

Identifying and Investigating Global News Coverage of Critical Events Such as Disasters and Terrorist Attacks

Erica Cai^{1†}, Xi Chen^{1†}, Reagan Grey Keeney¹, Ethan Zuckerman¹, Brendan O’Connor¹, Przemyslaw A. Grabowicz^{1,2}

¹University of Massachusetts Amherst

²University College Dublin

ecai@umass.edu, xchen4@umass.edu, reagankeeney@gmail.com, ethanz@umass.edu, brenocon@cs.umass.edu, grabowicz@cs.umass.edu

Abstract

Comparative studies of news coverage are challenging to conduct because methods to identify news articles about the same event in different languages require expertise that is difficult to scale. We introduce an AI-powered method for identifying news articles based on an event *fingerprint*, which is a minimal set of metadata required to identify critical events. Our event coverage identification method, FINGERPRINT TO ARTICLE MATCHING FOR EVENTS (FAME), efficiently identifies news articles about critical world events, specifically terrorist attacks and several types of natural disasters. FAME does not require training data and is able to automatically and efficiently identify news articles that discuss an event given its fingerprint: time, location, and class (such as storm or flood). The method achieves state-of-the-art performance and scales to massive databases of tens of millions of news articles and hundreds of events happening globally. We use FAME to identify 27,441 articles that cover 470 natural disaster and terrorist attack events that happened in 2020. To this end, we use a massive database of news articles in three languages from MediaCloud, and three widely used, expert-curated databases of critical events: EM-DAT, USGS, and GTD. Our case study reveals patterns consistent with prior literature: coverage of disasters and terrorist attacks correlates to death counts, to the GDP of a country where the event occurs, and to trade volume between the reporting country and the country where the event occurred. We share our NLP annotations and cross-country media attention data to support the efforts of researchers and media monitoring organizations.

1 Introduction

News organizations help shape the public’s perception of disasters and terrorist attacks by choosing how, and how much, to discuss them. The coverage an event receives influences what events the public concerns itself with, a phenomenon called “agenda setting”.

Previous research suggests that coverage biases — increased coverage of certain events based on the political or economic characteristics of the location in which they transpired — have significant effects on media attention toward natural disasters (Eisensee and Strömberg 2007), civil war (Baum and Zhukov 2015), terrorism (L. Hellmueller and

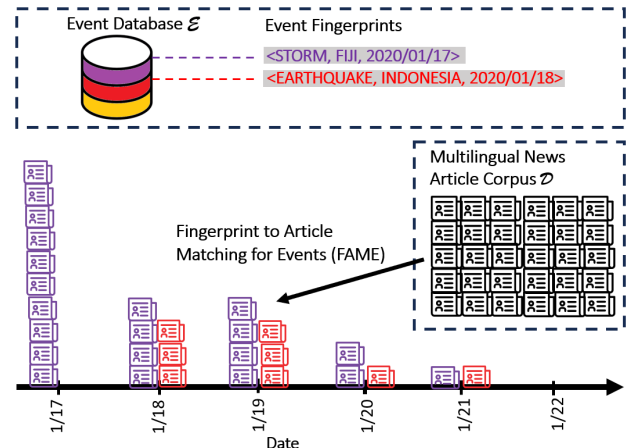


Figure 1: Our event linking method maps events in a knowledge base to articles that discuss corresponding events, at the granularity of their event fingerprints (EventClass (c), CountryLocation (l), Date (t)). The figure shows the real distribution over time of articles corresponding to event fingerprints \langle STORM, FIJI, 2020/01/17 \rangle and \langle EARTHQUAKE, INDONESIA, 2020/01/18 \rangle (purple and red, respectively).

Lindner 2022), international political violence (Dietrich and Eck 2020), etc.—for example, a disaster in Africa may require a death toll 40 times higher than one in Eastern Europe to receive equivalent US television news coverage, and such skews in media attention even affect the amount of provided humanitarian aid (Eisensee and Strömberg 2007).

Comprehensive investigation of biases in global news coverage is challenging in part because it requires efficient, accurate identification of news articles about the same events in different languages. Manual article labeling, prevalent in social science research, is costly and difficult to accomplish on a large scale, limiting the size of corpora that can be studied (Krippendorff 2019; Grimmer, Roberts, and Stewart 2022). On the other hand, automated natural language processing (NLP) methods for cross-document event analysis often require vast amounts of labeled data for training, which is costly (Eirew, Caciularu, and Dagan 2022; Eirew, Cattán, and Dagan 2021), or they have task specifications incompatible with ours — for example, operating only on

single sentences (Huang et al. 2018; Zhang et al. 2022).

In this paper, we introduce a novel task of linking events from an event database with news articles. Then, we contribute a novel method, FAME, which efficiently identifies full articles about specific events across different languages. This method leverages a minimal set of information from a database about an event—date, country location, and event class—as a *fingerprint* and accurately identifies relevant articles without requiring training data (Fingerprint to Article Matching for Events). Our multilingual approach includes two steps: matching on keywords and publication date to assemble candidate articles per event, followed by a more refined semantic filter using large language model question-answering (LLM QA) to eliminate false positives.

We rigorously evaluate FAME’s precision and recall for linking news articles in three languages against events of eight classes (terrorist attacks, and seven types of natural disasters) from three different databases. Unlike conventional NLP techniques that rely on large training datasets, our method achieves superior performance, with average F1 of 94.0 on linking events of the eight classes to English articles, without any training data by using GPT-3.5-turbo (Ouyang et al. 2022). We collect manual annotations across dozens of events, and thousands of articles, finding a significant variation in performance between events. We release data of 6,468 annotated event-news pairs to support the development of future event-to-news linking methods.³ To facilitate replication of this study we release its code.⁴

In a pilot study applying FAME to a massive corpus of tens of millions of multilingual news articles, we identified news in English, Spanish, and French about 470 disaster and terrorist attack events happening globally in the first six months of 2020. Through regression analysis, we quantify the role of various country-level factors for media coverage of disasters and terrorist attacks, spanning events in over a hundred countries. We find that the number of deaths associated with an event and the economic characteristics of the affected countries correlate most strongly with media coverage, consistent with previous literature (Eisensee and Strömberg 2007; Chen et al. 2024), underscoring the validity of this methodological approach. In comparison to existing media studies, our research examines patterns in news coverage of a large number of critical events across 128 countries, while adding nuance to existing knowledge of biases in news coverage. For instance, our findings suggest that GDP, rather than geography, explains decreased news coverage of disasters in Africa. We anticipate future studies to build upon these results by applying FAME to datasets spanning much longer time periods.

2 Related Work

Extracting articles discussing events. Here, we describe the different approaches to event extraction that have been developed in the areas of NLP and social sciences.

On the one hand, many information extraction methods from the NLP literature can enable event extraction, but

most face significant challenges when deployed at scale. High performing methods for event extraction require large amounts of training data (Wadden et al. 2019; Chen et al. 2020; Du and Cardie 2020; Liu et al. 2020; Li et al. 2020; Lin et al. 2020; Li, Ji, and Han 2021; Ahmad, Peng, and Chang 2021; Lu et al. 2021), and existing zero-shot methods tend to perform poorly and lack robustness (Huang et al. 2018; Zhang, Wang, and Roth 2021; Lyu et al. 2021; Zhang et al. 2022; Cai and O’Connor 2024). Further, most existing methods aim to extract events from individual sentences, while real-world applications require extractions from paragraphs or full documents.

A relatively small NLP literature has examined extracting events as they are discussed *across* many documents, which is necessary for media attention analysis. One related task is cross-document event coreference, which involves clustering mentions of the same event across a set of documents. In comparison to our approach, it effectively builds an event database from the corpus. However, performance is typically poor and most methods cannot scale to large datasets (Eirew, Cattan, and Dagan 2021); indeed, many earlier methods are trained and evaluated on extremely small datasets, e.g. 502 articles in the study of Cybulska and Vossen (Chen et al. 2023; Cattan et al. 2021; Barhom et al. 2019; Caciularu et al. 2021; Choubey and Huang 2017; Held, Iter, and Jurafsky 2021; Hsu and Horwood 2022; Allaway, Wang, and Balles-teros 2021). While some newer datasets are larger, they are still orders of magnitude smaller than our article repository (Eirew, Cattan, and Dagan 2021; Eirew, Caciularu, and Dagan 2022).

On the other hand, when social science researchers seek to analyze news bias, they typically use much simpler manual or heuristic methods to identify news articles discussing key events (Eisensee and Strömberg 2007; Mac Carron and Kenna 2013; L. Hellmueller and Lindner 2022; Dietrich and Eck 2020; Baum and Zhukov 2015), which is time consuming and hard to scale. Further, we find that previously used heuristic keyword matching approaches, when adapted to our task, are prone to significant error.

As the field of computational social science advances (Zuckerman 2003; Kwak and An 2014; Hamborg, Donnay, and Gipp 2019), new methods are being used to analyze large datasets, offering a more comprehensive view of international news coverage. Such studies promise to deepen our understanding of the mechanisms behind news selection and dissemination globally. To this end, here we develop and evaluate a novel method for event identification, FAME, bridging the NLP and social science approaches.

Agenda setting and news about disasters and terrorism. The media plays a crucial agenda-setting role in reporting disasters and terrorist attacks. News organizations act as gatekeepers, deciding which events gain public attention, thereby shaping the “agenda” of public concerns. (Shoemaker and Vos 2009; McCombs and Shaw 1972; McCombs and Valenzuela 2021).

