emotion Analysis. We extended our analysis beyond sentiment to explore the emotional tones embedded within the coverage. We investigated whether different privacy-related developments are reported in distinguishable emotional tones such as joy, sadness, fear, anger, or disgust. Figure 21 provides temporal regional snapshots of the average emotion scores for each emotion, demonstrating a consistent pattern of emotional tones over time. Intriguingly, Joy and Sadness are represented in roughly equal proportions across the time series, suggesting a balanced interplay of these emotions in the privacy discourse. Conversely, Anger, Disgust, and Fear register significantly lower scores, suggesting that the high-quality newspapers' commitment to measured, balanced reporting may limit the amplification of these more intense negative emotions in privacy-related coverage.

The emotional tones associated with different privacy-related topics provide valuable insights into how the public emotionally responds to specific privacy concerns and policy discussions. When the discourse revolves around topics like *Government Surveillance* or *Surveillance Technologies*, emotions of fear and anger emerge prominently (Figure 22, Appendix H). Articles discussing *Online Abuse* evoke a strong sense of disgust, reflecting the public's emotional response to the disturbing nature of online harassment, cyberstalking, and other forms of abusive behavior on digital platforms (Figure 9). On the other hand, *Regulation and Governance* topics elicit the most joy in the tone. Interestingly, sadness is most observed in articles discussing *Consumer Tracking and Tracing*.

6 DISCUSSION

Next, we examine the strengths and limitations of our methodological approach. We then delve into the implications of our findings for various stakeholders and propose avenues for future research.

6.1 Methodological Insights & Limitations

Our methodology has surfaced several insights that emphasize the efficacy of our approach and highlight areas of future improvement.

Dataset Curation and Model Performance. Our research was constrained by the availability of datasets encompassing major newspapers from key countries for the full duration of the study. While our results are reported with the granularity of six regions, our primary interest lies in examining the divide between the Global North and South. In this regard, our dataset was adequately representative and sufficient. In curating our dataset, we found that the GPT family of LLMs outperforms conventional supervised methods, such as fine-tuned BERT models for zero-shot privacy text classification. When prompted with carefully constructed domain context, GPT models are comparable to human annotators, an insight in line with recent work for hate speech and genre classification [15, 19]. Our application of GPT-3.5-turbo in text classification showcases the utility of these models in varied research contexts, echoing recent studies on news summarization [50] and text annotation [14]. Yet, their effectiveness varies by task, highlighting the need for precise validation for each application, a practice supported by our results and recent studies in the field [29]. For our annotation task, testing other open-source models such as Falcon [1] or Llama [46], while valuable, was deemed beyond our study's focused scope, which is not centered on model evaluation.

Multi-Language Analyses and Automated Translation. To conduct a cross-cultural analysis, capturing an international snapshot is challenging and language in particular can be a huge technical barrier. To address this, we employed Google Translate for processing non-English content. While necessary for a study of this scale, this approach may not fully capture the nuances, especially the sentiment and emotional tone, as effectively as native language analysis. However, our post-editing validation study with bilingual speakers (n=50) confirmed the translations' accuracy in preserving meaning and tone to be sufficient in this context. Our deliberate choice of well-resourced language pairs (such as French, German, Spanish, Arabic to English), where sufficiently large training data is available, contributes to the expected translation reliability. That said, disparities were observed, notably in Arabic, which exhibited lower agreement scores for tone and sentiment, in contrast to the near-perfect scores for the other languages. This variation highlights potential challenges in machine translation for low-resourced languages with limited training data.

LLM-based Topic Generation. Our user study validated the relevance and comprehensiveness of LLM-generated topics. Despite concerns about LLM hallucinations, recent research (Zhang et al. [50] and Pu et al. [32]) indicates GPT's text summarization capabilities are comparable to human performance, with similar proportions of 'extrinsic' hallucinations.

6.2 Takeaways for Stakeholders

Our analysis reveals a significant shift in privacy-related reporting, extending the conversation beyond Western borders to highlight the active engagement and concerns of the Global South. This engagement, marked by notable legal victories and vigorous privacy activism, signals a move towards a globally empowered civil society keenly aware of its digital rights. Such a shift not only challenges the traditional Western-centric narrative of privacy but also calls for the development of privacy policies that are truly inclusive, acknowledging the diverse cultural contexts and legal frameworks across the globe. The evolving narrative of privacy, now embracing a global viewpoint, stresses the need for universally relevant policies and dialogues, fostering a sense of digital solidarity that bridges economic and technological divides, pointing towards global digital solidarity. That said, the regional disparities in privacy coverage are crucial for policymakers and privacy advocates, as they highlight the need for more inclusive, culturally nuanced and globally representative privacy policies and discussions.

The landscape of privacy concerns has evolved far beyond the initial worries over government surveillance and data breaches, delving into deeply personal and distressing areas such as Child Sexual Abuse Material (CSAM), Intimate Partner Violence (IPV), and a myriad of online abuses. The broadening scope of privacy discourse highlights the urgent need for robust support for individuals at risk, while upholding the integrity of privacy for all. Stakeholders, including policymakers and technologists, are called upon to collaboratively design laws and technologies that address the full spectrum of digital harms without infringing on individual rights.

The discourse on corporate responsibility and public perception of tech companies is evolving, driven by significant incidents like Apple's CSAM scanning reversal, WhatsApp's privacy policy upheaval, and Facebook's data breaches, showcasing the influence of end users and privacy activists. Reporting around these events emphasizes the need for transparency, ethical data management, and user consent, urging companies to prioritize privacy and security to build user trust. Through these recent developments, privacy advocates and users have demonstrated their power to effect change.

The consistent negative sentiment in privacy-related media coverage signals a profound public concern and a general mistrust towards institutions handling personal data. Recognizing and addressing this sentiment trend is vital for stakeholders. Tech corporations should prioritize building public trust through enhanced transparency and accountability in data handling practices. Similarly, regulators and lawmakers are tasked with a critical role in clarifying data usage policies, enhancing consent protocols, and enforcing stricter data protection regulations across both public and private sectors. Together, these efforts can bridge the trust gap, ensuring that the guardians of personal data are perceived as responsible and trustworthy stewards in the eyes of the public.

