

Reframing Responsibility: Framing-Aware Event Causality Identification

Anonymous ACL submission

Abstract

Causal explanations in political narratives are often framed and contested. Different sources may explain the same event by assigning responsibility to different actors, expressing different levels of certainty. Standard Event Causality Identification (ECI) focuses on detecting causal links and does not capture these distinctions. We introduce **Framing-Aware Event Causality Identification (FrECI)**, a framing-aware extension of ECI that models causal explanations as structured claims including responsibility targets, evaluative framing, source type, and epistemic modality grounded in established framing theories. We construct a multilingual dataset aligned across English, Chinese, and Arabic narratives using shared event anchors. We evaluate FrECI using prompt-based large language model baselines and supervised neural models. Results show that prompt-based baselines struggle to recover complete framed causal claims, while joint supervised models perform substantially better. Finally, we demonstrate that FrECI enables quantitative analysis of divergent causal attribution across narratives. The dataset and code will be made publicly available.¹



Figure 1: Framing-aware causal explanations for a shared anchor event. The same effect event (*Israeli airstrikes*) is causally explained in three ways, differing in responsibility targets, framing, source, and certainty.

1 Introduction

Causal explanations play a central role in political narratives. When texts describe contentious events, they rarely present causality in a purely neutral way. Instead, they attribute responsibility to specific actors, frame those actors evaluatively, invoke sources to legitimize claims, and signal varying degrees of epistemic certainty. As a result, different sources often offer distinct causal explanations for the same real-world event, even when they appear to describe the same facts.

Figure 1 illustrates this phenomenon for a single anchor event, *Israeli airstrikes*. Across three

documents, the same effect event is causally explained in different ways. In the first, airstrikes are framed as a response to Hamas rocket fire, exonerating Israel while assigning blame to Hamas. In the second, the same airstrikes are attributed to Israel's long-standing occupation and blockade of Gaza, shifting blame toward Israel and attributing the claim to an opposing source (*Palestinian officials*). In the third, airstrikes are discussed through a third-party lens (*international observers*), with neutral epistemic stance and neutral framing. Although all three sentences describe the same event, they differ systematically in *who is held responsible*, *how responsibility is framed*, *who advances the causal claim*, and *how certain the claim is*.

¹Anonymized for the review.

Existing NLP work on Event Causality Identification (ECI) is not designed to capture these distinctions. Prior ECI formulations primarily focus on identifying whether a causal relation holds between two events, typically treating causality as a binary relation detached from attribution, framing evaluation, or epistemic stance. This puts limits on the applicability of ECI to political discourse and framing analysis, where causal explanations are inherently framed and contested.

In this paper, we introduce **Framing-Aware Event Causality Identification (FRECI)**, a novel task that extends ECI to explicitly model the interpretive structure of causal explanations. Rather than predicting causal links alone, FRECI requires models to recover *framed causal claims* that jointly encode: (i) a directed causal relation between events, (ii) epistemic modality capturing the certainty of the claim, (iii) source type indicating who advances the explanation, and (iv) responsibility targets paired with evaluative framing effects such as blame or credit. This formulation allows models to distinguish between causally similar but rhetorically divergent explanations.

In communication research, causal attribution is not treated as a neutral description of events, but as a central mechanism of framing through which responsible target (Iyengar, 1993), moral evaluation (Entman, 1993), epistemic stance (Bloome and Talwalkar, 1997) and source legitimacy (Tuchman, 1972) are constructed. FRECI operationalizes this insight by modeling causal explanations as framed claims rather than factual relations.

To support this task, we construct a new multilingual dataset using a semi-automated pipeline that bootstraps from high-quality event and causal annotations in MAVEN-ERE (Wang et al., 2022) and expands them into a politically contentious setting. Starting from English Wikipedia articles, we align parallel Chinese and Arabic narratives through Cross-Document Event Coreference (CDEC), recover candidate causal relations using a hybrid Large Language Model (LLM)-assisted approach, and annotate fine-grained framing attributes through expert adjudication. The resulting dataset contains 2,203 framed causal relations grounded in shared real-world event anchors.

We evaluate FRECI under an anchor-based setting using prompt-based large language models, a supervised joint neural architecture, and a Supervised Fine-Tuning (SFT) generative LLM. Our results show that while LLMs capture aspects of

causal reasoning, explicit joint supervision is crucial for reliably recovering complete framed causal claims, particularly for responsibility attribution and framing effects.

Finally, we demonstrate that FRECI enables a new class of analyses beyond extraction accuracy. Because framed causal claims are aligned across documents via coreference anchors, FRECI provides a principled basis for measuring how narratives propose competing causal explanations for the same real-world outcomes.

In summary, this paper makes three contributions: (1) we introduce FRECI, a framing-aware extension of ECI; (2) we construct a multilingual, anchor-aligned dataset for studying framed causal explanations in political narratives; and (3) we provide both modeling benchmarks and empirical analyses demonstrating how FRECI enables quantitative study of divergent causal attribution across narratives.

2 FRECI

2.1 Task Overview

We introduce FRECI, a task that extends standard ECI by requiring models not only to identify causal links between events, but also to characterize the sociopolitical interpretation of those links.

Given a document and a set of event mentions corresponding to real-world events of interest (anchors), the model is required to identify causal relations involving those anchors and to predict structured attribution and framing information associated with each causal explanation.

2.2 Task Definition

Given a document D and a set of query event mentions $M_Q \subset \mathcal{S}(D)$, where $\mathcal{S}(D)$ denotes the set of all event mentions in D of a set of real-world event anchors A , the goal of FRECI is to predict a set of framed causal assertions expressed in the document.

A causal explanation may involve multiple responsibility targets. We therefore define the atomic unit of prediction as a **target-specific causal assertion**, represented as a tuple:

$$(e_c, e_e, s, a_s, t, f_t)$$

where:

- $e_c \in M_Q$ is the **cause event mention**,
- $e_e \in M_Q$ is the **effect event mention**,

- $s \in \mathcal{S}$ is the **source type** indicating who advances the causal claim (Author, Target, Ally, Opponent, Third_Party),
- $a_s \in \mathcal{A}$ is the **epistemic modality** capturing the certainty (Full_Affirmative, Partial_Affirmative, Neutral, Partial_Negative, Full_Negative) with which the causal relation $\langle e_c, e_e \rangle$ is asserted by the source s ,
- $t \in \mathcal{T}$ is a **responsibility target** held responsible for the causal outcome,
- $f_t \in \mathcal{F}$ is the **framing effect** assigned to target t (Credit, Blame, Undermine_Credit, Exonerate_Blame, and Framing_Neutral).

For a given causal relation between events (e_c, e_e) , multiple responsibility targets may be present. In such cases, the causal explanation is represented as a *set of tuples*, one per target:

$$\{(e_c, e_e, s, a_s, t_1, f_{t_1}), \dots, (e_c, e_e, s, a_s, t_k, f_{t_k})\}$$

source type s and epistemic modality a_s are **claim-level attributes** shared across all targets associated with the same causal relation, while framing effects f_t are **target-specific**.

A model prediction is considered correct if it correctly identifies the causal event pair (e_c, e_e) and recovers the complete set of target-specific tuples with the correct source type, epistemic modality, responsibility targets, and framing effects.

3 Dataset Construction Pipeline

We construct our dataset using a semi-automated pipeline that bootstraps from high-quality event annotations in existing resources and systematically expands them into a multilingual, adversarial political setting. The core idea is to anchor annotation on reliable event and causal structures, while exposing them to divergent narrative framings across languages and media contexts. The dataset construction pipeline consists of four stages: (1) multilingual corpus sourcing and expansion, (2) event extraction and CDEC, (3) hybrid causal relation identification, and (4) fine-grained attributive annotation of causal links.