Research highlights several factors influencing the news-worthiness of disasters and terrorist attacks. While death tolls are important predictors, proximity—cultural, geographic, and economic—also plays a key role in determin-

³<https://zenodo.org/records/14667152>

⁴https://github.com/social-info-lab/disaster_event_analysis

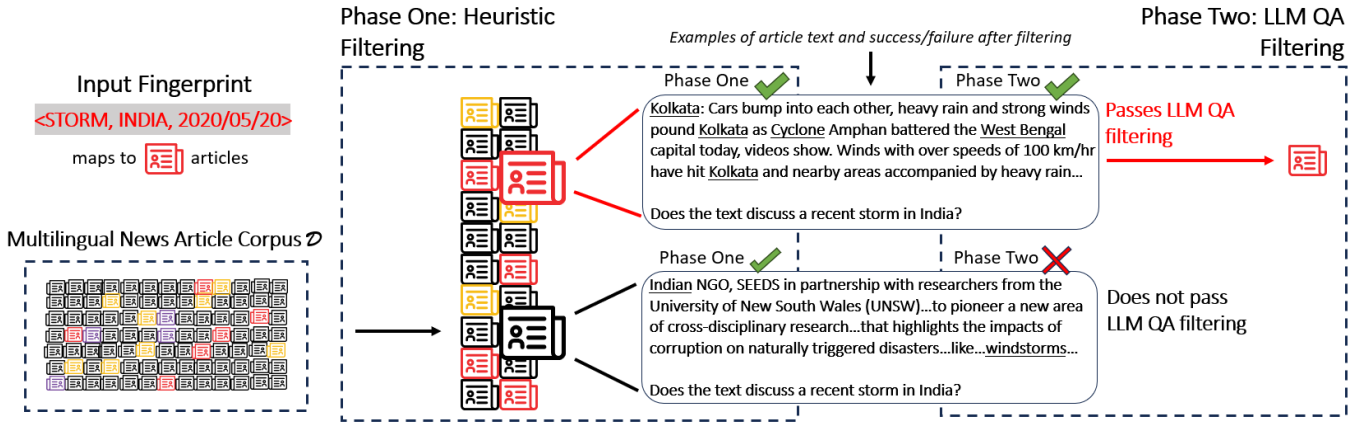


Figure 2: Examples of two articles from our dataset analyzed as candidates to link to specific event fingerprint (STORM, INDIA, 2020/05/20). For Phase One, matched keywords are underlined: from $K_{\text{STORM}} = \{\text{storm, hailstorm, cyclone, snowstorm, typhoon, windstorm, rainstorm...}\} + \text{suffixes}$; and $K_{\text{INDIA}} = \{\text{India, Indian, West Bengal, Uttar Pradesh, Tamil Nadu, Sikkim, Rajasthan, Kolkata, Odisha...}\} + \text{suffixes}$. LLM QA prompt: *Does the text discuss a recent storm in India?*

ing media attention (Wu 2007; Greer 2003). Cultural affinity, geographic closeness, and economic ties have all been shown to influence coverage (Adams 1986; Berlemann and Thomas 2019; Joye 2010). While these findings are robust, most studies examine US media outlets, with limited research on non-American sources.

Terrorism reporting has unique dynamics, as terrorism inherently aims to attract media attention. Cultural factors significantly affect agenda, as seen in US media’s disproportionate focus on Muslim perpetrators (Kearns, Betus, and Lemieux 2019). Comparative studies reveal differences in news selection criteria; for example, Chinese media are more selective in covering terrorist attacks compared to US outlets (Zhang, Shoemaker, and Wang 2013).

Most related work. In this study, we focus on conceptual replication and broadening of Eisensee and Strömberg (2007)’s seminal analysis of US television news attention toward global disasters. This study shows that a disaster in Africa requires significantly higher death tolls to gain equivalent coverage to one in Eastern Europe in US media. We extend this research to a much larger and multilingual corpus, analyze the news coverage in the US and outside of it, and study both coverage of disasters and terrorist attack events to illustrate our method’s flexibility. Finally, we provide the first evaluation of Eisensee and Strömberg (2007)’s keyword-based article matching method, finding significant errors, motivating the development of an alternative.

3 Linking News Articles to Event Records

In this section, we first introduce the fundamental task of identifying news discussing a particular event, using event metadata from databases of event records. Given the large datasets of event records and articles that we aim to link, a methodology that can solve this fundamental task should scale to hundreds of millions of news articles across different languages and thousands of events that naturally happen in our world.

3.1 General Problem Statement

Event information exists in an event database, \mathcal{E} , whose records contain metadata about each event—in particular, what it is, and where and when it happened. We formulate an event’s metadata as a *fingerprint*, $e \in \mathcal{E}$. Given event fingerprints, and a large corpus of time-stamped news articles, \mathcal{D} , the fundamental task of any event coverage analysis is to identify news articles that describe each of the events—that is, for every $e \in \mathcal{E}$, to retrieve the set of documents, $D_e \subset \mathcal{D}$, that discuss an event specified by fingerprint e , for any $e \in \mathcal{E}$. Figure 1 shows an example of the input and desired output for this task.

We next describe the news document dataset (\mathcal{D}) and event databases (\mathcal{E}) that we focus on in this study.

3.2 Large Dataset of News Articles

As a large corpus of news articles, \mathcal{D} , we use data from Media Cloud, a platform that collects articles from news feeds of 1.1 million distinct media outlets since 2008 (Roberts et al. 2021). The platform maintains a list of news outlets organized by country⁵, which has been carefully curated over years by international media scholars to cover important news outlets (Roberts et al. 2021). In this study, we exclude any URLs from popular social media platforms (such as twitter.com, facebook.com, reddit.com). We collected metadata and full text of all, ~ 42 million, news articles published from January 1, 2020 to June 30, 2020 in three languages (Chen et al. 2022, 2024): English (31M articles), Spanish (8.2M), and French (3.2M). We chose these languages because they are among the top 6 languages with most speakers worldwide (Eberhard, Simons, and Fennig 2021) and our team includes members proficient in these languages, which facilitated data exploration.

⁵<https://sources.mediacloud.org/#/collections/country-and-state>

| Lang. | Total | Passing Phase 1 | | | Passing Phase 2 | | |
|-------|-------|-----------------|-------------|------|-----------------|-------------|------|
| | | Total | over events | | Total | over events | |
| | | | Median | Max | | Median | Max |
| EN | 31 M | 82268 | 89 | 6082 | 18658 | 13 | 2405 |
| ES | 8.2 M | 61460 | 63 | 2771 | 7506 | 1.5 | 1000 |
| FR | 3.2 M | 8571 | 8 | 427 | 1277 | 0 | 167 |

Table 1: For each language’s subcorpus: the total number of initial articles, and how many pass each phase of FAME’s pipeline; and the median and maximum number of articles matched per event.

3.3 Databases of Event Records

We aim to identify news articles about critical events based on event records from databases, \mathcal{E} , of disasters (EM-DAT, USGS) and terrorist attacks (GTD). EM-DAT⁶ is a comprehensive database storing details about natural disasters, compiled from UN agencies, research institutes, and press agencies, among others. We focus on natural disasters that include hydrological, geophysical, and meteorological disasters. The United States Geological Survey⁷ (USGS) is a database storing details about earthquakes of magnitude 4.5 or higher throughout the world as part of the National Earthquake Hazards Reduction Program (NEHRP). The Global Terrorism Database⁸ (GTD) is a comprehensive and open-source database that contains information about a variety of terrorist incident events around the world from 1970 through 2020, where information is based on a variety of open media sources. Since GTD includes many terrorist incidents, we focus only on the most salient events that result in more than 10 casualties or constitute one of the top 3 terrorist attacks in terms of casualties for each unique country.

After this selection from the event databases, we accumulate 188 natural disaster and 282 attack events that started and finished in the first six months of 2020 for our news coverage study.

4 Method Identifying News on Critical Events

We introduce and use a method that efficiently links multilingual news articles in \mathcal{D} to records that are listed in an event database, \mathcal{E} .

4.1 Input and Fingerprint Specification

FAME takes as input an event fingerprint, $e \in \mathcal{E}$, which is a narrow subset of event record metadata, and a dataset of news articles \mathcal{D} . The fingerprint is a triplet $e = \langle c, l, t \rangle$ (e.g., $\langle \text{STORM}, \text{INDIA}, 2020/05/20 \rangle$), where: $c \in \mathcal{C}$ encodes an event class (e.g., STORM), $l \in \mathcal{L}$ is a location of an event (e.g., CHILE), and t is the starting date of an event (e.g., 2020/05/17). Here, \mathcal{C} is a set of discrete event classes (e.g., ATTACK, STORM) and \mathcal{L} is a set of discrete location classes (e.g., INDIA, ITALY). Such basic metadata is widely available across event record databases. More detailed metadata

⁶<https://www.emdat.be/>

⁷<https://earthquake.usgs.gov/earthquakes/map/>

⁸<http://apps.start.umd.edu/gtd/>

or types of event arguments, such as an ATTACK’s perpetrator group, is not in the fingerprint because they may be too class-specific, and can raise broader semantic issues; we leave incorporation of these for future work.

These are the same inputs as the ones used in the study of Eisensee and Strömberg (2007), which we seek to replicate and expand. This input specification assumes that there are no duplicate event fingerprints, i.e., two events having the same fingerprint. This assumption holds true for all natural disasters and for 98% of the terrorist attacks for the databases we work with.

4.2 Method

Our method for extracting multilingual articles about events follows the following steps. (1) *Heuristic filtering* selects a set of candidate articles that might discuss the event of interest. We apply a rule-based approach to select these candidate articles and aim to preserve recall in this step. (2) *LLM filtering* selects a final subset of articles. Here, we apply LLM QA. We illustrate an example application of heuristic filtering and LLM filtering in Figure 2.