6.3 Directions for Future Work

To deepen our understanding of the dynamic nature of the privacy discourse, future research should expand its analytical lens beyond traditional news outlets to encompass a diverse array of platforms, including social media, blogs, and forums. Such an expansion is crucial for capturing the multifaceted ways in which privacy concerns manifest and evolve across news consumption mediums.

An essential avenue for enriching privacy research lies in fostering interdisciplinary collaborations. By bringing together expertise from legal studies, sociology, computer science, and beyond, researchers can construct a more nuanced picture of the regulatory changes, media narratives, and societal impacts surrounding privacy issues. These collaborative efforts promise to reveal the complex interplay between technological advancements, legislative frameworks, and public discourse.

Future research could refine our understanding of privacy discourse by applying advanced methods like Interrupted Time Series (ITS) analysis and quasi-experimental designs to delineate and study the impact of significant incidents over time.

7 CONCLUSION

Our study offers a global, longitudinal view of the privacy discourse evolution, marking a shift from government surveillance to data breaches, and intensifying focus on online abuse and corporate accountability. The study extends the dialogue to the Global South, challenging the prevailing Western-centric privacy narrative and advocating for globally inclusive and culturally attuned privacy policies. The pervasive negative sentiment in media coverage signals a deep-seated public mistrust towards organizations handling personal data, emphasizing a critical demand for enhanced transparency and accountability in data practices. This observation emphasizes the need for policymakers, tech companies, and regulators to create trust-building strategies that strike a balance between protecting individual rights and fostering technological advancement.

ACKNOWLEDGMENTS

We express our appreciation to Raluca-Georgia Diugan, Yashaswi Malla, and Shantanu Bhatia for their involvement at various stages of this research and their contributions to the experimental setups. Special thanks to Nizar Habash, whose expertise and guidance were critical in refining our methodologies for the evaluation of machine translations. This work was supported by the Center for Cyber Security at New York University Abu Dhabi (NYUAD). We acknowledge the LexisNexis REST API and NYU Libraries for providing access to the data that was essential to our research.

REFERENCES

[1] Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Cojocaru, et al. 2023. The falcon series of open language models. *arXiv preprint arXiv:2311.16867* (2023).

[2] Annie I Antón and Julia B Earp. 2004. A requirements taxonomy for reducing Web site privacy vulnerabilities. *Requirements Engineering* 9, 3 (2004), 169–185.

[3] Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. Findings of the 2019 Conference on Machine Translation (WMT19). ACL, Florence, Italy. https://doi.org/10.5167/uzh-176407

[4] Jens Branum and Jonathan Charteris-Black. 2015. The Edward Snowden affair: A corpus study of the British press. *Discourse & Communication* 9, 2 (2015).

[5] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* (2023).

[6] Yin-Wong Cheung and Kon S Lai. 1995. Lag order and critical values of the augmented Dickey–Fuller test. *Journal of Business & Economic Statistics* 13, 3 (1995), 277–280.

[7] Mengyao Cui et al. 2020. Introduction to the k-means clustering algorithm based on the elbow method. *Accounting, Auditing and Finance* 1, 1 (2020), 5–8.

[8] Sauvik Das, Joanne Lo, Laura Dabbish, and Jason I Hong. 2018. Breaking! A typology of security and privacy news and how it's shared. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.

[9] Ralf De Wolf and Stijn Joye. 2019. Control Responsibility: The Discursive Construction of Privacy, Teens, and Facebook in Flemish Newspapers. *International Journal of Communication* 13 (2019), 20.

[10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[11] Philip Di Salvo and Gianluigi Negro. 2016. Framing Edward Snowden: A comparative analysis of four newspapers in China, United Kingdom and United States. *Journalism* 17, 7 (2016), 805–822.

[12] Benjamin D Douglas, Patrick J Ewell, and Markus Brauer. 2023. Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *Plos one* 18, 3 (2023), e0279720.

[13] James N. Druckman and Michael Parkin. 2005. The Impact of Media Bias: How Editorial Slant Affects Voters. *The Journal of Politics* 67, 4 (2005), 1030–1049.

[14] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. Chatgpt outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv:2303.15056* (2023).

[15] Fan Huang, Haewoon Kwak, and Jisun An. 2023. Is ChatGPT better than human annotators? potential and limitations of ChatGPT in explaining implicit hate speech. *arXiv preprint arXiv:2302.07736* (2023).

[16] Maurice George Kendall. 1948. Rank correlation methods. (1948).

[17] Brian Hyeongseok Kim, Shujaat Mirza, and Christina Pöpper. 2023. Extending Browser Extension Fingerprinting to Mobile Devices. In *Proceedings of the 22nd Workshop on Privacy in the Electronic Society*. 141–146.

[18] Kathleen M Kuehn. 2018. Framing mass surveillance: Analyzing New Zealand's media coverage of the early Snowden files. *Journalism* 19, 3 (2018), 402–419.

[19] Taja Kuzman, Nikola Ljubešić, and Igor Mozetič. 2023. Chatgpt: beginning of an end of manual annotation? Use case of automatic genre identification. *arXiv preprint arXiv:2303.03953* (2023).

[20] Susan Landau. 2013. Making sense from Snowden: What's significant in the NSA surveillance revelations. *IEEE Security & Privacy* 11, 4 (2013), 54–63.

[21] Carl Landwehr. 2019. 2018: A Big Year for Privacy. *Commun. ACM* 62, 2 (2019).

[22] Stanislav Mamonov and Marios Koufaris. 2016. The impact of exposure to news about electronic government surveillance on concerns about government intrusion, privacy self-efficacy, and privacy protective behavior. *Journal of Information Privacy and Security* 12, 2 (2016), 56–67.

[23] Henry B Mann. 1945. Nonparametric tests against trend. *Econometrica: Journal of the econometric society* (1945), 245–259.

[24] Maxwell E McCombs and Donald L Shaw. 1972. The agenda-setting function of mass media. *Public opinion quarterly* 36, 2 (1972), 176–187.