3.1 Corpus Sourcing and Cross-Lingual Expansion

We use MAVEN-ERE as the seed corpus. MAVEN-ERE is a large-scale English Wikipedia dataset

with high-quality annotations of event mentions, event causal relations and event coreference. To target politically contentious events, we manually select 51 MAVEN-ERE training articles whose topics are likely to elicit divergent interpretations across linguistic and cultural communities. For each English article, we identify the corresponding Chinese and Arabic Wikipedia articles, yielding parallel narratives of the same real-world events.

Because MAVEN-ERE consists of curated English excerpts while the Chinese and Arabic articles are substantially longer, we segment the non-English texts into multiple documents. We then translate them into English using gpt-4o, with prompts designed to preserve original causal wording and rhetorical structure (see Appendix 6). All translations are validated by fluent annotators. This process enables annotation in a unified English interface while preserving cross-lingual signals of framing divergence in causal attribution.

3.2 Event Extraction and CDEC

MAVEN-ERE provides gold event mentions and causal relations for the English seed articles. For the translated Chinese and Arabic documents, we extract event mentions using OmniEvent (Peng et al., 2023). We then select MAVEN-ERE events that participate in annotated causal relations and use them as anchors for cross-lingual alignment.

Using a pairwise cross-encoder CDEC model (Yu et al., 2022), we link event mentions in the translated documents to their corresponding MAVEN-ERE anchor events, forming cross-lingual CDEC clusters. These clusters connect mentions of the same real-world *cause* and *effect* events across articles and languages.

For each target event, the resulting structure captures multiple, potentially competing, causal attributions proposed by different sources. This alignment enables dataset to quantify causal attribution divergence by comparing how identical outcomes are causally explained across narratives.

3.3 ECI

We identify causal relations using a hybrid approach that combines LLM-based generative ECI with human refinement. Given the extracted event mentions and full document context, we prompt gpt-4o to propose candidate cause-effect relations, enabling the recovery of long-range and implicitly framed causal links that are difficult to capture with trained pairwise classifiers alone

Statistic	English	Chinese [†]	Arabic [†]	Total
Topics	49+2 [‡]	49	2	51
Documents	51	612	21	661
Tokens	22,098	406,911	10,986	439,995
Events	817	4,479	224	5520
Coref. Chains		–		756
Non-S. Chains		–		569
Causal Relations	845	1,244	114	2,203
Responsible Targets	1774	2736	265	4775

Table 1: Dataset statistics. [†] denotes translated data. [‡] indicates that 49 English topics align with Chinese, while 2 English topics align with Arabic (topics are listed in Appendix 5). Coreference Chains and Non-Singleton Coreference Chains cross-document only and therefore reported as corpus-level totals.

(Cheng et al., 2025). The prompt is provided in Appendix 7.

The LLM outputs are treated as high-recall candidates. Human annotators subsequently validate, correct, and refine these proposed causal relations by removing spurious links and resolving ambiguities. This hybrid strategy balances scalability with annotation quality.

3.4 Causal Framing Attributes Annotation

For each validated causal relation, annotators apply a structured annotation schema to capture how responsibility are framed within sociopolitical narratives. First, we use Semantic Role Labeling (SRL)² to automatically identify candidate agents associated with the cause and effect events. Annotators then validate or revise these agent annotations. When an event has no explicit agent and denotes a broad societal outcome (e.g., economic decline, public health), the target defaults to the Incumbent.

Second, for each identified target, annotators label the framing effect (e.g., Blame, Undermine_Credit). For modality annotation, we assist annotation with rule-based detectors that flag candidate quoted sources and epistemic certainty cues (e.g., “allegedly”). Annotators validate these suggestions and finalize source type labels (e.g., Author, Ally, Opponent) and certainty levels.

3.5 Annotation Process

We employ a multi-stage annotation pipeline that combines model-assisted pre-processing with

²https://huggingface.co/cu-kairos/propbank_srl_seq2seq_t5_large

Value	Framing		Source		Certainty	
	Value	%	Value	%	Value	%
Blame	21.4		Author	92.3	Full Aff.	90.9
Credit	24.2		3rd Party	2.5	Partial Aff.	6.4
Undermine Credit	9.8		Target	1.8	Neutral	1.2
Exonerate Blame	13.9		Opponent	2.2	Partial Neg.	0.3
Neutral	30.7		Ally	1.2	Full Neg.	1.1

Table 2: Framing-aware causal attributes distribution.

structured human adjudication to ensure annotation quality while managing the complexity of end-to-end causal framing analysis (see Figure 3 in Appendix for annotation interface).

3.5.1 Annotator Recruitment and Training

We recruited three annotators with backgrounds in computational linguistics. Annotators completed a calibration phase using a detailed guideline (see Table 2 in Appendix). A pilot annotation of 40 documents was jointly reviewed to resolve disagreements and align interpretations.

3.5.2 Two-Pass Annotation Workflow

We adopt a verify-and-refine workflow. For the causal relation verification in the first phase, annotators reviewed LLM-proposed causal relations, validating correct links, rejecting false positives, and adding missed causal relations.

For framing attributes annotation in the second phase, for each validated causal relation, annotators verified target and source spans and labeled framing attributes (Framing Effect, Source Category, and Certainty). Annotators were instructed to rely strictly on text-internal attribution signals, except when determining Ally/Opponent relations.

3.5.3 Adjudication and Agreement

To ensure consistency, 20% of topics were double-annotated. Disagreements were resolved through expert adjudication by a senior annotator, and recurring ambiguities informed iterative guideline refinement.

Inter-annotator agreement was evaluated separately for causal link identification, responsibility target identification, and framing attribution. For causal link identification, we report pairwise link F1. For framing annotations, agreement is computed only on causal links accepted by both annotators. Responsibility targets are treated as a set of text spans associated with each causal link. Target identification agreement is measured using overlap-based F1 with normalized span matching. Given

Component	Unit	Agreement
Causal Link	Event Pair	0.65
Target Identification	Span Set	0.71
Framing Effect	Target-level	0.62
Source Type	Claim-level	0.74
Epistemic Modality	Claim-level	0.88

Table 3: Inter-annotator agreement for FRECI. Agreement is measured using pairwise link F1 for causal link identification, span-level F1 for responsibility target identification, and Cohen’s κ for categorical framing, source, and epistemic modality annotations.

a matched causal link and responsibility target, annotators assign a target-specific framing effect. Agreement for the framing effect is computed per matched target using Cohen’s κ . In contrast, source type and epistemic modality are treated as causal claim level attributes shared across all targets associated with the same causal link. Agreement for source and modality is therefore computed once per causal link using Cohen’s κ .

Table 3 shows consistent inter-annotator agreement across FRECI components. Lower agreement for framing effects primarily reflects boundary cases in which weak or implicit attribution—such as subtle blame, credit, or exoneration—is difficult to distinguish from neutral framing, rather than systematic annotation error.

4 Experiments

We evaluate FRECI under an anchor-based setting, where gold anchor event mentions are provided as input. Given a document and a set of anchor event mentions corresponding to real-world events of interest, models are required to identify framed causal relations among these anchors and predict all associated attribution attributes.

4.1 Experimental Setup

Data Split. To avoid topic and event leakage, we construct topic-held-out splits. All documents associated with the same topic—including English articles and their translated Chinese or Arabic counterparts—are assigned to the same split. Specifically, we partition the 51 topics into 36 topics for training, 5 topics for development, 10 topics for testing. This ensures that models are evaluated on previously unseen political events and controversies (details in Appendix 5).