Phase One: Heuristic filtering For each event fingerprint, the heuristic filtering step aims to prune the original set of articles to a small subset which includes all articles that discuss the event specified by the fingerprint, as in Figure 2. Given fingerprint $\langle c, l, t \rangle$, the filtered documents must satisfy three rules:

- (1) Be within seven days of the start date t of the event,
- (2) Contain a keyword indicating the location of the event ($k_l \in K_l$ where K_l is a set of keywords that may indicate the occurrence of an event in location $l \in \mathcal{L}$ (e.g., Kolkata, Odisha... for INDIA)),
- (3) Contain a keyword indicating the class of the event ($k_c \in K_c$ where K_c is a set of keywords that may indicate the occurrence of an event of class $c \in \mathcal{C}$ (e.g., blizzard, typhoon... for STORM)).

Figure 2 shows examples of document text that passes Phase One heuristic filtering for the fingerprint, $\langle \text{STORM}, \text{INDIA}, 2020/05/20 \rangle$.

To increase the chance that all articles discussing an event described by fingerprint $\langle c, l, t \rangle$ pass this Phase One step, K_c and K_l must be comprehensive keyword sets, containing as many keywords as possible that could indicate the occurrence of the event. To construct K_c and K_l , we used glossaries and lexical databases such as WordNet (Miller 1995), which stores synonyms and hyponyms—words that store

more specific semantic fields. Then, following the approach of Cai and O’Connor (2024), we expanded the sets using an LLM by setting its temperature hyperparameter—which controls the extent of output randomness—to 0.5 and repeatedly prompting for candidate synonyms and hyponyms of c and l . We ultimately added any words that appeared in at least 50% of the LLM outputs to the keyword sets.

Phase Two: LLM filtering Since Phase One extracts documents that may potentially discuss an event specified by a fingerprint, aiming for perfect recall but not high precision, a second filtering step uses LLM QA to link event fingerprints with articles, with high precision and recall. For each candidate article that passes initial filtering, FAME sends to GPT-3.5-turbo the following *LLM prompt*:

```
[article title + first 3 sentences]

Does the text discuss a recent
[event class  $c$ ] in [location  $l$ ]?
```

Answers from the LLM that include a ‘Yes’ response indicate that an article discusses an event; otherwise, the article is dropped. We finalized the question posed to the LLM after evaluating four versions of questions that varied based on specificity (Appendix B). The LLM prompt contains only the article title and first 3 sentences of article text because we found that out of a sample of 150 articles across various event classes, if an article positively mentioned an event, it mentioned the event in either the title or first 3 sentences.

We use GPT-3.5-turbo for the LLM QA step, even though other large language models, including later versions of this one, are available because it performs well and is cheap. In our experiments, open source models did not perform as well (results are in Appendix B), but we hope to explore this possibility in our future research.

4.3 Handling of Multiple Languages

FAME can extract articles that discuss an event specified by a fingerprint for any language. For the Phase One heuristic filtering step, the keyword sets K_c and K_l consist of words that belong to the same language as those in the text documents. We populate the keyword sets in all languages (English, Spanish, and French) using the same procedure (§5.1).

While the Phase One heuristic filtering step requires keyword engineering for each language, the Phase Two LLM QA step applies cross-lingual QA, with no customization for each language. Specifically, the LLM prompt uses the article title and first 3 sentences of article text, word-for-word, and poses a question in English irrespective of the language of the article text. This cross-lingual approach does not face prompt and wording sensitivity issues for questions in languages that are not English, and follows a vast literature supporting cross-lingual QA (Liu et al. 2019; Lewis et al. 2020; Muller et al. 2022; Zhou et al. 2021). We additionally find that it performs well for linking events with Spanish- and French-language articles in §5.4.

5 Identifying News about Disasters and Terrorist Attacks

Here, we first describe how FAME is adjusted to the event databases by choosing relevant keywords sets K_l and K_c . Then, we apply FAME to identify news about natural disasters and terrorist attacks. We label 6,468 event-news pairs passing Phase One as either correct or incorrect matches. Using these labeled samples as a ground truth, we evaluate FAME’s performance across eight event classes.

5.1 Adjusting FAME to the Event Databases

For our evaluation and social science study, the event fingerprints $\langle c, l, t \rangle$ are from EM-DAT, USGS, and GTD databases, introduced in §3. The set of discrete event classes \mathcal{C} includes natural disaster and attack classes: EARTHQUAKE, FLOOD, AVALANCHE, STORM, LANDSLIDE, VOLCANO, WILDFIRE, and ATTACK, and the set \mathcal{L} of discrete location classes corresponds to countries, all derived from EM-DAT, USGS, and GTD.

Keyword sets K_l and K_c consist of location and event class keywords respectively, and aim to be comprehensive. Therefore, K_l , which is the keyword set corresponding to country $l \in \mathcal{L}$, consists of the country name and demonyms from the Open Event Data Alliance’s `CountryInfo.txt`,⁹ as used in (O’Connor, Stewart, and Smith 2013) province names as indicated by ‘admin1’ in the Geonames database¹⁰, and any city in the top 5000 most populated cities in the world as per the GeoNames database. These databases are all in English; we used Google Translate¹¹ to convert database entries to Spanish and French languages, and verified a sample of 200 of the translations with fluent speakers. K_c , which is the keyword set corresponding to event class c , consists of synonyms or hyponyms which are from thesauruses¹², are sampled from GPT-3.5-turbo, or are derived from the WordNet lexical database. These are stemmed and appended to different affixes (e.g., -s, -ing, etc.). Since WordNet is an English lexical database, we used Google Translate and verified the Spanish and French keyword sets with fluent speakers.

5.2 How Many Articles Pass FAME Filters?

Next, we apply FAME to the three databases. In Table 1, we present the number of articles passing each Phase of FAME, including their statistics per event.

Phase One drastically filters down articles by more than two orders of magnitude, highlighting its importance for computational efficiency—the initial corpus is far too large to run through an LLM. Then in Phase Two, LLM QA again removes a large number of articles (more than three quarters), motivating the next section’s manual evaluation to assess if it is removing false positives as intended. Consistent with

⁹<https://github.com/openeventdata/CountryInfo>

¹⁰<https://www.geonames.org/>

¹¹<https://translate.google.com/>

¹²<https://www.merriam-webster.com/thesaurus/>,
<https://dictionary.cambridge.org/dictionary/english-french/>,
<https://dictionary.cambridge.org/dictionary/english-spanish/>

Zuckerman et al. (2019), we also find that most events have very low news coverage—in fact, for all phases and each language’s subcorpus, there are events with zero linked articles (i.e., our method finds no news coverage), while a small number of events have significantly higher news coverage.

5.3 Ground Truth Event-news Pairs

We label pairs of events and news identified by FAME in English, Spanish and French over all 8 event classes. If more than 3 event fingerprints in \mathcal{E} have the same event class, we randomly selected pairs corresponding to 3 events from that class for labeling (4 for the ATTACK event class, of which there are significantly more events), where each event fingerprint corresponds to at least 30 articles passing Phase One, and at least 5 articles passing Phase Two. If fewer than 3 event fingerprints in \mathcal{E} have the same event class, we labeled pairs of events and news corresponding to all events of that class. We compute precision, recall and F1 scores for each language and each event fingerprint that we labeled news for.

We construct ground truth labels of articles corresponding to event fingerprint records $\langle c, l, t \rangle$ using annotators hired from Upwork who are fluent in the target languages.¹³ For each language, two annotators labeled all or the first 150 Media Cloud articles that passed the Phase One heuristic filtering step, which aimed to have perfect recall, for each event record. An article received a positive label if any part of its title or first three sentences of text discusses the input event record and a negative label otherwise, indicating no or peripheral discussion of the event. Annotators for English-, Spanish-, and French-language text agreed 94.5%, 95.6%, and 95.8% of the time respectively (Cohen’s kappa of 0.88, 0.91, and 0.91 respectively). Most disagreements (approximately 90%) stemmed from inadvertent errors, such as overlooking relevant text, inconsistent application of the annotation criteria, and skipping annotation of certain articles entirely. Disagreements were resolved by the authors through a systematic review of the annotation guidelines. To ensure consistency, all resolved labels were carefully aligned with the predefined annotation criteria.

Given that the ground truth data consists of articles that pass Phase One, we seek as complete a set of articles as possible, using a very large set of keywords containing synonyms and hyponyms. We test whether the keyword sets are expansive enough by performing an ablation that considers even more distant synonyms in the keyword sets (details in Appendix B). Because we find that FAME links the same articles, we believe our original keyword sets are as large, relevant and varied as they need to be, though we cannot guarantee perfect recall in retrieval of relevant articles from Media Cloud.

5.4 Performance Comparison

The averaged results of our evaluation in Table 2 highlight the high performance of our method for linking English-,

¹³Each annotator has between 98 and 100 percent satisfaction over their previous jobs. Compensation was roughly \$25 per hour. The UMass Human Research Protection Office determined this annotation study is not human subjects research.

| Lang. | Method | Prec. | Recall | F1 |
|-------|----------------|-------|--------|------|
| EN | KW: Title-only | 88.1 | 48.4 | 57.6 |
| EN | KW: Title+Body | 84.1 | 78.0 | 79.3 |
| EN | FAME | 93.2 | 95.4 | 94.0 |
| FR | FAME | 96.9 | 94.1 | 95.2 |
| ES | FAME | 98.1 | 96.0 | 96.8 |

Table 2: Averaged precision, recall, and F1 scores over expert-annotated articles matched to events for FAME and two previously used keyword classifier baselines (“KW”) adapted to our task on English data (title-only, and on first three sentences of article text). Results for English (FAME and the baselines), Spanish (FAME), and French (FAME).