[25] Shujaat Mirza and Christina Pöpper. 2021. My Past Dictates my Present: Relevance, Exposure, and Influence of Longitudinal Data on Facebook. *Proceedings of the Workshop on Usable Security and Privacy* (2021).

[26] Victor Morel and Raúl Pardo. 2020. SoK: Three facets of privacy policies. In *Proceedings of the 19th Workshop on Privacy in the Electronic Society*. 41–56.

[27] Ethan Morrow. 2022. Priming Privacy: The Effect of Privacy News Consumption on Privacy Attitudes, Beliefs, and Knowledge. *Journal of Broadcasting & Electronic Media* 66, 5 (2022), 772–793.

[28] Stefan Palan and Christian Schitter. 2018. Prolific. ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance* 17 (2018), 22–27.

[29] Nicholas Pangakis, Samuel Wolken, and Neil Fasching. 2023. Automated Annotation with Generative AI Requires Validation. *arXiv preprint arXiv:2306.00176* (2023).

[30] Nicolas Papernot, Patrick McDaniel, Arunesh Sinha, and Michael P Wellman. 2018. Sok: Security and privacy in machine learning. In *2018 IEEE European symposium on security and privacy (EuroS&P)*. IEEE, 399–414.

[31] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.

[32] Xiao Pu, Mingqi Gao, and Xiaojun Wan. 2023. Summarization is (almost) dead. *arXiv preprint arXiv:2309.09558* (2023).

[33] Ramzi Rizk, Daniel Marx, Matthias Schrepfer, Janina Zimmerman, and Oliver Guenther. 2009. Media Coverage of Online Social Network Privacy Issues in Germany: A Thematic Analysis. *AMCIS 2009 proceedings* (2009), 342.

[34] Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining*. 399–408.

[35] Peter J Rousseeuw. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics* 20 (1987), 53–65.

[36] Mary Sanford and Taha Yasseri. 2021. The Kaleidoscope of Privacy: Differences across French, German, UK, and US GDPR Media Discourse. *arXiv preprint arXiv:2104.04074* (2021).

[37] Theodor Schnitzler, Shujaat Mirza, Markus Dürmuth, and Christina Pöpper. 2021. Sok: Managing longitudinal privacy of publicly shared personal online data. *Proceedings on Privacy Enhancing Technologies* (2021).

[38] Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. BLEURT: Learning Robust Metrics for Text Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 7881–7892. https://doi.org/10.18653/v1/2020.acl-main.704

[39] Pranab Kumar Sen. 1968. Estimates of the regression coefficient based on Kendall's tau. *Journal of the American statistical association* 63, 324 (1968), 1379–1389.

[40] Karthik Sheshadri, Nirav Ajmeri, and Jessica Staddon. 2017. No (Privacy) News is Good News: An Analysis of New York Times and Guardian Privacy News from 2010–2016. In *2017 15th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 159–15909.

[41] Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*. 223–231.

[42] Daniel J Solove. 2005. A Taxonomy of Privacy. *U. Pa. L. Rev.* 154 (2005).

[43] Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune BERT for text classification?. In *China National Conference on Chinese Computational Linguistics*. Springer, 194–206.

[44] Doris Teutsch and Julia Niemann. 2016. Social network sites as a threat to users' self-determination and security: A framing analysis of German newspapers. *The Journal of International Communication* 22, 1 (2016), 22–41.

[45] Minna Tiainen. 2017. (De) legitimating electronic surveillance: a critical discourse analysis of the Finnish news coverage of the Edward Snowden revelations. *Critical Discourse Studies* 14, 4 (2017), 402–419.

[46] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, et al. 2023. Llama: Open and efficient foundation language models. *arXiv* (2023).

[47] Lucas Nunes Vieira. 2019. Post-editing of machine translation. In *The Routledge handbook of translation and technology*. Routledge, 319–336.

[48] Karin Wahl-Jorgensen, Lucy Bennett, and Gregory Taylor. 2017. The Normalization of Surveillance and the Invisibility of Digital Citizenship: Media Debates After the Snowden Revelations. *International Journal of Communication* 11 (2017).

[49] Rob Whitley and JiaWei Wang. 2017. Good News? A Longitudinal Analysis of Newspaper Portrayals of Mental Illness in Canada 2005 to 2015. *The Canadian Journal of Psychiatry* 62, 4 (2017), 278–285.

[50] Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B Hashimoto. 2023. Benchmarking large language models for news summarization. *arXiv preprint arXiv:2301.13848* (2023).

A LLM-BASED TEXT CLASSIFICATION

Listing 1 provides the input prompt that was used for the GPT-3.5-turbo model for the classification of news articles with privacy focus. To explore the effect of ChatGPT's temperature parameter, which controls the degree of randomness of the output, we experimented on the validation set by varying the temperature between 0 and 2 and recording its impact on the quality of annotations. As observed in Figure 10, a temperature value of 0 yields the least number of misclassifications (i.e., # of false positives + # of false negatives), which is what we utilize to evaluate the model on the test set.

Listing 1: Prompt used for GPT-3.5-turbo model.

```
You are a helpful assistant that takes in a
newspaper article and extracts the following
information:

summary: Extract a summary of the article in
1-2 sentences alone.

keywords: What 3-5 keywords would best describe
the focus of the article?

digital_privacy_focus: Has the article discussed
aspects of digital privacy?
Answer 1 if True, 0 if False or unknown.

argument: Argue succinctly in 1-2 sentences.

Format the output as JSON with the following
keys:
    summary
    keywords
    digital_privacy_focus
    argument

Before you perform the task, revisit your
understanding of the digital privacy concept
and stages of data life cycle by reading
this definition: {definition}
```

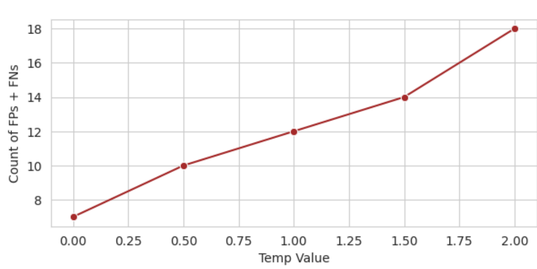

Figure 10: Effect of the temperature parameter on GPT-3.5-turbo model's misclassification (FP + FN) for the privacy filter.