Candidate Anchor Pair Generation. For each document, models form candidate cause–effect

pairs by pairing each anchor-linked event mention with all event mentions in the same document. To control computational complexity, we restrict candidates to ordered mention pairs whose two event mentions occur within a fixed sentence window. All models are evaluated on the same candidate set to ensure fair comparison.

4.2 Models

Prompt-based LLM Baselines. We implement prompt-based LLM baselines using GPT-4o in zero-shot and chain-of-thought (CoT) settings. The model is prompted to identify framed causal claims involving gold anchor events and to output structured predictions in a constrained JSON format. Following the FRECI schema, GPT-4o predicts causal links between anchor events, assigns claim-level attributes (source type and Epistemic Modality), and generates a set of responsibility targets with target-specific Framing Effects.

In the zero-shot setting, the model is provided only with a task description, label definitions, and output schema. In the CoT setting, the model is additionally instructed to perform step-by-step reasoning prior to producing the final structured output; however, only the final output is evaluated. Both settings use identical output schemas and evaluation protocols. Full prompt templates 8 and 9 are provided in Appendix.

Neural Joint FRECI Model. We implement a joint neural model based on a pretrained RoBERTa encoder that factorizes causal reasoning and framing attribution into task-specific heads while sharing a common encoder and being trained end-to-end.

Given a document and an anchor-linked event mention, together with a candidate event mention from the same document, the model encodes the document using a pretrained RoBERTa_{LARGE} with explicit anchor markers. Contextual representations of the anchor mentions and the document are combined into a shared pair representation.

On top of this representation, the model jointly predicts: (i) a directional causal link between the event mention pair, (ii) claim-level attributes including source type and Epistemic Modality, and (iii) responsibility targets and their associated Framing Effects. Responsibility targets are predicted by classifying text spans of candidate agent identified via SRL, and framing effects are predicted conditionally for each selected target. All

408 components are trained jointly with a multi-task
409 objective, enabling structured information sharing
410 across causal and framing decisions.

411 **SFT LLM Baseline.** We additionally report
412 an open-model baseline obtained by SFT
413 LLaMA-3.1-8B-Instruct with QLoRA to directly
414 generate framed causal claims. The model is
415 trained to map from a document and a set of gold
416 anchor event mentions to a JSON-formatted out-
417 put representing complete causal claims, including
418 claim-level attributes and target-specific framing.
419 This approach performs *implicit joint modeling*, as
420 all components of FRECI are predicted simulta-
421 neously via sequence generation without explicit
422 task-specific heads. All predictions are evaluated
423 using the same metrics as other models.

424 4.3 Evaluation Metrics

425 Models predict a set of target-specific causal asser-
426 tion tuples of the form $(e_c, e_e, s, a_s, t, f_t)$. Evalu-
427 ation is performed at both the tuple level and the
428 causal-relation level.

429 **Causal Link Identification.** We evaluate causal
430 link detection independently of framing attributes
431 by computing precision, recall, and F1 over di-
432 rected anchor event pairs (e_c, e_e) . A predicted link
433 is considered correct if both event mentions and
434 direction match a gold causal link.

435 **Responsibility Target Identification.** For cor-
436 rectly predicted causal links, responsibility targets
437 t are evaluated as a set of text spans. We report
438 span-level micro-F1 using one-to-one matching
439 with an IoU threshold of 0.5, where IoU is defined
440 as the token-level intersection over union between
441 predicted and gold target spans. Each matched
442 (e_c, e_e, t) corresponds to a target-specific causal
443 assertion tuple.

444 **Framing Effect Prediction.** Given a correctly
445 predicted causal link and matched responsibility
446 target, the framing effect f_t is evaluated at the
447 tuple level using macro-F1 over matched (e_c, e_e, t)
448 instances.

449 **Claim-Level Attributes.** source type s and epis-
450 temic modality a_s are evaluated once per causal
451 relation (e_c, e_e) , since they are shared across all
452 target-specific tuples derived from the same causal
453 explanation. We report macro-F1 over these claim-
454 level attributes.

455 **Full-Claim Exact Match (EM).** We addition-
456 ally report Full-Claim EM, a strict end-to-end
457 metric under which a predicted causal explana-
458 tion is counted as correct if and only if the com-
459 plete set of target-specific causal assertion tuples
460 $(e_c, e_e, s, a_s, t, f_t)$ associated with a causal relation
461 is recovered exactly. This requires correct identi-
462 fication of the causal event pair, the shared source
463 type and epistemic modality, the full responsibility
464 target set, and the framing effect assigned to each
465 target.

466 4.4 Results

467 Table 4 reports anchor-based FRECI results on
468 held-out test topics. Overall, the results show that
469 recovering complete framed causal claims is chal-
470 lenging, and that explicit joint modeling substan-
471 tially outperforms prompt-based and generative
472 baselines.

473 **Causal Link Identification.** Prompt-based
474 GPT-4o exhibits a clear precision–recall trade-off.
475 Zero-shot prompting yields moderate recall but
476 low precision, while Chain-of-Thoughts (CoT)
477 prompting substantially increases precision at the
478 cost of recall, resulting in lower F1. In contrast,
479 supervised models achieve much higher causal
480 link F1, with the Joint FRECI model performing
481 best (59.0), followed closely by the SFT LLaMA
482 baseline (56.5). These results indicate the
483 importance of task-specific supervision for robust
484 causal link detection.

485 **Attribution Components.** Responsibility target
486 identification achieves relatively high performance
487 across models when conditioned on correct causal
488 links. However, framing effect prediction re-
489 mains the most difficult component. GPT-4o base-
490 lines perform poorly on framing effect ($f \approx 28$ –
491 29 macro-F1), whereas the Joint FRECI model
492 achieves a substantial improvement (73.1), outper-
493 forming the SFT LLaMA baseline (61.2). Source
494 type and epistemic modality achieve uniformly
495 high scores across models, reflecting both their re-
496 lative stability and strong label skew in the data,
497 where most causal claims are attributed to the
498 Author and expressed with Full_Affirmative.

499 **End-to-End Evaluation.** EM provides a strict
500 end-to-end measure, requiring correct prediction
501 of the causal link, all responsibility targets, and
502 all framing attributes. As expected, EM scores
503 are much lower than individual component scores.

Model	Causal Link			Attribution Components				Full Claim
	P	R	F1	T	f	s	m	EM
GPT-4o (Zero-shot)	34.0	41.5	37.4	73.1	28.4	74.4	91.5	14.6
GPT-4o (CoT)	61.5	17.0	26.6	81.4	29.1	79.4	98.5	15.2
Joint FRECI (RoBERTa)	49.7	72.6	59.0	77.3	73.1	75.6	96.0	27.4
SFT LLaMA-3.1-8B	47.7	69.3	56.5	80.8	61.2	72.3	96.0	26.7

Table 4: Anchor-based FRECI results on held-out test topics. Causal link performance is reported as precision (P), recall (R), and F1 over directed anchor event pairs. Responsibility Target (T) is evaluated using span-level micro-F1, conditioned on correct causal links. Framing Effect (f) is evaluated using macro-F1 over matched targets. source type (s) and Epistemic Modality (m) are evaluated using macro-F1 at the claim level. Full-Claim EM requires correct prediction of the causal link, complete target set, claim-level attributes, and all target-specific framing effects.

Prompt-based GPT-4o achieves low EM (14.6–15.2), while both supervised models substantially improve performance. The Joint FRECI model achieves the highest EM (27.4), narrowly outperforming SFT LLaMA (26.7), indicating that explicit joint modeling yields more consistent structured predictions.