Spanish-, and French-language articles to events.¹⁴ Since many social science studies use keyword matching to identify event coverage (L. Hellmueller and Lindner 2022), we adapt the method in Eisensee and Strömberg (2007), which has the most similar task definition as FAME, as a baseline, using the same location and event class keywords that they use. Since Eisensee and Strömberg (2007) does not analyze or have keywords for ATTACK events, we construct the keyword set based on the attack hierarchy and definitions in the GTD codebook. We do not use existing NLP methods as baselines since they either rely on large amounts of training data to perform well or face other inconsistencies with our task as discussed in §2.

Specifically, for English-language articles, our baselines consist of keyword matching on Media Cloud article titles, and on article titles along with the first three sentences of article text. In both cases, the approach produces significant errors and performs more poorly than FAME.

While aggregate results as in Table 2 are typically used in NLP, they may mask significant variation across events, which are ultimately our units of analysis. Therefore, we show models’ precision and recall performance broken out over each event in Figure 3 (the results for Spanish in French are in Appendix A). We note that while the averaged KW:Title+Body result, of 79.3 F1, may sound reasonably accurate, this masks massive variation—some events see precision low as 21.4 (and recall, 26.9)! This could have negative consequences when trying to make qualitative or social scientific interpretations from LLM/NLP output, highlighting the need for more NLP evaluation—and perhaps new modeling approaches—that better target the finer levels of substantively relevant granularity that are important for computational social science.

5.5 Ranking of Events with Most News Coverage

To qualitatively evaluate the results of FAME, we examine the events that experienced the greatest news coverage in the first half of 2020. Table 3 presents the 10 disasters that garnered the highest number of news reports as retrieved by FAME. The majority of these events are documented

¹⁴Averaged events must have at least one article discussing it to properly calculate precision and recall.

| Event | Class (<i>c</i>) | Location (<i>l</i>) | Time (<i>t</i>) | News | Death |
|-----------------------|--------------------|-----------------------|-------------------|------|-------|
| Cyclone Amphan | Storm | India | May 20 | 2723 | 90 |
| Taal Volcano eruption | Volcano | Philippines | Jan 12 | 2617 | 1 |
| Magnitude 6.4 quake | Earthq. | Puerto Rico | Jan 7 | 1690 | 2 |
| Cyclone Amphan | Storm | Bangladesh | May 20 | 1437 | 26 |
| Storm Ciara | Storm | UK | Feb 8 | 1314 | - |
| Severe thunderstorms | Storm | USA | Mar 2 | 1268 | - |
| 2020 Elazığ quake | Earthq. | Turkey | Jan 24 | 1100 | 41 |
| Storm Gloria | Storm | Spain | Jan 19 | 1046 | 4 |
| Severe Storm | Storm | USA | Apr 10 | 857 | - |
| Storm Amanda | Storm | El Salvador | May 31 | 835 | - |

Table 3: Top 10 disasters with the most news in the first five months of 2020. Event names were derived from the headlines of news or titles of respective Wikipedia pages.

on Wikipedia. Notably, our method successfully identifies events associated with the same disaster occurring in different locations, such as Cyclone Amphan which affected India and Bangladesh. For space considerations, the top 10 terrorist attacks are provided in Appendix C.

6 Investigating Bias in News Coverage

Event identification allows us to examine how news coverage differs across countries. Prior work found that a disaster in Africa requires a death toll that is 40 times higher than one in Eastern Europe to receive equivalent news coverage in US media (Eisensee and Strömberg 2007). However, it is unclear whether such differences in media coverage are due to geographic distance between the countries where the event occurred and where it was reported, or due to other factors, such as cultural, societal, or economic proximity.

Using articles extracted from Media Cloud from between January and May of 2020 using FAME, we perform a large-scale investigation of patterns in news coverage of natural disasters and terrorist attacks. First, we compute the average number of news stories per event in a country reported in another country.¹⁵ Then, to identify patterns in media coverage of such events, we regress this average against a set of 47 country characteristics proposed as determining factors by media scholars. We were able to gather these characteristics for 128 countries, which form the base for our regression analysis, since samples in our regression models correspond to country pairs.

We describe details of our regression models in the next subsection (§6.1). Then, we analyze the US media coverage of critical events (§6.2). Finally, we expand the scope to investigate global media coverage of disasters and terrorist attacks (§6.2), focusing on reports from countries where either English, Spanish, or French is an official language.

6.1 Details of Regression Models

Sample selection Overall, our set of events happen in 136 unique countries. If a country did not report on an event, then the dependent variable for that country pair is zero. To limit the number of zeroes, we restrict reporting countries either to the US (§6.2), or to the countries that cover critical events

happening in at least ten countries (§6.3). There are 10 such reporting countries in English, 4 in Spanish, and 2 in French (full list in Appendix C).

Choice of candidate factors To identify factors that might influence news coverage of disasters, we conducted a literature review (§2). We classified the factors into five categories: event characteristics, economic, societal, geographic and cultural (Lee 2007; Segev 2016). We consider key characteristics of the specific event, such as the number of deaths, and the characteristics of the country where that event happened: its continent, GDP, Gini index of income inequality, and Democracy index (The Economist Intelligence Unit 2024). Country indices are sourced from the World Bank and The Organization for Economic Co-operation and Development (OECD). Most of the factors we consider are relational, that is they describe the relationship between the country where the critical event happened and the country where it might be covered, e.g., whether these countries are neighbors, share an official language, their levels of trade and direct investment in each other. All 47 considered factors are listed in Appendix D with their detailed explanations, definitions, and value ranges.

Feature and model selection Next, we evaluate which of the 47 candidate factors may be related to the news coverage of disaster and terrorist attack events. Given the possibility of collinearity and irrelevance among these factors, we employed a feature selection algorithm to discern the most impactful predictors. Specifically, we identify the features that are the most significant through forward selection to greedily optimize the Akaike Information Criterion (AIC), adding factors one by one. The AIC method assists in identifying the optimal set of predictors that most effectively explain the variations in the dependent variables, ensuring a robust and effective model. We also experimented with Variance Inflation Factor (VIF) feature selection, which yielded similar results.

Factor preprocessing To ensure comparability across factors, each factor was normalized using a min/max scaler, which rescales every factor to a numeric value between 0 and 1. Additionally, logarithmic scaling was applied to the death count for smoothing the skewness of the data. Death count information is missing for some events: in such cases

¹⁵We consider only the events with unique fingerprints.

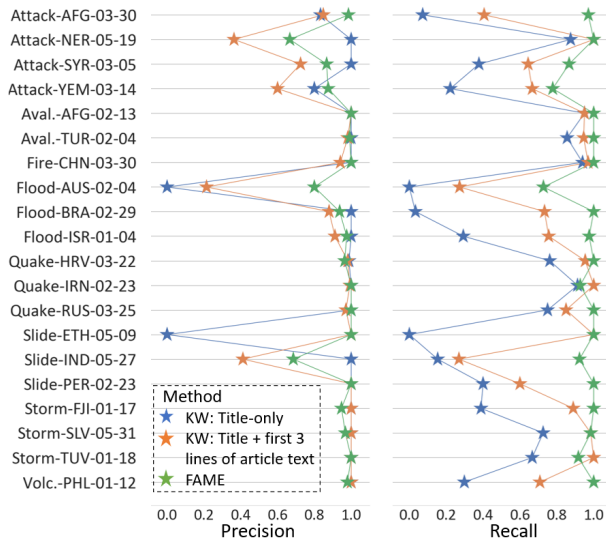


Figure 3: Precision and recall of FAME and the two base-lines (“KW”) for random events of each event class.

we imputed it with the average death count per event, following our manual examination of a subset of 20 disasters without death counts.¹⁶

6.2 US News Coverage

We regress the average number of news articles published in the US about critical events happening in another country against country-level factors. Table 4 lists the coefficients of the factors identified by feature selection as predictive of the outcome variable.

First, the resulting models for disasters and terrorist attacks achieve high R^2 values, 0.603 and 0.440, meaning that these models explain 60% and 44% of the variance in the average number of news coverage per disaster, respectively. We find that the factor that impacts the news coverage of disasters the most is the number of deaths (left column coefficients in Table 4). This factor is also the most predictive for the news coverage of terrorist attacks (right column coefficients in Table 4). This common-sense relationship between mortality and coverage of disasters was described in prior works (Eisensee and Strömberg 2007) and, thus, suggests FAME is retrieving relevant data.

Second, Eisensee and Strömberg (ES) find that geography matters significantly in the coverage of disasters in US news in the late 20th century. We do not observe a significant bias in news coverage of disasters in Europe or Africa (left column in Table 4), but we see that disasters in high GDP countries received significantly more attention. This is consistent with ES’s provocative finding that disasters in Eastern Europe receive proportionally more coverage than disasters in Africa, since national GDP and GDP per capita are generally higher in Eastern European countries than in African

¹⁶For about a half of them we inferred non-zero death counts based on information available in the Internet, but we could find the exact death count only for three of them.

| Type | Variable | Disaster | Attack |
|------|---------------------|-----------------|-----------------|
| Eve. | #Deaths | 3.401*** | 3.221*** |
| Geo. | Neighbor | -2.234 | -3.479 |
| | Asia | 0.382 | 0.728 |
| | Europe | 0.472 | 0.415 |
| | Africa | 0.335 | 0.204 |
| | Oceania | 1.615*** | 0.310 |
| | North America | 0.420* | 0.099 |
| | South America | 0.207 | 0.158 |
| Eco. | GDP | 0.315** | 0.379** |
| | Trade | 2.050 | 4.427 |
| | Investment | -0.254 | 0.350 |
| Pol. | Democracy in. | 0.060 | -0.011 |
| | Diplomatic relation | 0.080 | -0.108 |
| Soc. | Gini in. | -0.070 | 0.032 |
| Cul. | Same language | 0.042 | -0.024 |
| | const | -0.442 | -0.095 |
| | #Observations | 128 | 128 |
| | Adjust R^2 | 0.603 | 0.440 |

Table 4: Regression results of average US news coverage per disaster or terrorist attack happening in a foreign country. The top 3 significant coefficients with the largest absolute value are marked in bold font. We indicate statistical significance at levels $p < 0.001$ (***), $p < 0.01$ (**), and $p < 0.05$ (*).

countries. In our data, a disaster in Nigeria requires 10 times more deaths than in Russia to receive the same level of attention, while a terrorist attack requires 21 times more deaths. While this similarity to ES is encouraging, there’s a danger in comparing news from 2020 in our set to theirs, from 1968 to 2002. For most of ES’s data, the Cold War meant that Eastern Europe was of particular interest to US audiences. In our data, we note a high degree of attention to disasters in North America and Oceania: wealthy nations (the US, Canada, Australia, New Zealand) are well-represented on these continents.