B MACHINE TRANSLATION VALIDATION

To rigorously evaluate the quality of our article translations in each language, we conducted a thorough validation study involving 50 participants. These bilingual annotators, each proficient in English and another language featured in our study, were tasked with refining a random sample of five translated articles per language. They fine-tuned the translations to ensure fidelity in tone and meaning, implementing minimal edits for accuracy. For article selection, we collected a random sampling from the only news source available in the respective language, with the exception of Spanish, where we randomly chose articles from the three sources available: ECP, EUM, LNA.

In Listing 2, we provide an abridged version of the survey that assesses the machine translation reliability.

Listing 2: Survey questionnaire - Post editing task.

```
Edit Translation: Your task is to edit the
    translation making as few changes as possible
    so that it matches the meaning, tone and
    sentiment of the text in original language.

Meaning: How accurately does the translation
    convey the meaning of the original text?
    Please rate on a scale from 1 (Not at all
    accurate) to 5 (Extremely accurate).

Tone & Sentiment: Does the translation maintain
    the tone & sentiment of the original text?
    Rate its effectiveness from 1 (Completely
    different tone) to 5 (Perfectly maintains tone
    ).

Naturalness & Fluency: How natural and fluent does
     the translated text sound in English? Rate
    from 1 (Very unnatural) to 5 (
    Indistinguishable from native English).

Grammatical Correctness: Assess the grammatical
    correctness of the translation. Rate from 1 (
    Many errors) to 5 (Free of errors).

Before & After Editing: How much improvement do
    you perceive in the translation after your
    edits? Rate from 1 (No improvement) to 5 (
    Significant improvement).

Compared to Expectations: How did the translation
    quality compare to your expectations? Rate
    from 1 (Far below expectations) to 5 (Exceeded
     expectations).

Please provide any additional comments or
    observations about the translation quality
```

The study's annotators hailed from diverse linguistic backgrounds, with 72% being native speakers of the language they reviewed. Their

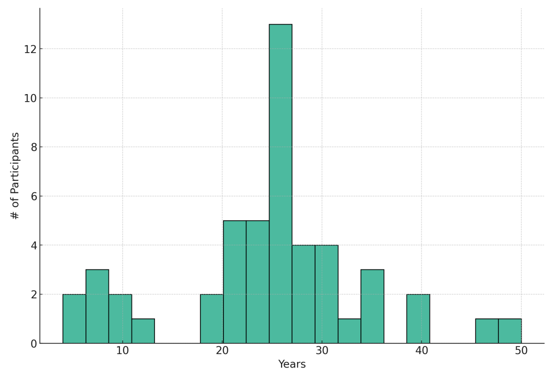

Figure 11: Annotator Language Experience. (Min: 4. Max: 50. Mean: 25)

expertise was crucial in ensuring the reliability of our translations, as depicted in Figure 11, which details their years of language experience.

The participants then rated the initial translations on a scale of 1 to 5—where a higher score denoted better quality—across four dimensions: accuracy, tone and sentiment, naturalness and fluency, and grammatical correctness. The aggregated results of these assessments are presented in Figure 12.

C LLM GENERATED TOPICS VALIDATION

In Listing 3, we provide an abridged version of the survey that assesses the relevance and comprehensiveness of focus topics assigned by GPT-3.5. Figure 13 shows the most frequently occurring tech companies featured as the main subjects in articles.

Listing 3: Survey questionnaire - LLM generated focus topics.

```
Task: For the article you have read,please review
    the list of keywords/topics below that have
    been generated to capture its focus. Your task
     is to assess the relevance ofthese keywords/
    topics to the article and suggest any
    improvements.

Please assess the relevance of the following
    topics to the article on a scale where 1 is "
    Not Relevant" and 5 is "Relevant". Please use
    the scale to indicate how relevant you find
    each topic to the article.

For any topics marked as "Not Relevant," please
    explain your decision.

Comprehensiveness: Do you feel that the provided
    focus topics/keywords comprehensively cover
    the key points of the article? Please rate
    from 1 (Not at all comprehensive) to 5 (Highly
     comprehensive).
```

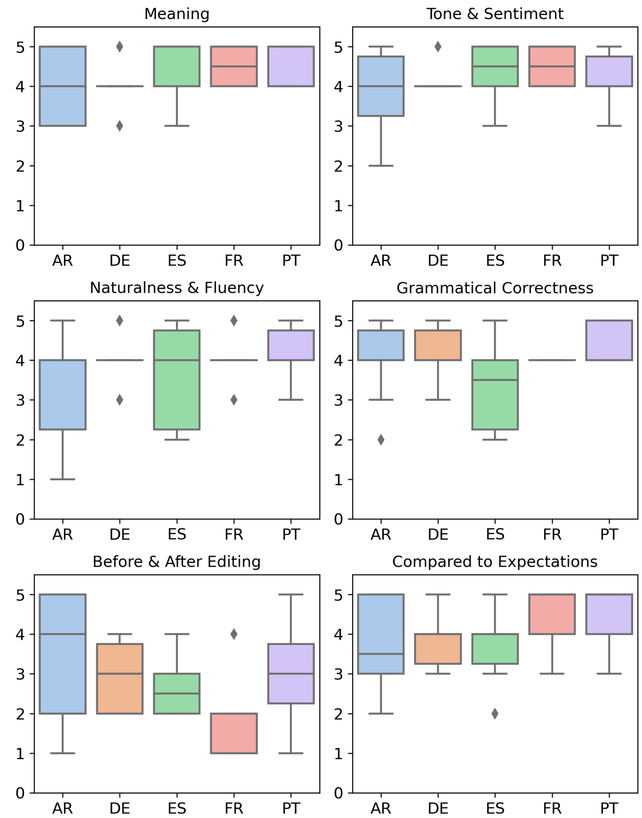

Figure 12: Translation Metrics Across Languages. Singular lines indicate the span of the Interquartile Range (IQR) falls on one value. Outliers, represented as dots, are defined as values that fall beyond $1.5 \times$ IQR.