Together, these results show that recovering complete framed causal claims requires models to jointly capture causal structure and responsibility attribution. While LLMs capture aspects of causal reasoning under prompting, explicit joint supervision is crucial for achieving consistent end-to-end performance under the strict EM metric.

5 Competing Causes for Shared Outcomes

FRECI enables quantitative analysis of how narratives propose competing causal explanations for the same real-world outcome. Because event mentions are linked into CDEC clusters, we analyze causal competition at the level of event clusters.

Let E denote an effect event cluster representing a shared outcome. We collect all annotated causal relations whose effect event belongs to E . Each causal relation proposes a cause event that belongs to some cause cluster C , inducing an empirical distribution over cause clusters:

$$p(C | E) = \frac{n(C, E)}{\sum_{C'} n(C', E)},$$

where $n(C, E)$ counts causal relations linking a cause from cluster C to an effect in cluster E . Each causal relation contributes one vote, independent of the number of responsibility targets.

Causal Fragmentation. We quantify the degree to which an outcome admits competing causal ex-

planations using *causal fragmentation*, defined as the normalized entropy of the cause distribution:

$$\text{Frag}(E) = \frac{H(p(C | E))}{\log |C(E)|},$$

where $H(\cdot)$ is Shannon entropy and $C(E)$ is the set of distinct cause clusters observed for E . $\text{Frag}(E) \in [0, 1]$ is 0 when all explanations concentrate on a single cause cluster and increases as explanations spread more evenly across alternative causes.

Our dataset contains 569 non-singleton event clusters and 2,203 annotated causal relations. We compute causal fragmentation for effect clusters with at least two associated causal relations to exclude degenerate single-claim cases. Figure 1 illustrates this phenomenon: the same outcome (*Israeli airstrikes*) is explained by multiple distinct cause clusters (e.g., *rocket fire*, *occupation/blockade*, *regional developments*), resulting in high causal fragmentation. Across eligible effect clusters, the average causal fragmentation is 0.44, indicating that many outcomes are associated with multiple competing causal explanations rather than a single dominant cause.

6 Related Work

6.1 Event Causality Identification

ECI has been extensively studied as the task of detecting binary cause–effect relations between events in text. Widely used benchmarks such as CausalTimeBank (Mirza et al., 2014), the Event StoryLine Corpus (Caselli and Vossen, 2017), and MAVEN-ERE (Wang et al., 2022) provide annotations of event-level causality within documents. These datasets do not model who advances a causal explanation, how responsibility is framed, or how

certainty is expressed, and they largely assume a neutral, context-independent notion of causality.

Methodologically, early neural approaches framed ECI as pairwise classification using contextual encoders (Zhang et al., 2015; Kadowaki et al., 2019). Subsequent work introduced document-level modeling, constructing event graphs and applying structured reasoning with graph neural networks or transformers to capture long-range dependencies and global consistency (Gao et al., 2019; Chen et al., 2022, 2023). More recently, prompt-based and generative approaches using large language models have shown strong few-shot and zero-shot performance (Shen et al., 2022; Xiang et al., 2023; Man et al., 2024). However, multiple studies report that LLM-based methods are sensitive to prompt design and prone to over-attributing causality, particularly in long or ambiguous contexts (Gao et al., 2023; Kiciman et al., 2024). As a result, recent surveys and evaluations suggest that supervised models with explicit event representations remain the most reliable approach for ECI, with LLMs best used as auxiliary components (Wang et al., 2022; Chen et al., 2023; Cai et al., 2025).

FRECI builds on this line of work but departs from standard ECI by treating causal explanations as *framed claims*. In addition to identifying causal links, FRECI models responsibility attribution, evaluative framing, source type, and epistemic modality—dimensions not represented in existing ECI datasets or models.

6.2 Computational Framing

Computational approaches to framing have primarily focused on identifying issue- or topic-level frames in news and political discourse. Prominent datasets include the Media Frames Corpus (Card et al., 2015) and its extensions (Piskorski et al., 2023; Ajour et al., 2019; Mendelsohn et al., 2021), BU-NEmo (Reardon et al., 2022), and issue-specific resources such as the Gun Violence Frame Corpus (Liu et al., 2019), VoynaSlov (Park et al., 2022), and stereomigrants (Sánchez-Junquera et al., 2021). Related work also addresses media attitude and bias detection (Zhao et al., 2024; Hamborg et al., 2019). These datasets have supported a range of computational methods, including topic modeling (DiMaggio et al., 2013; Nguyen et al., 2015), unsupervised learning (Burscher et al., 2016), semantic parsing (Ziems and Yang, 2021), and fine-tuned language models (Mendelsohn et al.,

2021).

More recent work adopts an event-centric perspective on framing. Liu et al. (2023) aligns news articles covering the same story to analyze partisan event selection, while Das et al. (2024) clusters event relations into narrative structures. Zhao et al. (2024) examines contextual events surrounding a main event to identify framing through selection and omission. While these approaches move framing analysis closer to event structure, they do not explicitly extract causal claims of the form “X caused Y”, nor do they model whether such claims assign blame, credit, or justification.

6.3 Cross-Lingual Narrative Analysis

Our work is also related to studies of narrative divergence across languages. Prior research using multilingual Wikipedia and parallel news corpora has shown that articles describing the same historical or political events often reflect language-specific or community-specific perspectives (Hecht and Gergle, 2010; Bao et al., 2012; Hale, 2014). Despite efforts such as Wikipedia’s Neutral Point of View policy, cross-lingual analyses consistently find differences in emphasis, interpretation, and framing.

However, existing cross-lingual studies typically focus on high-level content differences, such as fact selection, topic emphasis, or sentiment, rather than aligning specific causal explanations. They do not provide structured representations that allow side-by-side comparison of how different sources attribute causality and responsibility for the same event. FRECI enables the first structured analysis of divergent causal explanations for the same events across sources.

7 Conclusion

We introduced FRECI, a framing-aware event causality task that models not only causal relations, but how they are attributed and framed in political narratives. By anchoring explanations on shared real-world events, FRECI enables controlled cross-source and cross-lingual comparison of divergent causal attributions. Experiments show that recovering complete framed causal claims is challenging, and that joint supervision clearly outperforms prompt-based baselines. Overall, FRECI provides both a challenging NLP benchmark and a practical framework for analyzing framed causality in political discourse.

672 Limitations

673 Our study focuses on Wikipedia articles, which
674 differ from time-sensitive news reporting in both
675 authorship and editorial process. Consequently,
676 the causal narratives captured in our dataset re-
677 flect stabilized, negotiated representations rather
678 than immediate journalistic reactions. This makes
679 Wikipedia a conservative testbed for studying fram-
680 ing: although a Neutral Point of View policy is
681 enforced, neutrality is negotiated separately within
682 each language edition, allowing framing differ-
683 ences to persist despite editorial normalization. As
684 a result, our findings may not directly generalize
685 to real-time media dynamics.

686 Our framing and source annotations rely on a
687 discrete label set that enables clear operationaliza-
688 tion and evaluation. However, real-world causal
689 explanations may exhibit mixed, graded, or evol-
690 ving framing that is not fully captured by categorical
691 labels. Future work could explore continuous or
692 probabilistic representations of causal responsibil-
693 ity and framing.

694 We evaluate FRECI under an *anchor-based* set-
695 ting, where gold anchor event mentions are pro-
696 vided as input, isolating framed causal reason-
697 ing from upstream extraction and alignment er-
698 rors. FrECI adopts an anchor-based formulation
699 to enable controlled comparison of causal expla-
700 nations across documents while avoiding unre-
701 stricted event pairing, making it well suited for
702 cross-source and cross-lingual analysis. End-to-
703 end and cross-narrative settings introduce addi-
704 tional challenges and are left for future work.