Third, terrorist attacks achieve much more coverage if there is a large volume of trade between the reporting country and the country affected. This factor was by far the most predictive of news coverage of terrorist attack events in the US, even more so than the number of deaths. This finding echoes findings in existing studies of news coverage (Segev 2016; Wu 2000; Chen et al. 2024).

6.3 News Coverage across Countries and Languages

We also performed a regression analysis at a global scale, examining the average number of news articles about an event in a country against various country-level factors. Table 5 shows the coefficients of the factors identified as predictive for news coverage of disaster events, categorized by language. Similar to the analysis of US news coverage, the global analysis finds that the number of deaths, GDP, and event location in Oceania are among the most significant factors influencing global media attention. A higher death toll corresponds to greater news coverage, and Oceania’s high

| Type | Variable | English | Spanish | French |
|------|---------------------|-----------------|-----------------|-----------------|
| Eve. | #Deaths | 2.961*** | 3.485*** | 2.343*** |
| Geo. | Continent sim | 0.004 | 0.128 | 0.106* |
| | Neighbor | 0.271** | 0.171 | -0.084 |
| | Asia | 0.157* | 0.222 | 0.048 |
| | Europe | 0.213** | 0.302** | -0.043 |
| | Africa | 0.074 | 0.216 | -0.017 |
| | Oceania | 0.830*** | 0.574*** | 0.303** |
| | North America | 0.141 | 0.282* | -0.041 |
| | South America | 0.073 | -0.040 | 0.001 |
| Eco. | GDP (high) | | | -0.012 |
| | GDP (high-high) | 0.351*** | 0.556*** | -0.012 |
| | GDP (high-low) | 0.057 | -0.070 | |
| | Investment | -0.300 | -0.622 | -0.023 |
| | Trade | 0.499 | -0.643 | -0.038 |
| Pol. | Democracy in. | 0.044 | 0.051 | 0.011 |
| | Democracy in. (h-l) | 0.060 | | |
| | Diplomatic relation | 0.031 | -0.043 | -0.009 |
| Soc. | Gini in. | -0.059 | -0.076 | -0.019 |
| | Gini in. (h-h) | 0.004 | | |
| Cul. | Same language | 0.025 | 0.344*** | -0.006 |
| | const | -0.133 | -0.197 | 0.017 |
| | #Observations | 1281 | 511 | 256 |
| | Adjust R^2 | 0.623 | 0.639 | 0.740 |

Table 5: Regression results of average news coverage between countries per disaster in different languages. Missing regression coefficient indicates that the respective factor is irrelevant according to feature selection.

coefficient across languages likely reflects the prominence of high-GDP countries such as Australia and New Zealand, which play significant roles in the geopolitical landscape. This pattern aligns with studies on systemic biases in global media coverage, suggesting that wealthier nations receive more attention due to their economic influence, media infrastructure, and geopolitical importance (Wu 2000; Segev 2016; Chen et al. 2024). Unlike the US-based study by Eisensee, which identifies a bias against Africa in disaster coverage, the global analysis indicates more balanced attention toward disasters in Africa, Europe, and Asia.

Table 6 presents the coefficients for terrorist attack events, demonstrating that the number of deaths remains the most predictive factor for news coverage across languages. Economic factors, however, exhibit a nuanced role in the context of terrorist attacks. While GDP remains significant, trade volume emerges as an equally predictive factor, underscoring the economic interdependence between nations and its influence on media agendas. Diplomatic relations also appear as a significant factor for terrorist attack events but not for disasters, suggesting that geopolitical alliances and strategic partnerships amplify media attention to certain attacks. South America exhibits a notable coefficient exclusively in Spanish-language media, while Africa appears prominently in French-language media. This pattern may reflect the heightened attention of the Spanish-speaking and French-speaking worlds to the respective continents’ political instability and humanitarian crises.

| Type | Variable | English | Spanish | French |
|------|---------------------|-----------------|-----------------|-----------------|
| Eve. | #Deaths | 2.695*** | 2.646*** | 2.423*** |
| Geo. | Continent sim | -0.015 | -0.099 | -0.016 |
| | Neighbor | 0.298** | 0.024 | -0.056 |
| | Asia | 0.311*** | 0.345* | 0.280 |
| | Europe | 0.251** | 0.390** | 0.261 |
| | Africa | 0.076 | 0.261 | 0.407* |
| | Oceania | 0.171 | 0.618** | 0.656** |
| | North America | 0.192 | 0.330 | 0.287 |
| | South America | 0.106 | 1.106*** | 0.326 |
| Eco. | GDP (high) | 0.324*** | 1.012*** | 0.587*** |
| | GDP (high-high) | 0.561*** | 0.490*** | 0.567*** |
| | GDP (high-low) | | -0.489*** | |
| | Investment | 0.127 | 0.916* | 0.560 |
| | Trade | 1.967*** | 3.858*** | 1.798** |
| Pol. | Democracy in. (h-h) | -0.073 | -0.025 | -0.000 |
| | Democracy in. (h-l) | 0.155* | -0.489 | |
| | Diplomatic relation | 0.217*** | 0.318*** | 0.049 |
| Soc. | Gini in. (h-h) | 0.031 | 0.087 | -0.049 |
| | Gini in. (h-l) | 0.002 | | -0.049 |
| Cul. | Same language | 0.108** | -0.205 | -0.247*** |
| | const | -0.206* | -0.400** | -0.197 |
| | #Observations | 1281 | 511 | 256 |
| | Adjust R^2 | 0.484 | 0.518 | 0.499 |

Table 6: Regression results of average news coverage between countries per terrorist attack in different languages.

7 Limitations

In this study, we proposed minimalist fingerprints, which did not result in any fingerprint collisions for disasters, but resulted in 2% of terrorist attacks having the same fingerprint. For other datasets, future studies may want to use longer fingerprints or account for the potential fingerprint collisions.

Our findings are inherently limited by the datasets we focus on. Terrorism is a somewhat subjective label, and an attack that is not considered a terrorist attack by the GTD will not register in our set. Similar caveats apply regarding EM-DAT: errors in the database will propagate to our analysis. Media Cloud, while far reaching, does not have universal coverage.

One serious limitation is the short duration of our study. Unlike Eisensee and Strömberg, who examine 34 years of data, we consider only six months. Those six months in question were quite unusual: COVID-19 was spreading rapidly, and our study period includes the global “shutdown” associated with the pandemic. It is possible that results during our study period might be significantly different than other time periods due to rapid changes in global newsgathering and the global focus on the pandemic. We are reassured though that our method seems to emphasize similarities in media attention to studies conducted during other time periods.

8 Conclusion

In this study, we present a novel methodology for systematically identifying and investigating global news coverage of critical events, specifically disasters and terrorist attacks. Our method, FAME (FINGERPRINT TO ARTICLE MATCH-

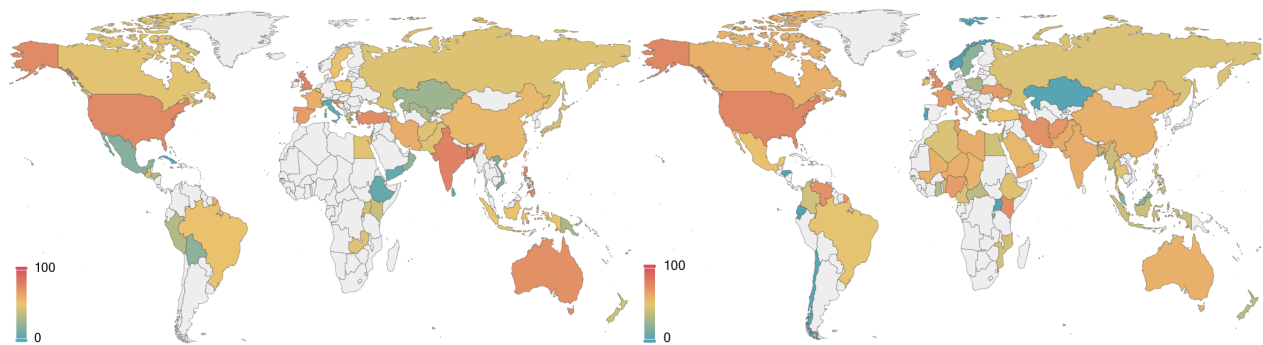


Figure 4: The average number of news in English per disaster (left) and per terrorist attack (right) happening in a given country.