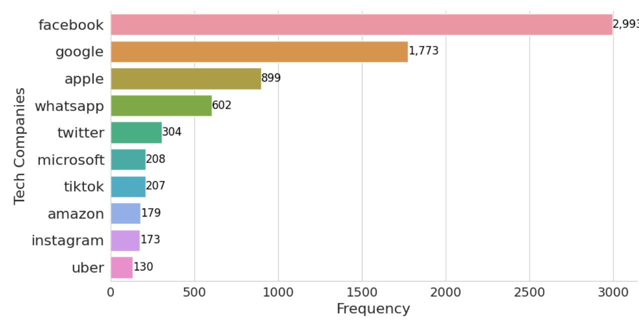

Figure 13: Tech Companies by Focus.

D NOTION OF PRIVACY

We utilize conceptions of privacy put forth by Solove [42] and Antón & Earp [2] to formulate definitions of privacy that guide the process of filtering articles for privacy. Figure 14 details the definition that was generated based on both taxonomies. The definition served as reference for us throughout this work and was provided to LLM as context too in its input prompt.

Figure 14: Definition of Digital Privacy based on Solove's [42] and Antón & Earp [2] taxonomies that are provided to GPT-3.5-turbo model as part of the prompt.

```
In the digital realm, privacy is primarily concerned with the protection and management of personal data. This encompasses
the right or ability of individuals to control the collection, processing, and dissemination of their personal information
by digital platforms and services. One work categorizes intrusions into digital privacy into four main groups: information
collection, information processing, information dissemination, and invasion, as detailed below:

Information Collection deals exclusively with privacy problems resulting from gathering information.
- Surveillance consists of methods of watching, listening & recording a subject's activities.
- Interrogation describes methods used to ask or elicit information from a subject.

Information Processing describes methods to store, modify or manipulate a subject's information.
- Aggregation combines individual and previously separate pieces of data about a subject.
- Identification depicts an organization's methods for determining which individual is described by a set of data.
- Insecurity is a failure to properly protect stored data.
- Secondary Use reflects the use of data for a purpose other than that for which it was originally provided.
- Exclusion is inability of a subject to have knowledge of how their data is being used.

Information Dissemination consists of privacy harms resulting from the release of information about a subject.
- Breach of Confidentiality contains those harms based on the violation of a trust agreement to maintain confidentiality of
   a subject's information.
- Disclosure describes harms related to release of truthful information about a data subject.
- Exposure describes the dissemination of information about a subject's grief, body or bodily functions.
- Increased Accessibility consists of the ways that a subject's public information may be made available to a wider audience
   than before.
- Blackmail involves a threat made to a subject about potential release of their information.
- Appropriation describes the use of a subject's identity or information to serve the purposes of the organization rather
   than the subject.
- Distortion consists of harms related to release of falsified information about a subject.

Invasion consists of the various intrusions on an individual's private life.
- Intrusion is a form of invasion to describe all harms resulting from the disturbance of an individual's peace & solitude.
- Decisional Interference is an invasion into a subject's decisions about their private affairs.

Another taxonomy was developed by applying grounded theory to online privacy policies revealing 12 categories of privacy
elements spread across two broad classifications (Privacy protection & Vulnerabilities), as shown below:

Privacy Protection Goals safeguard the privacy of a customer's data and there are five categories as follows:
- Notice and Awareness goals describe how a customer is informed about an organization's practices regarding their data.
- Choice & Consent goals describe a customer's ability to choose how they want their data to be managed by an organization.
- Access & Participation reflects a customer's ability to challenge, correct or modify their data as used by an organization.
- Integrity & Security goals describe measures an organization takes to protect the accuracy & security of a customer's data.
- Enforcement & Redress goals describe the ways that organization  approaches internal policy violations by their employees.

Vulnerabilities reflect a potential privacy violation and there are 7 categories as follows:
- Information Monitoring describes how an organization tracks customers' interaction with their website.
- Information Aggregation reflects the ways that an organization will combine customer data with third-party data sources.
- Information Storage reflects an organization's practices regarding what/how customer records are stored in the
  organization's database.
- Information Transfer describes how an organization may share their collected customer information with affiliates and
   third-parties.
- Information Collection shows what types of information an organization may collect and how that organization collects the
    specified information.
- Information Personalization reflects the methods an organization uses to tailor the presentation of their website to their
    customers.
- Solicitation shows the purposes and methods an organization would use to contact their customers.

Key to these definitions is the concept of informed consent, where individuals have right to know what data is being
collected about them, how it's being used, and who it's being shared with. Beyond just being informed, individuals should
also have the ability to prevent, restrict, or alter the collection, use, or sharing of their personal information.
```

E DUPLICATE REMOVAL

To determine duplicates we aggregate articles by week of publishing and compare their titles pair-wise. We capture small editorial changes, besides perfect overlaps, by applying the cosine similarity over each set of titles. To set the similarity threshold, we analyzed data from two randomly chosen time periods of six months each (Jul 1st–Dec 31st 2015 and Jan 1st–Jun 30th 2018) and aggregated the encountered similarity scores.

Figure 15 captures the distribution of similarity scores for the second time period; a similar trend was observed for the first time range. With a threshold of .5 usually indicating sufficient similarity between documents, we further tuned our threshold by manually investigating title pairs with similarity scores in the .5–.7 range. Differences between titles stem from minor punctuation or spelling fixes, rewordings, or extensions of the titles. Other pairs in the range were reporting on similar issues, hence the higher overlap in titles. Following this investigation, we set our similarity threshold conservatively to $\theta = .7$ to avoid mislabeling articles that report on similar issues within the same week as duplicates.