705 Because non-English articles are translated into
706 English for annotation, some language-specific
707 framing cues may be altered or attenuated despite
708 human validation. As a result, certain cross-lingual
709 framing differences may be understated or over-
710 regularized.

711 Our dataset covers three languages and does
712 not capture framing variation in other linguistic or
713 geopolitical contexts. Extending FRECI to addi-
714 tional languages remains future work.

715 Ethical Considerations

716 This work analyzes political narratives and
717 causal explanations drawn from publicly available
718 Wikipedia articles. The dataset contains no private,
719 personal, or sensitive individual-level information,
720 and all source texts are released under Wikipedia’s
721 licensing terms. As such, the work poses minimal

risk to individual privacy.

Our annotations model how causal explanations
assign responsibility and evaluative framing within
text. These labels are intended to describe narra-
tive structure rather than to endorse any particu-
lar political interpretation. Nevertheless, framing
annotations necessarily involve interpretive judg-
ment. To mitigate annotator bias, we provide de-
tailed annotation guidelines, conduct calibration
and adjudication, and restrict annotators to labeling
text-internal signals rather than external beliefs.

The dataset reflects the perspectives and norms
of Wikipedia editor communities, which may not
represent broader public discourse. As a result,
analyses derived from this resource should not
be interpreted as measuring societal consensus or
ground truth about political events. Users of the
dataset should exercise caution when drawing nor-
mative or policy conclusions.

Finally, while the proposed task and models can
support analysis of political framing, they could
also be misused to amplify or target particular nar-
ratives. We release this dataset and framework
for research purposes and encourage responsible
use, transparency, and critical interpretation when
applying FRECI to real-world political content.

Acknowledgements

References

- Yamen Ajjour, Milad Alshomary, Henning Wachsmuth,
and Benno Stein. 2019. [Modeling frames in argu-
mentation](#). In *Conference on Empirical Methods in
Natural Language Processing*.
- Patti Bao, Brent J. Hecht, Samuel Carton, Mahmood
Quaderi, Michael S. Horn, and Darren Gergle. 2012.
[Omnipedia: bridging the wikipedia language gap](#).
*Proceedings of the SIGCHI Conference on Human
Factors in Computing Systems*.
- David Bloome and Susan Talwalkar. 1997. [Book re-
views: Critical discourse analysis and the study of
reading and writing](#). *Reading Research Quarterly*,
32(1):104–112.
- Bjorn Burscher, Rens Vliegthart, and Claes H. de
Vreese. 2016. [Frames beyond words: Applying clus-
ter and sentiment analysis to news coverage of the
nuclear power issue](#). *Social Science Computer Re-
view*, 34(5):530–545.
- Ruichu Cai, Shengyin Yu, Jiahao Zhang, Wei Chen,
Boyan Xu, and Keli Zhang. 2025. [Dr.ECI: Infusing
large language models with causal knowledge for
decomposed reasoning in event causality identifica-
tion](#). In *Proceedings of the 31st International Con-
ference on Computational Linguistics*, pages 9346–

774	9375, Abu Dhabi, UAE. Association for Computational Linguistics.		
775			
776	Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues . In <i>Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)</i> , pages 438–444, Beijing, China. Association for Computational Linguistics.		830
777			831
778			832
779			833
780			834
781			835
782			836
783			
784			
785	Tommaso Caselli and Piek Vossen. 2017. The event StoryLine corpus: A new benchmark for causal and temporal relation extraction . In <i>Proceedings of the Events and Stories in the News Workshop</i> , pages 77–86, Vancouver, Canada. Association for Computational Linguistics.		837
786			838
787			839
788			
789			
790			
791	Meiqi Chen, Yixin Cao, Kunquan Deng, Mukai Li, Kun Wang, Jing Shao, and Yan Zhang. 2022. ERGO: Event relational graph transformer for document-level event causality identification . In <i>Proceedings of the 29th International Conference on Computational Linguistics</i> , pages 2118–2128, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.		840
792			841
793			842
794			843
795			844
796			
797			
798			
799	Meiqi Chen, Yixin Cao, Yan Zhang, and Zhiwei Liu. 2023. CHEER: Centrality-aware high-order event reasoning network for document-level event causality identification . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 10804–10816, Toronto, Canada. Association for Computational Linguistics.		845
800			846
801			847
802			848
803			849
804			
805			
806			
807	Qing Cheng, Zefan Zeng, Xingchen Hu, Yuehang Si, and Zhong Liu. 2025. A survey of event causality identification: Taxonomy, challenges, assessment, and prospects .		850
808			851
809			852
810			
811	Rohan Das, Aditya Chandra, I-Ta Lee, and Maria Leonor Pacheco. 2024. Media framing through the lens of event-centric narratives. Note: media framing through event perspective: Rohan used event temporal relations .		853
812			854
813			855
814			856
815			857
816	Paul DiMaggio, Manish Nag, and David Blei. 2013. Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of u.s. government arts funding . <i>Poetics</i> , 41(6):570–606. Topic Models and the Cultural Sciences.		858
817			859
818			860
819			861
820			
821			
822	Robert Entman. 1993. Framing: Toward clarification of a fractured paradigm . <i>The Journal of Communication</i> , 43:51–58.		862
823			863
824			864
825	Jinglong Gao, Xiao Ding, Bing Qin, and Ting Liu. 2023. Is chatgpt a good causal reasoner? a comprehensive evaluation .		865
826			866
827			867
828	Lei Gao, Prafulla Kumar Choubey, and Ruihong Huang. 2019. Modeling document-level causal structures for event causal relation identification . In <i>Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)</i> , pages 1808–1817, Minneapolis, Minnesota. Association for Computational Linguistics.		868
829			869
830			870
831			871
832			872
833			
834			
835			
836			
837	Scott A. Hale. 2014. Multilinguals and wikipedia editing . In <i>Proceedings of the 2014 ACM conference on Web science, WebSci '14</i> , page 99–108. ACM.		873
838			874
839			875
840			876
841	Felix Hamborg, Anastasia Zhukova, and Bela Gipp. 2019. Automated identification of media bias by word choice and labeling in news articles . In <i>Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL)</i> .		877
842			878
843			879
844			880
845			881
846			882
847			883
848			884
849			885
850	Brent Hecht and Darren Gergle. 2010. The tower of babel meets web 2.0: user-generated content and its applications in a multilingual context . In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10</i> , page 291–300. ACM.		
851			
852			
853	Shanto Iyengar. 1993. Is anyone responsible?: How television frames political issues . <i>American Journal of Sociology</i> , 98(6):1459–1462.		
854			
855			
856			
857			
858			
859			
860			
861			
862	Kazuma Kadowaki, Ryu Iida, Kentaro Torisawa, Jong-Hoon Oh, and Julien Kloetzer. 2019. Event causality recognition exploiting multiple annotators' judgments and background knowledge . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 5816–5822, Hong Kong, China. Association for Computational Linguistics.		
863			
864			
865			
866			
867			
868			
869			
870			
871			
872			
873	Emre Kıcıman, Robert Ness, Amit Sharma, and Chenhao Tan. 2024. Causal reasoning and large language models: Opening a new frontier for causality .		
874			
875			
876			
877			
878			
879			
880			
881			
882			
883			
884			
885			
886			
887			
888			
889			
890			
891			
892			
893			
894			
895			
896			
897			
898			
899			
900			
901			
902			
903			
904			
905			
906			
907			
908			
909			
910			
911			
912			
913			
914			
915			
916			
917			
918			
919			
920			
921			
922			
923			
924			
925			
926			
927			
928			
929			
930			
931			
932			
933			
934			
935			
936			
937			
938			
939			
940			
941			
942			
943			
944			
945			
946			
947			
948			
949			
950			
951			
952			
953			
954			
955			
956			
957			
958			
959			
960			
961			
962			
963			
964			
965			
966			
967			
968			
969			
970			
971			
972			
973			
974			
975			
976			
977			
978			
979			
980			
981			
982			
983			
984			
985			
986			
987			
988			
989			
990			
991			
992			
993			
994			
995			
996			
997			
998			
999			
1000			