ING FOR EVENTS), provides a scalable, efficient, and robust approach to match event fingerprints—time, location, and class—to news articles without relying on training data. By leveraging massive databases, including event databases and multilingual news articles from Media Cloud, FAME achieved excellent performance, with average F1 of 94.0 on English articles, and of 95.3 across all three languages. It linked 470 events to over 27,441 news articles (out of 42.4 million) across 3 languages and 8 event classes. The introduced novel dataset and task of linking events from databases to news articles, will stimulate the development of methods that will enable unprecedented agenda setting studies. Future works can extend our approach, e.g., by using more complex fingerprints, or supervised learning.

Through our large-scale analysis, we identified patterns in global news coverage that align with findings from prior research and add nuance to them. Specifically, we observed that events with higher death tolls, occurring in countries with higher GDP, and those with greater trade volumes are more likely to receive extensive global media attention (noticeable in Figure 4). Furthermore, this methodology enables us to investigate such patterns at scale, overcoming the limitations of prior approaches.

This work provides a robust foundation for further exploration of patterns in global news coverage. By offering a scalable and automated solution, FAME opens avenues for examining critical events across centuries and their coverage in diverse languages, regions, and contexts. Future research can build on this approach to deepen our understanding of how global media frames and prioritizes events.

Acknowledgments

We thank Paul Musgrave for extensive feedback and preliminary collaboration, and additionally the UMass NLP group and the anonymous reviewers for helpful comments and feedback. This work was supported by the Provost’s Interdisciplinary Research Grant 2023 (University of Massachusetts Amherst) titled “Analyzing cross-country bias in news coverage of international conflicts and disasters”, as well as the National Science Foundation award 1845576. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science

Foundation. For annotations, we thank Ana Lucia Hatsumi Yamashiro Miyahira, Luciano Redondo, Romaric Zoungan, Olivier Gaurat, and Houndemikon Hintekpo David.

References

- Adams, W. C. 1986. Whose lives count? TV coverage of natural disasters. *Journal of communication*, 36(2): 113–122.
- Ahmad, W.; Peng, N.; and Chang, K.-W. 2021. GATE: Graph Attention Transformer Encoder for Cross-lingual Relation and Event Extraction. In *AAAI*.
- Allaway, E.; Wang, S.; and Ballesteros, M. 2021. Sequential Cross-Document Coreference Resolution. In *EMNLP*, 4659–4671.
- Barhom, S.; Shwartz, V.; Eirew, A.; Bugert, M.; Reimers, N.; and Dagan, I. 2019. Revisiting Joint Modeling of Cross-document Entity and Event Coreference Resolution. In *ACL*, 4179–4189.
- Baum, M. A.; and Zhukov, Y. M. 2015. Filtering revolution: Reporting bias in international newspaper coverage of the Libyan civil war. *Journal of Peace Research*, 52(3): 384–400.
- Berlemann, M.; and Thomas, T. 2019. The distance bias in natural disaster reporting - empirical evidence for the United States. *Applied Economics Letters*, 26(12): 1026 – 1032.
- Caciularu, A.; Cohan, A.; Beltagy, I.; Peters, M.; Cattan, A.; and Dagan, I. 2021. CDLM: Cross-Document Language Modeling. In *Findings of EMNLP*, 2648–2662.
- Cai, E.; and O’Connor, B. 2024. A Monte Carlo Language Model Pipeline for Zero-Shot Sociopolitical Event Extraction. arXiv:2305.15051.
- Cattan, A.; Eirew, A.; Stanovsky, G.; Joshi, M.; and Dagan, I. 2021. Cross-document Coreference Resolution over Predicted Mentions. In *Findings of ACL*, 5100–5107.
- Chen, X.; Hale, S. A.; Jurgens, D.; Samory, M.; Zuckerman, E.; and Grabowicz, P. A. 2024. Global News Synchrony and Diversity During the Start of the COVID-19 Pandemic. In *WWW*, 2639–2650.
- Chen, X.; Xu, S.; Li, P.; and Zhu, Q. 2023. Cross-Document Event Coreference Resolution on Discourse Structure. In *EMNLP*, 4833–4843.

- Chen, X.; Zeynali, A.; Camargo, C.; Flöck, F.; Gaffney, D.; Grabowicz, P.; Hale, S.; Jurgens, D.; and Samory, M. 2022. SemEval-2022 Task 8: Multilingual news article similarity. In *SemEval-2022*, 1094–1106. Seattle, United States: Association for Computational Linguistics.
- Chen, Y.; Chen, T.; Ebner, S.; White, A. S.; and Van Durme, B. 2020. Reading the Manual: Event Extraction as Definition Comprehension. In *Proceedings of the Fourth Workshop on Structured Prediction for NLP*, 74–83.
- Choubey, P. K.; and Huang, R. 2017. Event Coreference Resolution by Iteratively Unfolding Inter-dependencies among Events. In *EMNLP*, 2124–2133.
- Cybulska, A.; and Vossen, P. 2014. Using a sledgehammer to crack a nut? Lexical diversity and event coreference resolution. In *LREC*, 4545–4552.
- Dietrich, N.; and Eck, K. 2020. Known unknowns: media bias in the reporting of political violence. *International Interactions*, 46(6): 1043–1060.
- Du, X.; and Cardie, C. 2020. Event Extraction by Answering (Almost) Natural Questions. In *EMNLP*, 671–683.
- Eberhard, D. M.; Simons, G. F.; and Fennig, C. D. 2021. What are the top 200 most spoken languages. *Ethnologue: Languages of the world*.
- Eirew, A.; Caciularu, A.; and Dagan, I. 2022. Cross-document Event Coreference Search: Task, Dataset and Modeling. In *EMNLP*, 900–913.
- Eirew, A.; Cattan, A.; and Dagan, I. 2021. WEC: Deriving a Large-scale Cross-document Event Coreference dataset from Wikipedia. In *NAACL*, 2498–2510.
- Eisensee, T.; and Strömberg, D. 2007. News droughts, news floods, and US disaster relief. *The Quarterly Journal of Economics*, 122(2): 693–728.
- FORCE11. 2020. The FAIR Data principles. <https://force11.org/info/the-fair-data-principles/>. Accessed: 2024-01-15.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Greer, C. 2003. *Sex Crime and the Media: Sex Offending and the Press in a Divided Society*. Willan.
- Grimmer, J.; Roberts, M.; and Stewart, B. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton University Press.
- Hamborg, F.; Donnay, K.; and Gipp, B. 2019. Automated identification of media bias in news articles: an interdisciplinary literature review. *International Journal on Digital Libraries*, 20(4): 391–415.
- Held, W.; Iyer, D.; and Jurafsky, D. 2021. Focus on what matters: Applying Discourse Coherence Theory to Cross Document Coreference. In *EMNLP*, 1406–1417.
- Hsu, B.; and Horwood, G. 2022. Contrastive Representation Learning for Cross-Document Coreference Resolution of Events and Entities. In *NAACL*, 3644–3655.
- Huang, L.; Ji, H.; Cho, K.; Dagan, I.; Riedel, S.; and Voss, C. 2018. Zero-Shot Transfer Learning for Event Extraction. In *ACL*, 2160–2170.
- Joye, S. 2010. News media and the (de) construction of risk: how Flemish newspapers select and cover international disasters. *Catalan Journal of Communication & Cultural Studies*, 2(2): 253–266.
- Kearns, E. M.; Betus, A. E.; and Lemieux, A. F. 2019. Why Do Some Terrorist Attacks Receive More Media Attention Than Others? *Justice Quarterly*, 36(6): 985–1022.
- Krippendorff, K. 2019. *Content Analysis: An Introduction to Its Methodology*.
- Kwak, H.; and An, J. 2014. Understanding News Geography and Major Determinants of Global News Coverage of Disasters. *ArXiv*, abs/1410.3710.
- L. Hellmueller, V. H.; and Lindner, P. 2022. Terrorist Organizations in the News: A Computational Approach to Measure Media Attention Toward Terrorism. *Mass Communication and Society*, 25(1): 134–157.
- Lee, S. 2007. International public relations as a predictor of prominence of US news coverage. *Public Relations Review*, 33(2): 158–165.
- Lewis, P.; Oguz, B.; Rinott, R.; Riedel, S.; and Schwenk, H. 2020. MLQA: Evaluating Cross-lingual Extractive Question Answering. In *ACL*, 7315–7330.
- Li, F.; Peng, W.; Chen, Y.; Wang, Q.; Pan, L.; Lyu, Y.; and Zhu, Y. 2020. Event Extraction as Multi-turn Question Answering. In *Findings of EMNLP*, 829–838.
- Li, S.; Ji, H.; and Han, J. 2021. Document-Level Event Argument Extraction by Conditional Generation. In *NAACL*.
- Lin, Y.; Ji, H.; Huang, F.; and Wu, L. 2020. A Joint Neural Model for Information Extraction with Global Features. In *ACL*, 7999–8009.
- Liu, J.; Chen, Y.; Liu, K.; Bi, W.; and Liu, X. 2020. Event Extraction as Machine Reading Comprehension. In *EMNLP*, 1641–1651.
- Liu, J.; Lin, Y.; Liu, Z.; and Sun, M. 2019. XQA: A Cross-lingual Open-domain Question Answering Dataset. In *ACL*.
- Lu, Y.; Lin, H.; Xu, J.; Han, X.; Tang, J.; Li, A.; Sun, L.; Liao, M.; and Chen, S. 2021. Text2Event: Controllable Sequence-to-Structure Generation for End-to-end Event Extraction. In *ACL*, 2795–2806.
- Lyu, Q.; Zhang, H.; Sulem, E.; and Roth, D. 2021. Zero-shot Event Extraction via Transfer Learning: Challenges and Insights. In *ACL*, 322–332.
- Mac Carron, P.; and Kenna, R. 2013. Viking sagas: Six degrees of Icelandic separation Social networks from the Viking era. *Significance*, 10(6): 12–17.
- McCombs, M.; and Valenzuela, S. 2021. *Setting the Agenda: Mass Media and Public Opinion*. Wiley.
- McCombs, M. E.; and Shaw, D. L. 1972. The agenda-setting function of mass media. *Public opinion quarterly*, 36(2): 176–187.
- Miller, G. A. 1995. WordNet: A Lexical Database for English. *Commun. ACM*, 38(11): 39–41.
- Muller, B.; Soldaini, L.; Koncel-Kedziorski, R.; Lind, E.; and Moschitti, A. 2022. Cross-Lingual Open-Domain Question Answering with Answer Sentence Generation. In *ACL*, 337–353.