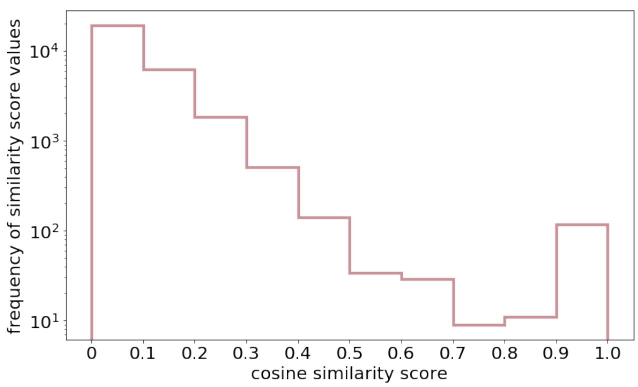

Figure 15: Aggregated similarity scores of article titles published within a calendar week for all newspapers in the time range Jan 1st–Jun 30th 2018. Based on our related analysis, we selected a similarity score of .7 as indicative of duplicates.

F CLUSTERING ANALYSIS

To group the regions based on the similarity of their time series data, we performed time series clustering using a k-means clustering algorithm, which resulted in three clusters. The objective was to identify patterns and similarities in the coverage of privacy-related topics across different regions. To determine the optimal number of clusters, we used two empirical methods: the Elbow method and Silhouette analysis. Figure 16 shows experiments to determine the optimal number of clusters empirically for K-means clustering of privacy-related coverage time series using the Elbow method (Left) and Silhouette analysis (Right). The Elbow method indicates that the optimal number of clusters is either 3 or 4, as the sum of squared distances starts to decrease more slowly after this point, forming an 'elbow' in the curve. On the other hand, Silhouette Analysis suggests that 2 or 3 clusters provide the highest Silhouette scores, indicating a better separation of data points within clusters. The regions in the same cluster have more similar time series data compared to regions in different clusters. Asia has been grouped into its own cluster. This is likely due to its unique pattern compared to other regions, as observed in Figure 3, as the volume of privacy articles increases significantly over time. Africa and Latin America are grouped together in a separate cluster as evidenced by their similar coverage. Oceania, Europe, and Americas - Northern are grouped in their own cluster.

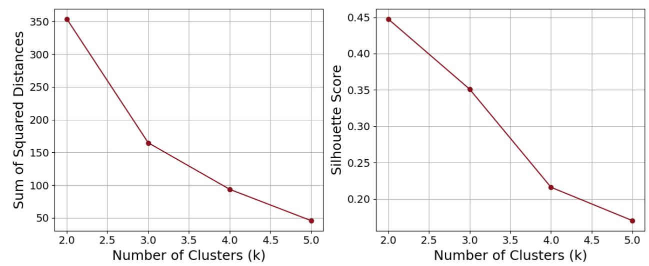

Figure 16: Determine the optimal number of clusters empirically for K-means clustering of privacy-related coverage time series using the Elbow method (Left) and Silhouette analysis (Right). The Elbow method [7] indicates that the optimal number of clusters is either 3 or 4, as the sum of squared distances starts to decrease more slowly after this point, forming an 'elbow' in the curve. On the other hand, Silhouette analysis [35] suggests that 2 or 3 clusters provide the highest Silhouette scores, indicating a better separation of data points within clusters.

Figure 17: Word cloud of focus topics.

G REGIONAL ANALYSIS

Figure 18 presents a stacked area plot comparing topic popularity over time between the Global North and Global South. Each colored region in the plot represents a distinct topic, with the height of each region at any given time indicating its relative popularity.

Figure 19 is a heat map of privacy topics by region, using a color gradient for correlation strength; darker shades represent stronger correlations.