886	Julia Mendelsohn, Ceren Budak, and David Jurgens.	International Committee on Computational Linguistics.	944
887	2021. Modeling framing in immigration discourse on social media .		945
888	In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 2219–2263, Online.		
889	Association for Computational Linguistics.		
890			
891			
892			
893	Paramita Mirza, Rachele Sprugnoli, Sara Tonelli, and Manuela Speranza. 2014. Annotating causality in the TempEval-3 corpus .		949
894	In <i>Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CAtoCL)</i> , pages 10–19, Gothenburg, Sweden.		950
895	Association for Computational Linguistics.		951
896			952
897			953
898			954
899			955
900			956
901			957
902			
903			
904			
905			
906			
907			
908			
909			
910			
911			
912			
913			
914			
915			
916			
917			
918			
919			
920			
921			
922			
923			
924			
925			
926			
927			
928			
929			
930			
931			
932			
933			
934			
935			
936			
937			
938			
939			
940			
941			
942			
943			
944			
945			
946			
947			
948			
949			
950			
951			
952			
953			
954			
955			
956			
957			
958			
959			
960			
961			
962			
963			
964			
965			
966			
967			
968			
969			
970			
971			
972			
973			
974			
975			
976			
977			
978			
979			
980			
981			
982			
983			
984			
985			

A Appendix

A.1 FRECI Annotation Guidelines

See Figure 2

A.2 Annotation and Analysis Interfaces

To support consistent and efficient annotation of framed causal claims, we developed a web-based interface that presents anchor events, candidate causal relations, and structured framing attributes in a unified view. The interface enforces the FRECI schema and reduces annotator error by constraining label choices and maintaining shared claim-level attributes. See Figure 3 We additionally built a visualization interface to support qualitative analysis of framing-divergent causal explanations across aligned narratives. See Figure 4

A.3 Data Split

See Table 5

A.4 Prompt Templates for Translation

See Table 6

A.5 Prompt Templates for ECI

See Table 7

A.6 Prompt Templates for Prompt-based LLM Baselines

See zero-shot prompt at Table 8. See CoT prompt at Table 9

A.7 Training Details for Joint FRECI model.

The Joint FRECI model is initialized with a pre-trained RoBERTa_{LARGE} and fine-tuned end-to-end on the FRECI task. Input documents are truncated to a maximum length of 512 tokens. Special marker tokens are inserted to indicate the spans of anchor event mentions.

Candidate responsibility targets are extracted using semantic role labeling and restricted to agentive argument roles.

The model is trained using a multi-task objective that jointly optimizes all prediction heads. Directional causal link prediction and claim-level attributes (source type and Epistemic Modality) are optimized using cross-entropy loss. Responsibility target selection is optimized using binary cross-entropy loss over candidate targets. Target-specific Framing Effects are optimized using cross-entropy loss conditioned on gold responsibility targets. Attribute losses are masked for non-causal event pairs.

We use the AdamW optimizer with a learning rate of 2×10^{-5} . Models are trained for up to 10 epochs with early stopping based on development set Full-Claim EM. Batch size is 8.

A.8 Training Details for SFT LLM Baseline

We fine-tune LLaMA-3.1-8B-Instruct using SFT with QLoRA to produce structured FRECI outputs. Each training instance consists of a single document paired with its gold anchor event mentions as input, and a JSON-formatted list of framed causal claims as output. The output schema matches the prompt-based LLM baselines, including causal links, claim-level attributes, and target-specific framing.

Model and Optimization. We apply 4-bit quantization with NF4 and train LoRA adapters on attention projection layers. Training is performed using the AdamW optimizer with a learning rate of 2×10^{-4} . Models are trained for up to 3 epochs with early stopping based on development set Full-Claim EM.

LoRA Configuration. We use LoRA rank $r = 32$, scaling factor $\alpha = 64$, and dropout of 0.05. The maximum input sequence length is set to 2048 tokens. Batch size is adjusted via gradient accumulation to achieve a global batch size between 64 and 128.

Decoding and Output Constraints. At inference time, decoding is performed with greedy or low-temperature sampling. The model is instructed to output valid JSON only. Generated outputs are post-processed to remove extraneous text and validated for schema correctness prior to evaluation. Invalid or unparsable outputs are treated as empty predictions.

Evaluation Protocol. All outputs from the SFT LLM baseline are evaluated using the same evaluation protocol as other models, including causal link F1, attribution component scores, and Full-Claim EM. This ensures a fair comparison between implicit generative joint modeling and explicitly structured neural approaches.

Split	Topics
Train (n=36)	1953 Iranian Coup d’État, 2011 Chinese Pro-Democracy Protests, 2013 Tiananmen Square Attack, Annexation Of Crimea By The Russian Federation, Atomic Bombings Of Hiroshima And Nagasaki, British Expedition To Tibet, Caesar’s Civil War, Cambodian-Vietnamese War, Finnish Civil War, Greek Civil War, History Of Xinjiang, Hungarian Revolution Of 1956, Indonesian Occupation Of East Timor, Iranian Revolution, Iran-Iraq War, Iraq War, Japanese Invasion Of Taiwan (1895), Korean Air Lines Flight 007, Lebanese Civil War, Libyan Civil War 2014-Present, Mukden Incident, Nepalese Civil War, Pacific War, Polish-Russian War Of 1792, Sino-Indian War, Sino-Vietnamese War, Somali Civil War, Soviet Invasion Of Poland, Spanish Civil War, Sri Lankan Civil War, Thallium Poisoning Case Of Zhu Ling, Turkish War Of Independence, War On Terror, Warsaw Pact Invasion Of Czechoslovakia, 2014 Gaza-Israel Conflict, Gaza War 2008-09
Dev (n=5)	2010 Hong Kong Electoral Reform, 2014 Pro-Russian Unrest In Ukraine, Korean War, Russo Japanese War, Slovak National Uprising
Test (n=10)	1963 Syrian Coup d’État, Air Raids On Japan, American Civil War, Angolan Civil War, Chinese Civil War, English Civil War, First Indochina War, Soviet-Afghan War, Warsaw Uprising (1794), Wukan Protests

Table 5: Train/Dev/Test Split of Topics

Component	Instruction
System Role	You are a computational linguistics expert assisting in a study of Adversarial Causal Attributions in Wikipedia articles. You will translate text from [SOURCE LANGUAGE] to English.
Objective	The translation must transparently preserve the original author’s causal logic, framing choices, and assignment of responsibility.
Causal Mapping	Every causal marker in the source text (connectors, resultative verbs, particles, or serial verb constructions) must have a 1:1 mapped equivalent in English. Do not paraphrase or restructure causality.
Agency Preservation	If the source uses passive voice to obscure or background an actor, retain the passive voice. If the source uses active or aggressive verbs to assign blame, do not soften or neutralize them.
Rhetorical Weight	Preserve the political or adversarial charge of terms. If the source uses a framing-loaded term (e.g., “martyred,” “liberated,” “betrayal”), select an English equivalent with the same framing intensity.
No Correction	Do not fix logical inconsistencies, balance viewpoints, or update facts to align with English Wikipedia. Faithfulness takes priority over accuracy or neutrality.
Language-Specific Instructions	
Chinese → English	Pay close attention to resultative verb compounds (e.g., , ,). Do not translate all as generic “caused.” Use “resulted in” (often negative), “gave rise to,” or “triggered” to match intensity. Preserve topic-comment structure and serial verb constructions when they imply causality without conjunctions. Maintain passive focus introduced by the <i>bei</i> () construction.
Arabic → English	Differentiate clause-based causality (<i>li-annahu</i>) from nominal causality (<i>bi-sababi</i>). If the particle <i>fa</i> implies immediate consequence, do not translate it as “and”; use “consequently,” “thus,” or “whereupon.” Preserve metaphorical or abstract agency (e.g., attributing actions to concepts or forces) without sanitization.
Output	[Translation]