O'Connor, B.; Stewart, B.; and Smith, N. A. 2013. Learning to Extract International Relations from Political Context. In *ACL*.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *NeurIPS*, 35: 27730–27744.

Roberts, H.; Bhargava, R.; Valiukas, L.; Jen, D.; Malik, M. M.; Bishop, C. S.; Ndulue, E. B.; Dave, A.; Clark, J.; Etlings, B.; et al. 2021. Media cloud: Massive open source collection of global news on the open web. *ICWSM*.

Segev, E. 2016. The group-sphere model of international news flow: A cross-national comparison of news sites. *International Communication Gazette*, 78(3): 200–222.

Shoemaker, P. J.; and Vos, T. 2009. *Gatekeeping theory*. Routledge.

The Economist Intelligence Unit. 2024. Democracy Index. [Accessed: 24-Dec-2024].

Wadden, D.; Wennberg, U.; Luan, Y.; and Hajishirzi, H. 2019. Entity, Relation, and Event Extraction with Contextualized Span Representations. In *EMNLP*, 5784–5789.

Wu, H. D. 2000. Systemic determinants of international news coverage: A comparison of 38 countries. *Journal of communication*, 50(2): 110–130.

Wu, H. D. 2007. A brave new world for international news? Exploring the determinants of the coverage of foreign news on US websites. *International Communication Gazette*, 69(6): 539–551.

Zhang, D.; Shoemaker, P.; and Wang, X. 2013. Reality and newsworthiness: Press coverage of International Terrorism by China and the United States. *Asian Journal of Communication*, 23: 449–471.

Zhang, H.; Wang, H.; and Roth, D. 2021. Zero-shot Label-Aware Event Trigger and Argument Classification. In *Findings of ACL*, 1331–1340.

Zhang, S.; Ji, T.; Ji, W.; and Wang, X. 2022. Zero-Shot Event Detection Based on Ordered Contrastive Learning and Prompt-Based Prediction. In *NAACL*, 2572–2580.

Zhou, Y.; Geng, X.; Shen, T.; Zhang, W.; and Jiang, D. 2021. Improving Zero-Shot Cross-lingual Transfer for Multilingual Question Answering over Knowledge Graph. In *NAACL*, 5822–5834.

Zuckerman, E. 2003. Global Attention Profiles-A working paper: First steps towards a quantitative approach to the study of media attention. *Berkman Center Research Publication*, (2003-06).

Zuckerman, E.; Matias, J. N.; Bhargava, R.; Bermejo, F.; and Ko, A. 2019. Whose Death Matters? A Quantitative Analysis of Media Attention to Deaths of Black Americans in Police Confrontations, 2013–2016. *International Journal of Communication*, 13(0).

Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes.**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes.**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes.**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **NA.**
 - (e) Did you describe the limitations of your work? **Yes.**
 - (f) Did you discuss any potential negative societal impacts of your work? **No, because we haven't found any potential negative societal impacts.**
 - (g) Did you discuss any potential misuse of your work? **NA.**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes.**
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes.**
2. Additionally, if your study involves hypotheses testing...
- (a) Did you clearly state the assumptions underlying all theoretical results? **NA.**
 - (b) Have you provided justifications for all theoretical results? **NA.**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA.**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **Yes.**
 - (e) Did you address potential biases or limitations in your theoretical framework? **NA.**
 - (f) Have you related your theoretical results to the existing literature in social science? **Yes.**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **Yes.**
3. Additionally, if you are including theoretical proofs...
- (a) Did you state the full set of assumptions of all theoretical results? **NA.**
 - (b) Did you include complete proofs of all theoretical results? **NA.**
4. Additionally, if you ran machine learning experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes.**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes.**

- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **Yes.**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **NA.**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes.**
 - (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? **Yes.**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity...**
- (a) If your work uses existing assets, did you cite the creators? **Yes.**
 - (b) Did you mention the license of the assets? **NA.**
 - (c) Did you include any new assets in the supplemental material or as a URL? **Yes.**
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **NA.**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes.**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? **Yes.**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? **Yes.**

6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**
- (a) Did you include the full text of instructions given to participants and screenshots? **Yes.**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **Yes.**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **Yes.**
 - (d) Did you discuss how data is stored, shared, and de-identified? **Yes.**

A Appendix A: Per-event Evaluation for Spanish- and French-language articles

Figures 5 and 6 present the per-event precision and recall scores of FAME, on linking disaster events of 8 classes with labeled Spanish- and French-language articles, where events must have at least one article discussing them to properly calculate precision and recall. In most cases, FAME can achieve more than 90% score on the metrics.




Figure 5: Per-event precision and recall of FAME, on linking Spanish-language articles to event fingerprints.

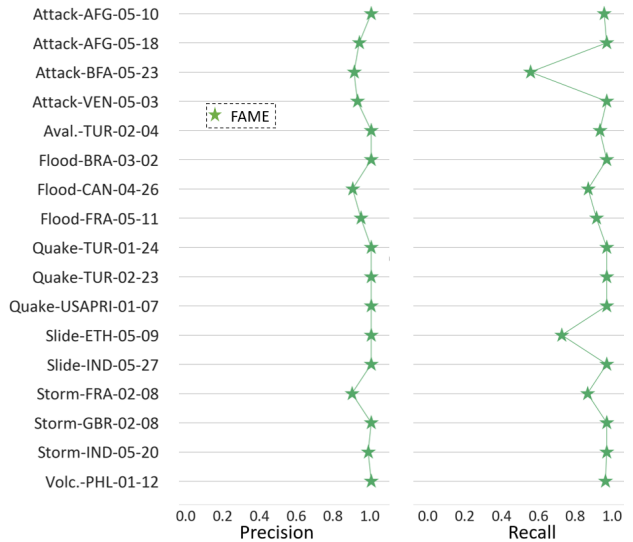


Figure 6: Per-event precision and recall of FAME, on linking French-language articles to event fingerprints.

| Lang. | Method | Prec. | Recall | F1 |
|-------|----------------------|-------|--------|------|
| EN | KW: Title-only | 88.1 | 48.4 | 57.6 |
| EN | KW: Title+Body | 84.1 | 78.0 | 79.3 |
| EN | FAME (llama-3-70b) | 89.0 | 84.3 | 85.5 |
| EN | FAME (llama-3-8b) | 90.9 | 86.1 | 87.0 |
| EN | FAME (GPT-3.5-turbo) | 93.2 | 95.4 | 94.0 |
| FR | FAME (llama-3-70b) | 97.6 | 90.5 | 93.3 |
| FR | FAME (llama-3-8b) | 94.8 | 93.9 | 94.0 |
| FR | FAME (GPT-3.5-turbo) | 96.9 | 94.1 | 95.2 |
| ES | FAME (llama-3-70b) | 98.2 | 91.4 | 94.2 |
| ES | FAME (llama-3-8b) | 96.8 | 93.0 | 94.5 |
| ES | FAME (GPT-3.5-turbo) | 98.1 | 96.0 | 96.8 |

Table 7: Averaged precision, recall, and F1 scores over expert-annotated articles corresponding to events that do not have zero prevalence, for English (FAME and baselines), Spanish (FAME), and French (FAME).

B Appendix B: Ablations

Increasing the size of K_c and K_l . We investigate whether the keyword sets that FAME uses are comprehensive enough. We further increase the size of K_c by including 5 to 10 more synonyms and hyponyms that GPT-3.5-turbo recommends for each event class c , even if they seem to have a very different meaning from c , but observed no difference in results compared to when using the original keyword sets.

Using open-source LLMs. We replace GPT-3.5-turbo in the Phase Two LLM QA step with recent open-source models. In Table 7, we show results of FAME using high-performing open-source models of llama-3-70b and llama-3-8b, which perform quite similarly to, but slightly worse, than FAME using GPT-3.5-turbo. The difference in number of model parameters does not seem to impact performance much.

Question specificity. We also investigate whether wording of the LLM prompt affects performance. First, we find that using a single prompt for the Phase Two LLM QA step has better performance than using a series of prompts. We therefore experiment with various single prompts:

```
[article title + first 3 sentences]

Does the text discuss a recent
[event class c] in [location l]?
```

This wording is the simplest, asking if the event class in the fingerprint has occurred in the country specified by the fingerprint recently.

```
[article title + first 3 sentences]

Does the text discuss a recent
[event class c] [broader category
of c] in [location l]?
```

```
[article title + first 3 sentences]

Does the text discuss a recent
[event class c] [(broader category
of c)] in [location l]?
```

These two prompts clarify the broader category of the event class (e.g., natural disaster, terrorist attack), using slightly different wordings.

```
[article title + first 3 sentences]

Does the text discuss a recent
[event class c], where [definition
of c], in [location l]?
```

The fourth prompt provides the definition of the event class, where the definition is from the Merriam-Webster dictionary.

The results corresponding to various prompts are all within one percentage point of each other for precision, recall, and F1 score. To avoid prompt customization for each event class, FAME’s Phase Two LLM QA step uses the simplest prompt without any category or definition clarification.