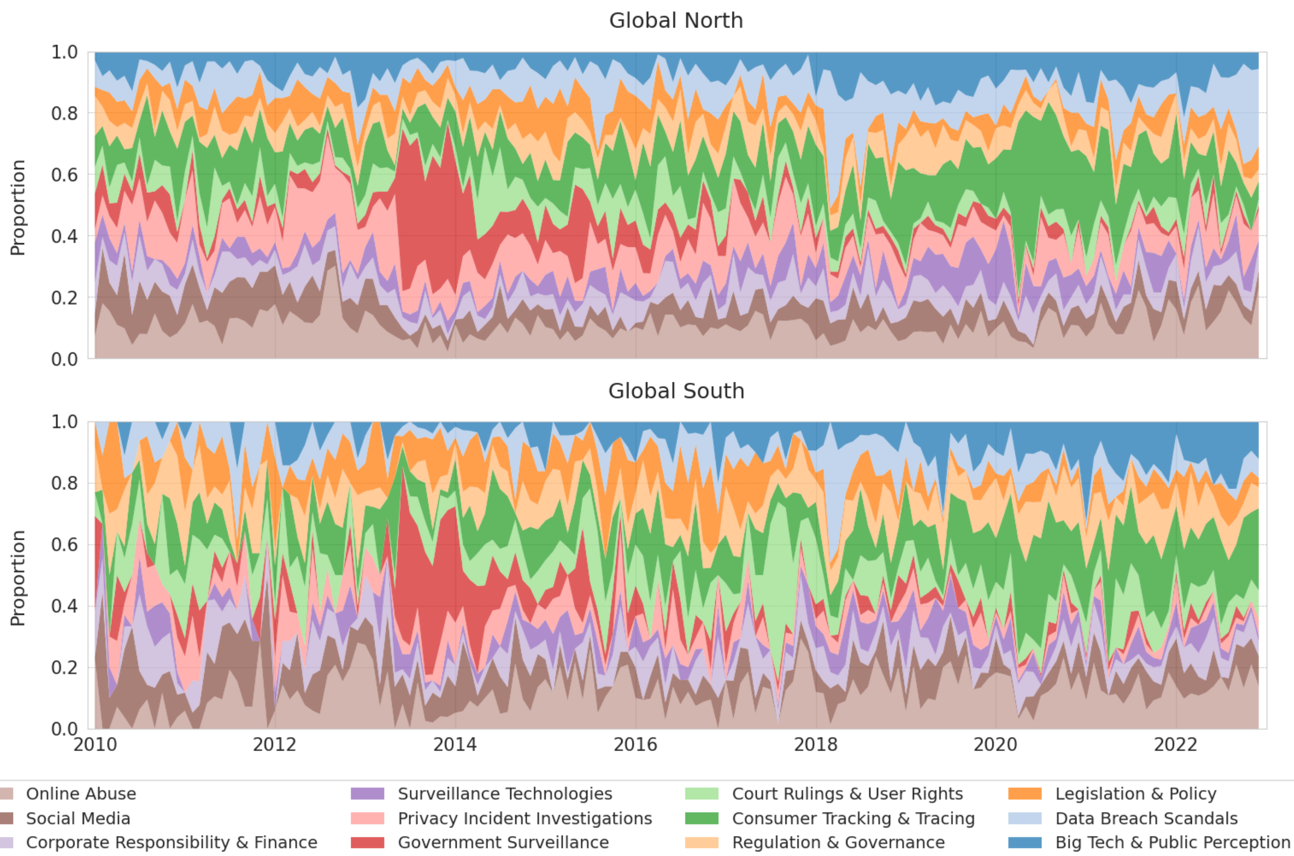

Figure 18: Stacked Area Plot illustrating the changing popularity of topics over time in the Global North and Global South. Each colored region represents a distinct topic, and the height of a region at any given time point reflects the proportion of articles dedicated to that topic during that period.

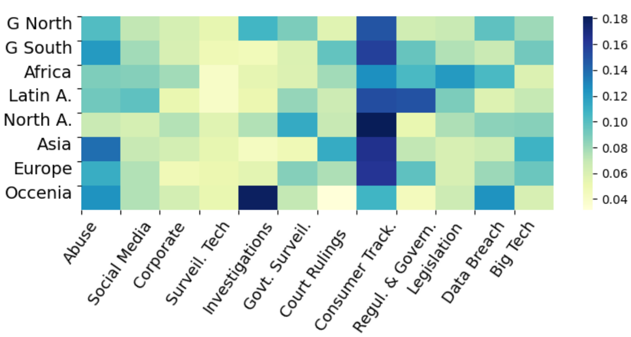

Figure 19: Heat-map of privacy topics by region, using a color gradient for correlation strength; darker shades represent stronger correlations.

H TONE ANALYSIS

Figure 20 shows the distribution of sentiments across regions and languages present in our dataset. Figure 21 provides a time-series representation of the average monthly scores for five key emotions – joy, sadness, fear, anger, and disgust – across two regions: Asia and Africa.

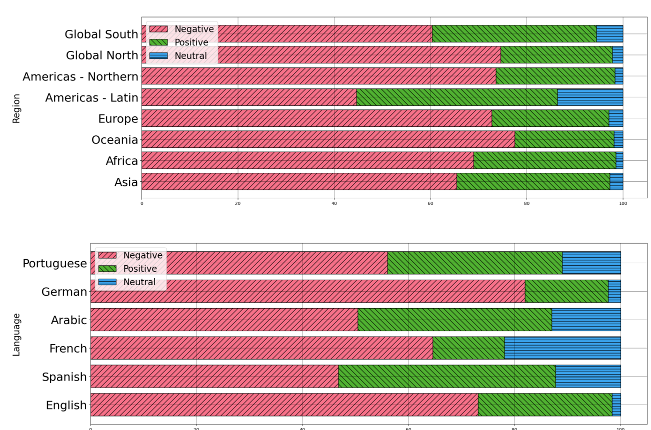

Figure 20: Distribution of sentiments across regions (top) and languages (bottom) present in our dataset.

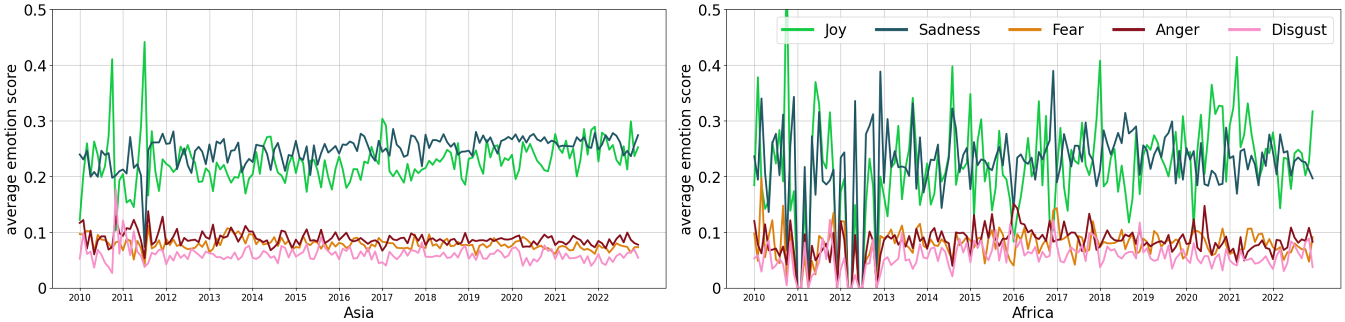

Figure 21: Time-series representation of the average monthly scores for five key emotions – joy, sadness, fear, anger, and disgust – across two regions: Asia and Africa.

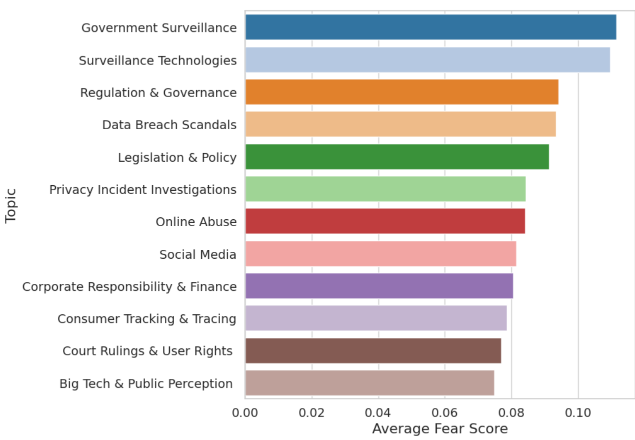

Figure 22: Average fear score per topic.