Table 6: Prompt for Analytically Faithful Translation in Adversarial Causal Attribution

Annotation Guidelines for Framing-Aware Event Causality Identification (FrECI)

Task	Given a document and a set of <i>anchor event mentions</i> corresponding to real-world events, annotate all framed causal claims involving these anchors. A framed causal claim specifies not only that one event causes another, but also how responsibility is assigned and framed in the narrative.
Causal Claim	<p>A causal claim is a directed relation between two anchor events:</p> <ul style="list-style-type: none"> • Cause: an anchor event presented as causally prior. • Effect: an anchor event presented as the outcome. <p>Annotate if the text asserts or implies that the Cause brings about, explains, or leads to the Effect. Do not annotate purely temporal, correlational, or descriptive relations.</p> <p>Hybrid annotation: Review LLM-proposed causal relations, validate correct links, reject false positives, and add any missed relations.</p>
Responsibility Targets	<p>For each causal claim, identify all responsibility targets: actors, institutions, or groups portrayed as responsible for causing or enabling the effect.</p> <ul style="list-style-type: none"> • Use SRL-extracted agents as candidates; validate or revise. • If responsibility is diffuse or no agent is named, annotate <code>Incumbent</code>.
Framing Effect	<p>Assign a framing effect <i>independently</i> to each responsibility target:</p> <ul style="list-style-type: none"> • <code>Blame</code>: target is negatively responsible or condemned. • <code>Credit</code>: target is positively responsible or praised. • <code>Undermine_Credit</code>: positive contribution is questioned or minimized. • <code>Exonerate_Blame</code>: target is justified or absolved of responsibility. • <code>Framing_Neutral</code>: responsibility stated without evaluation. <p>When uncertain, default to <code>Framing_Neutral</code>.</p>
Source Alignment	<p>Assign one source label per causal claim (shared across all targets):</p> <ul style="list-style-type: none"> • <code>Author</code>: asserted in the narrative voice (default if no attribution cue). • <code>Target</code>: attributed to a responsibility target. • <code>Ally</code>: attributed to a source aligned with the target. • <code>Opponent</code>: attributed to a source opposing the target. • <code>Third_Party</code>: attributed to external observers or analysts.
Epistemic Modality	<p>Assign one modality per causal claim (shared across all targets):</p> <ul style="list-style-type: none"> • <code>Full_Affirmative</code>: unhedged assertion. • <code>Partial_Affirmative</code>: hedged support (e.g., <i>likely</i>, <i>appears to</i>). • <code>Neutral</code>: reporting stance without commitment. • <code>Partial_Negative</code>: qualified doubt. • <code>Full_Negative</code>: explicit rejection of causality.
Annotation Scope	<ul style="list-style-type: none"> • Claim-level: Source Alignment and Epistemic Modality apply to the entire causal claim. • Target-level: Framing Effects are assigned independently to each responsibility target. • A single causal claim may have multiple targets with different framing effects.
Edge Cases	<ul style="list-style-type: none"> • If multiple causal claims appear, annotate each separately. • If multiple sources are cited, choose the source most directly asserting the claim. • If no causal relation exists between anchor events, annotate nothing. • Follow the text's assertion, not real-world plausibility.

Table: Human annotation guidelines for FrECI. Annotators identify framed causal claims between anchor events, annotate responsibility targets with framing effects, and assign claim-level source alignment and epistemic modality.

Figure 2: Annotation Guidelines for FrECI.

▼ **Causal Relation #24** MAVEN

Relation Type: PRECONDITION

Event A (Cause): elected
[EVENT_d6021d607a092333f1dc75fdb00d8255]

[S0] The 1953 Iranian coup d'état, known in Iran as the 28 Mordad coup d'état (), was the overthrow of the democratically elected Prime Minister Mohammad Mosaddegh in favour of strengthening the monarchical rule of Mohammad Reza Pahlavi on 19 August 1953, orchestrated by the United States (under the name "TPAJAX" Project or "Operation Ajax") and the United Kingdom (under the name "Operation Boot").

[S1] It was the first covert action of the United States to overthrow a foreign government during peacetime.

Event B (Effect): overthrow
[EVENT_955eb3c7a9aa1dee937635358d5a34]

[S5] Initially, Britain mobilized its military to seize control of the British-built Abadan oil refinery, then the world's largest, but Prime Minister Clement Attlee opted instead to tighten the economic boycott while using Iranian agents to undermine Mosaddegh's government.

[S6] Judging Mosaddegh to be unreliable and fearing a Communist takeover in Iran, UK prime minister Winston Churchill and the Eisenhower administration decided to overthrow Iran's government, though the predecessor Truman administration had opposed a coup, fearing the precedent that Central Intelligence Agency (CIA) involvement would set.

[S7] British intelligence officials' conclusions and the UK government's solicitations were instrumental in initiating and planning the coup, despite the fact that the U.S. government in 1952 had been considering unilateral action (without UK support) to assist the Mosaddegh government.

Article Text

[0] The 1953 Iranian coup d'état, known in Iran as the 28 Mordad coup d'état (), was the overthrow of the democratically elected Prime Minister Mohammad Mosaddegh in favour of strengthening the monarchical rule of Mohammad Reza Pahlavi on 19 August 1953, orchestrated by the United States (under the name "TPAJAX" Project or "Operation Ajax") and the United Kingdom (under the name "Operation Boot").

[1] It was the first covert action of the United States to overthrow a foreign government during

Target Annotations

◀ **Causal Relation #24** MAVEN Delete Relation

Event A (Cause) - Targets + Add Cause Target

▼ **Target A0** Remove Target

Target Text: (actor/agent in the event)

Capture Selected Text e.g., 'government', 'protesters', 'ea'

1. Framing Effects:

- Select...
- Credit
- Blame
- Undermine Credit
- Exonerate Blame
- Framing Neutral
- Select Source Type...

3. Certainty/Modality:

Select...

Event B (Effect) - Targets + Add Effect Target

▼ **Target B0** Remove Target

Target Text: (actor/agent in the event)

Capture Selected Text e.g., 'government', 'protesters', 'ea'

1. Framing Effects:

Select...

2. Source Attribution:

Capture Selected Text Type or capture text...

Select Source Type...

3. Certainty/Modality:

Select...

+ Add New Causal Relation

Click to manually add a causal relation not annotated in MAVEN

Save Annotations

Figure 3: Web-based annotation interface used for FRECI. The interface presents gold anchor events, candidate causal relations, and structured annotation fields for responsibility targets, framing effects, source type, and epistemic modality. Claim-level attributes are shared across targets, while framing effects are assigned per target, enforcing the FRECI schema during annotation.

Government Response to Ojwang Protest

All Events

Coreference Clusters

Unique Events

Causal Relations

? Help

Media Source 1

Following the fatal custody incident involving blogger Albert Ojwang, authorities deployed tear gas and blocked major roads to contain escalating unrest. The state emphasized that these measures were critical for maintaining public order during volatile demonstrations.