C Appendix C: Reporting countries and example of identified terrorist attacks

Table 8: Reporting Countries by Official Language

| English (en) | Spanish (es) | French (fr) |
|----------------------|-----------------|--------------|
| United States (USA) | Mexico (MEX) | Canada (CAN) |
| United Kingdom (GBR) | Colombia (COL) | France (FRA) |
| Canada (CAN) | Spain (ESP) | |
| Australia (AUS) | Argentina (ARG) | |
| Ghana (GHA) | | |
| India (IND) | | |
| Ireland (IRL) | | |
| Kenya (KEN) | | |
| Nigeria (NGA) | | |
| South Africa (ZAF) | | |

Table 8 presents the list of reporting countries we considered in the regression on global news coverage.

Similar to Table 3, we also examined example terrorist attacks identified by our method as part of a quantitative robustness check. Table 9 lists the top 10 terrorist attacks that garnered the highest number of news reports. Most of these identified terrorist attacks align with online news records, except for one false positive case. Specifically, our database includes two attack events, “Woman shot in Temple Terrace” and “Boogaloo shooting in Oakland,” which overlap with the period of the “George Floyd protests” series of events. All three events share the same fingerprint. As a result, our method erroneously categorized thousands of news articles on the “George Floyd protests” as false positive samples associated with the fingerprint of an attack event.

| Event | Class (<i>c</i>) | Location (<i>l</i>) | Time (<i>t</i>) | News | Deaths |
|--|-------------------------|----------------------------|------------------------|-------------|---------------|
| Woman shot in Temple Terrace | Attack | United States | 2020-05-29 | 15094 | 1 |
| Boogaloo shooting in Oakland | Attack | Venezuela | 2020-05-03 | 1741 | 6 |
| Foiled coup attempt | Attack | United States | 2020-05-21 | 1712 | 1 |
| Shooting at Corpus Christi Navy base | Attack | Germany | 2020-02-19 | 1677 | 2 |
| Hanau shootings | Attack | Nigeria | 2020-05-31 | 1379 | 13 |
| Fulani militants attacked in Kaduna State | Attack | Afghanistan | 2020-05-12 | 1216 | 24 |
| Massacre in a maternity ward | Attack | United Kingdom | 2020-02-02 | 1161 | 1 |
| Streatham stabbing | Attack | Nigeria | 2020-01-06 | 1024 | 30 |
| Gamboru bombing | Attack | Mexico | 2020-05-16 | 1022 | 2 |
| Jorge Armenta: Mexican journalist killed | Attack | Nigeria | 2020-02-09 | 980 | 30 |
| Boko Haram militants burn to death motorists | Attack | | | | |

Table 9: Top 10 terrorist attacks with the most news in the first five months of 2020.

D Appendix D: Country factors in regression analysis

We extensively examined literature of the country factors whose impact on media coverage has been studied. All considered factors are listed (color-coded by the class) with detailed explanation in Table 10.

| Type | Predictor | description |
|------|--------------------------|---|
| Eve. | #Deaths | The average death number of each event happening in the country |
| | GDP | event in a high GDP country? |
| Eco. | GDP (h-h) | event in a high GDP country and reported by a high GDP country? |
| | GDP (h-l) | event in a high GDP country and reported by a low GDP country? |
| | GDP (l-h) | event in a low GDP country and reported by a high GDP country? |
| | GDP (l-l) | event in a low GDP country and reported by a low GDP country? |
| | Trade | the trading volume between the reporting country and the event country |
| | Investment | the amount (\$ USD) that the reporting country invest into the event country |
| | Democracy in. | event in a high democracy index (DI) country? |
| | Democracy in. (h-h) | event in a high democracy index country and reported by a high democracy index country? |
| | Democracy in. (h-l) | event in a high democracy index country and reported by a low democracy index country? |
| | Democracy in. (l-h) | event in a low democracy index country and reported by a high democracy index country? |
| | Democracy in. (l-l) | event in a low democracy index country and reported by a low democracy index country? |
| | Press Freedom in. | event in a high press freedom index (PFI) country? |
| | Press freedom in. (h-h) | event in a high press freedom index country and reported by a high press freedom index country? |
| | Press freedom in. (h-l) | event in a high press freedom index country and reported by a low press freedom index country? |
| | Press freedom in. (l-h) | event in a low press freedom index country and reported by a high press freedom index country? |
| | Press freedom in. (l-l) | event in a low press freedom index country and reported by a low press freedom index country? |
| Pol. | Federalism | event in a federal country? |
| | Republic | event in a republic country? |
| | Other mode of government | event in a country neither federal or republic? |
| | Both Federalism | both the event country and the reporting country are federal? |
| | Both Republic | both the event country and the reporting country are republic? |
| | Federalism and Other | event in a federal country and reported by a country neither federal or republic? |
| | Republic and Other | event in a republic country and reported by a country neither federal or republic? |
| | Diplomatic relation | an integer measure with six categories: 1-Unknown; 2-Interests Served by; 3-Interest Desk; 4-d'affairs; 5-Minister/Envoy; 6-Ambassador, Nuncio, Secretary of the People's Bureau; |
| | Population | event country population |
| | Population density | the population density of the event country |
| | Immigration | the immigration amount from the reporting country to the event country |
| | Gini in. | event in a high Gini index country? |
| Soc. | Gini in. (h-h) | event in a high Gini index country and reported by a high Gini index country? |
| | Gini in. (h-l) | event in a high Gini index country and reported by a low Gini index country? |
| | Gini in. (l-h) | event in a low Gini index country and reported by a high Gini index country? |
| | Gini in. (l-l) | event in a low Gini index country and reported by a low Gini index country? |
| | Area | the area of the event country |
| Geo. | Asia | event country in Asia? |
| | Europe | event country in Europe? |
| | Africa | event country in Africa? |
| | Oceania | event country in Oceania? |
| | North America | event country in North America? |
| | South America | event country in South America? |
| | Neighbors | event country and the reporting country are neighbors? |
| | Continent sim | event country and the reporting country on the same continent? |
| | Same Language | the event country and the reporting country using the same official language? |
| | Religion Diversity | the shannon entropy of the distribution with respect to mainstream religions (including Christian, Muslim, Unaffil, Hindu, Buddhist, and Folk religion) in the event country |
| Cul. | Literacy rate | the literacy rate of the reporting country |
| | Internet user rate | the internet user rate of the reporting country |

Table 10: Descriptions for the predictors considered in the regressions of news coverage.

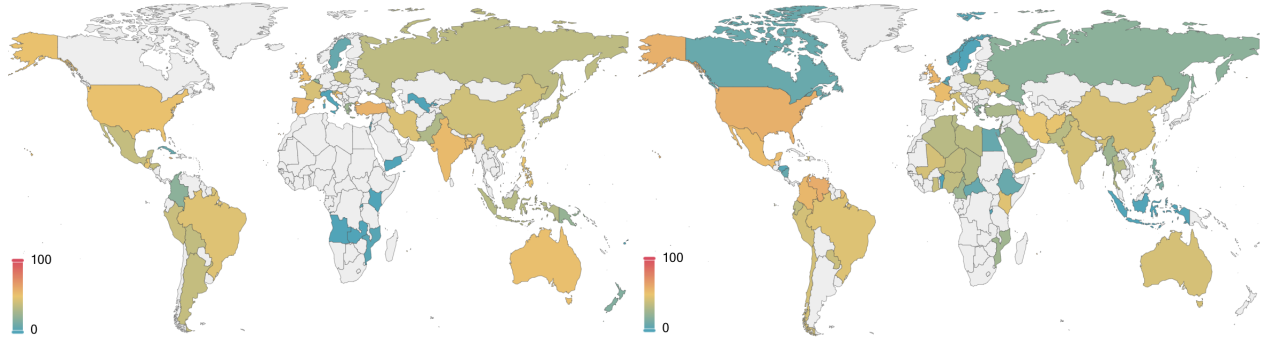


Figure 7: The average number of news in Spanish per disaster (left) and per terrorist attack (right) happening in a given country.

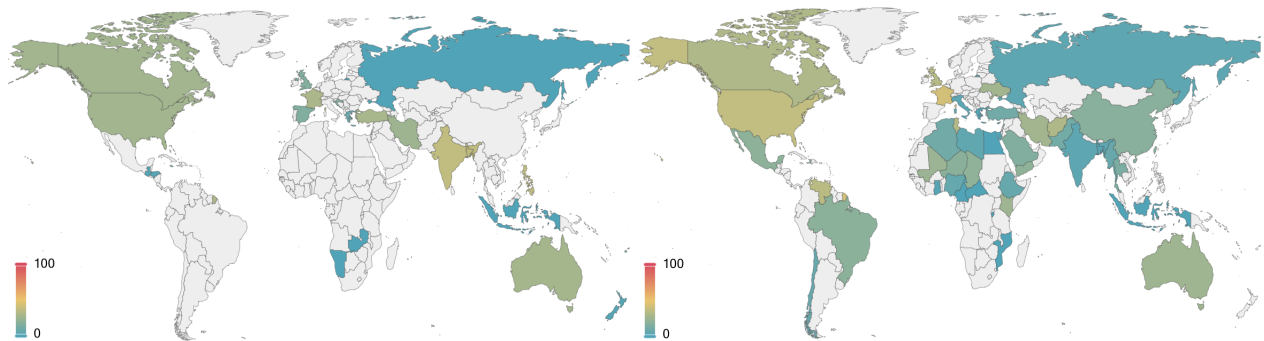


Figure 8: The average number of news in French per disaster (left) and per terrorist attack (right) happening in a given country.