Figure 22 captures fear scores for different topics in our LDA model.

I ADDITIONAL DETAILS

Figure 23 shows the distribution of LDA topic probabilities whereas figure 24 depicts accuracy and average training loss on validation set over 10 epochs during BERT model fine-tuning for the privacy filter.

Table 5 lists a breakdown of the number of articles on digital privacy per newspaper per year.

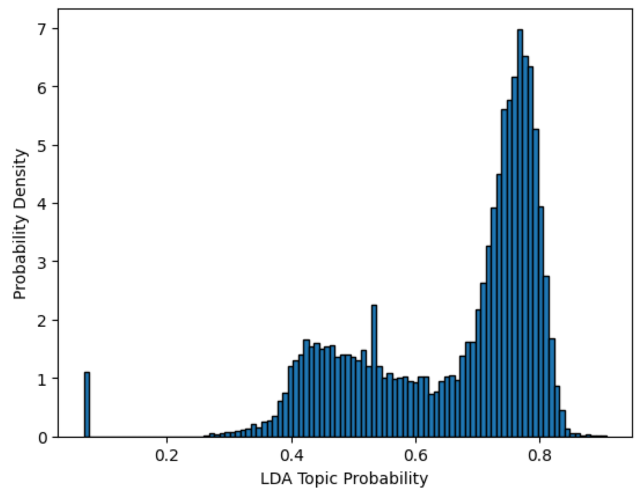

Figure 23: Distribution of LDA Topic Probabilities

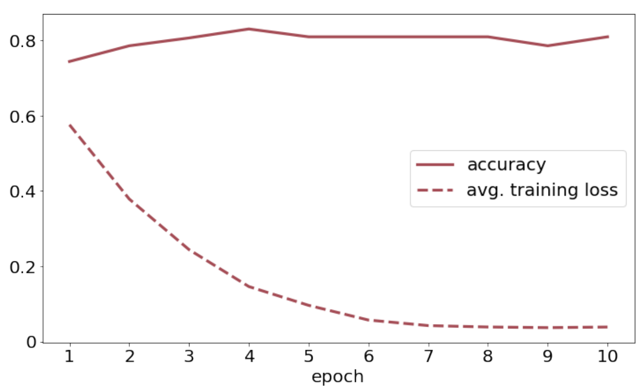

Figure 24: Accuracy and average training loss on validation set over 10 epochs during BERT model fine-tuning for the privacy filter.

Table 5: Number of Articles on Digital Privacy: A Breakdown by Year and Newspaper

ID	Name	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	Total
TOI	The Times of India	69	104	152	259	189	236	197	352	380	490	644	616	497	4185
NZH	New Zealand Herald	194	192	340	508	303	118	124	154	208	234	220	179	127	2901
FTL	Financial Times	149	155	157	316	266	198	171	163	351	347	245	243	133	2894
TS	The Toronto Star	169	178	160	216	429	254	257	134	305	303	202	138	92	2837
DT	The Daily Telegraph	169	185	174	195	182	142	154	118	281	260	279	233	122	2494
NYT	The New York Times	202	143	151	287	244	215	200	81	152	164	96	97	68	2100
SMH	Sydney Morning Herald	136	142	143	127	150	134	131	89	175	131	131	88	106	1683
SZG	Süddeutsche Zeitung	146	79	71	172	157	113	117	100	125	114	99	83	45	1421
USA	USA Today	105	103	93	177	123	98	138	80	119	96	96	60	55	1343
AFR	Financial Review	72	70	78	59	115	71	72	93	139	132	130	93	143	1267
TDP	The Dominion Post	41	44	161	159	99	72	92	100	117	107	90	97	48	1227
EPS	El Pais	108	101	102	175	113	69	67	55	100	97	96	56	48	1187
HSM	Herald Sun	139	118	101	69	68	65	60	76	143	95	77	77	72	1160
FPC	Financial Post	39	32	31	44	80	94	133	80	177	135	99	85	26	1055
ESP	O Estado de S.Paulo	106	72	90	191	119	42	28	28	35	55	58	57	31	912
EUM	El Universal	46	57	57	39	67	43	59	19	9	53	108	113	225	895
CD	China Daily	0	0	0	39	96	93	75	80	77	86	101	57	45	749
KT	Khaleej Times	9	10	20	17	21	23	7	24	124	121	78	93	54	601
MSN	Manawatu Standard	4	4	54	84	39	37	46	69	61	45	39	51	17	550
LNA	La Nacin	18	21	13	17	37	21	20	2	2	15	51	68	255	540
EEM	El Economista	0	0	0	0	0	0	0	0	102	82	87	126	96	493
THU	The Hill	9	47	33	49	37	31	19	20	32	48	11	15	9	360
DN	Dawn	0	0	0	24	23	34	37	19	38	52	35	40	53	355
DNK	Daily Nation	0	0	0	25	37	16	26	27	80	62	44	14	17	348
AAA	Asharq Alawsat	0	0	12	11	19	17	24	23	70	62	41	36	24	339
LPC	La Presse Canadienne	33	7	18	12	12	17	13	17	34	23	25	19	26	256
LFF	Le Figaro	18	13	19	26	23	10	12	5	36	22	25	30	13	252
TDM	The Daily Monitor	0	0	0	8	15	11	12	13	35	22	18	52	33	219
ECP	El Comercio	14	12	15	2	9	15	7	2	13	10	10	8	74	191
DNE	Daily News Egypt	11	7	4	13	14	10	15	5	37	14	15	8	8	161
TMT	The Moscow Times	3	15	4	29	18	18	4	7	21	10	14	12	4	159
TSN	The Sun	0	0	0	7	5	5	5	13	14	23	21	29	34	156
TDL	This Day (Lagos)	5	3	6	18	7	12	6	5	4	8	20	16	5	115
CAN	Caribbean News	0	0	1	9	3	4	15	12	20	9	8	4	11	96
BFT	Business & Finan. Times	0	0	0	0	0	0	5	7	5	15	26	14	17	89
NA	Nikkei Asia	2	3	8	5	1	1	3	4	7	10	9	10	2	65