Media Source 2

Despite the public outcry following Albert Ojwang's death in custody, security forces fired tear gas and enacted a citywide lockdown, undermining citizens' right to peaceful protest and intensifying national outrage, the reporter noted. The government's harsh response appeared disproportionate and politically driven.

Event Cluster Descriptors

Generated neutral descriptions of the framing-divergent event coreferential pairs

Albert Ojwang's death in custody deployment of tear gas road blocks and lockdown measures Ojwang Incident Protest

Attitudes

The attitude conveyed by this event mention toward the main event (Government Response to Ojwang Protest)

Supportive Skeptical Neutral

Media Source 1

Following the fatal custody incident involving blogger Albert Ojwang, authorities deployed tear gas and blocked major roads to contain escalating unrest. The state emphasized that these measures were critical for maintaining public order during volatile demonstrations.

Media Source 2

Despite the public outcry following Albert Ojwang's death in custody, security forces fired tear gas and enacted a citywide lockdown, undermining citizens' right to peaceful protest and intensifying national outrage, the reporter noted. The government's harsh response appeared disproportionate and politically driven.

Unique Context Events

Events unique to the article from Media Source 1 Events unique to the article from Media Source 2

Media Source 1

Following the fatal custody incident involving blogger Albert Ojwang, authorities deployed tear gas and blocked major roads to contain escalating unrest. The state emphasized that these measures were critical for maintaining public order during volatile demonstrations.

Media Source 2

Despite the public outcry following Albert Ojwang's death in custody, security forces fired tear gas and enacted a citywide lockdown, undermining citizens' right to peaceful protest and intensifying national outrage, the reporter noted. The government's harsh response appeared disproportionate and politically driven.

Causal Relations

Click to select • Ctrl/Cmd+Click to select multiple

Cause Effect Causal Relation

Western Media

Cause: 1_2: "deployed te... → Effect: 1_4: "contain escalatin...

Source: author

Certainty: positive Framing Type: intentional Framing Aware: deflect the blame

Cause: 1_3: "blocked majo... → Effect: 1_4: "contain escalatin...

Source: author

Certainty: positive Framing Type: intentional Framing Aware: deflect the blame

Russian State Media

Cause: 2_2: "fired tear... → Effect: 2_4: "undermining citizen...

Source: the reporter

Certainty: positive Framing Type: intentional Framing Aware: blame

Cause: 2_3: "enacted a city... → Effect: 2_4: "undermining ...

Source: the reporter

Certainty: positive Framing Type: intentional Framing Aware: blame

Cause: 2_2: "fired tea... → Effect: 2_6: "intensifying national ...

Source: the reporter

Certainty: positive Framing Type: intentional Framing Aware: blame

Figure 4: Visualization interface for analyzing framing-divergent causal explanations across aligned narratives. The interface highlights shared and unique events, causal relations, and responsibility attributions across sources, supporting qualitative analysis of divergent causal framing enabled by FRECI.

Component	Instruction
System Role	You are an expert in Causal Inference and Natural Language Processing, specialized in event extraction.
Task	Identify directed cause–effect relationships between event mentions drawn only from a provided event list, using the document context.
Definitions	Direct Causality: Event <i>A</i> explicitly triggers Event <i>B</i> (e.g., “The bombing caused the shutdown.”) Implicit / Long-range Causality: Event <i>A</i> creates necessary conditions, incentives, or motives for Event <i>B</i> , even when separated by paragraphs or mediated by discourse structure (e.g., a policy change early in the article leading to a protest later).
Input (User Template)	Document Context: [INSERT FULL WIKIPEDIA ARTICLE TEXT]
Output Format	Extracted Event Mentions: [ID: E1, Text: "..."] [ID: E2, Text: "..."] ... [ID: En, Text: "..."] Return a JSON list, where each item has the following schema: <pre>{ "cause_id": "EX", "effect_id": "EY", "evidence_quote": "Exact sentence(s) justifying the link." }</pre>

Table 7: Prompt for event-based causality extraction (ECI) from a document with a fixed event mention list.

Zero-shot Prompt for FRECI

Task.

You are given a news document and a list of anchor event mentions corresponding to real-world events of interest. Your task is to identify *framed causal claims* involving these anchor events.

A framed causal claim consists of:

- **cause**: an anchor event that causes another anchor event
- **effect**: an anchor event that is caused
- **source**: who expresses the causal claim, one of {Author, Target, Ally, Opponent, Third_Party}
- **modality**: certainty of the causal claim, one of {Full_Affirmative, Partial_Affirmative, Neutral, Partial_Negative, Full_Negative}
- **targets**: a list of responsibility targets, where each target has:
 - **target**: the actor held responsible
 - **framing**: one of {Credit, Blame, Undermine_Credit, Exonerate_Blame, Framing_Neutral}

Rules.

- Use only the provided anchor events as causes or effects.
- Source and modality apply to the entire causal claim and are shared across all targets.
- Framing effects may differ across targets.
- If no causal relation exists between anchor events, output an empty list.
- Output *valid JSON only*. Do not include explanations or extra text.

Document:

[DOCUMENT TEXT]

Anchor Events:

[ANCHOR EVENT LIST]

Output Format (JSON):

```
[
  {
    "cause": "...",
    "effect": "...",
    "source": "Author|Target|Ally|Opponent|Third_Party",
    "modality": "Full_Affirmative|Partial_Affirmative|Neutral|Partial_Negative|Full_Negative",
    "targets": [
      {"target": "...", "framing": "Credit|Blame|Undermine_Credit|Exonerate_Blame|Framing_Neutral"}
    ]
  }
]
```

Table 8: Zero-shot prompt used for the prompt-based LLM baseline in FRECI.

Chain-of-Thought Prompt for FRECI

Task.

You are given a news document and a list of anchor event mentions corresponding to real-world events of interest. Your task is to identify *framed causal claims* involving these anchor events.

A framed causal claim consists of:

- **cause**: an anchor event that causes another anchor event
- **effect**: an anchor event that is caused
- **source**: who expresses the causal claim, one of {Author, Target, Ally, Opponent, Third_Party}
- **modality**: certainty of the causal claim, one of {Full_Affirmative, Partial_Affirmative, Neutral, Partial_Negative, Full_Negative}
- **targets**: a list of responsibility targets, where each target has:
 - **target**: the actor held responsible
 - **framing**: one of {Credit, Blame, Undermine_Credit, Exonerate_Blame, Framing_Neutral}

Instructions.

Before producing the final answer, reason step by step about:

- whether any anchor event causally leads to another anchor event;
- who is held responsible for the outcome;
- how responsibility is framed for each target;
- who expresses the causal claim and how certain it is.

After completing the reasoning, output *only* the final structured result in valid JSON.

Rules.

- Use only the provided anchor events as causes or effects.
- Source and modality apply to the entire causal claim and are shared across all targets.
- Framing effects may differ across targets.
- If no causal relation exists between anchor events, output an empty list.
- Do not include reasoning steps or explanations in the final output.

Document:

[DOCUMENT TEXT]

Anchor Events:

[ANCHOR EVENT LIST]

Output Format (JSON):

```
[
  {
    "cause": "...",
    "effect": "...",
    "source": "Author|Target|Ally|Opponent|Third_Party",
    "modality": "Full_Affirmative|Partial_Affirmative|Neutral|Partial_Negative|Full_Negative",
    "targets": [
      {"target": "...", "framing": "Credit|Blame|Undermine_Credit|Exonerate_Blame|Framing_Neutral"}
    ]
  }
]
```

Table 9: Chain-of-thought prompt used for the prompt-based LLM baseline in FRECI. The model is instructed to perform internal reasoning before producing the final structured output, though only the final output is evaluated.