

FACTAPPEAL: Identifying Epistemic Factual Appeals in News Media

Guy Mor-Lan Tamir Sheafer Shaul R. Shenhav

Hebrew University of Jerusalem

{guy.mor, tamir.sheafer, shaul.shenhav}@mail.huji.ac.il

Abstract

How is a factual claim made credible? We propose the novel task of *Epistemic Appeal Identification*, which identifies whether and how factual statements have been anchored by external sources or evidence. To advance research on this task, we present FACTAPPEAL, a manually annotated dataset of 3,226 English-language news sentences. Unlike prior resources that focus solely on claim detection and verification, FACTAPPEAL identifies the nuanced epistemic structures and evidentiary basis underlying these claims and used to support them. FACTAPPEAL contains span-level annotations which identify factual statements and mentions of sources on which they rely. Moreover, the annotations include fine-grained characteristics of factual appeals such as the type of source (e.g. Active Participant, Witness, Expert, Direct Evidence), whether it is mentioned by name, mentions of the source’s role and epistemic credentials, attribution to the source via direct or indirect quotation, and other features. We model the task with a range of encoder models and generative decoder models in the 2B–9B parameter range. Our best performing model, based on Gemma 2 9B, achieves a macro- F_1 score of 0.73.¹

1 Introduction

In an era marked by pervasive misinformation and heightened skepticism of media reporting, understanding how factual claims are presented has become more important than ever. While research in claim detection and verification has made substantial progress (Saurí and Pustejovsky, 2009; Thorne et al., 2018; Hassan et al., 2017; Wadden et al., 2020; Aly et al., 2021), most existing methods focus on the content of the statements in isolation and overlook the epistemic structures that confer credibility and persuasive force to these claims. In news

media, for example, the credibility of a claim is not only determined by its content but also by the way it appeals to external sources of knowledge—be it through expert testimony, official statements, or direct empirical evidence. Understanding how factual claims are anchored by appeals to external sources is also important for broader tasks in discourse analysis, fact-checking, and the study of knowledge flows in the media.

To address this gap, we introduce FACTAPPEAL, a novel dataset designed to address the dual challenge of detecting both factuality and epistemic appeals within news statements. This task not only identifies whether a statement conveys a factual claim (i.e. a claim about the state of the external world) but also captures the underlying structure of how such claims are supported by sources such as experts, witnesses, and reports.

1.1 Epistemic Appeal Identification

An epistemic appeal is a factual claim supported by a reference to an authoritative source—whether genuinely authoritative or only purported to be—thereby providing a reason to accept the claim as true and enhancing its credibility. Epistemic appeals play a pivotal role in shaping how factual claims are constructed and perceived in public discourse, especially within news media. They are significant for analyzing epistemic justification structures for automatic fact verification, discourse analysis, and analyses of the social sources and dynamics of knowledge.

We propose the task of *Epistemic Appeal Identification*, which requires determining whether a sentence presents a factual claim and, if so, identifying how it invokes an external source or evidence to support that claim. This task requires identifying the source of epistemic authority, as well as classifying the type and method of appeal. This new task pushes the boundaries of traditional factuality detection by introducing a rich layer of epistemic

¹Data available at github.com/guymorlan/factappeal/, cc-by-4.0 license.

reasoning, crucial for understanding how information is conveyed and validated in public discourse.

FACTAPPEAL comprises 3,226 sentences from news articles manually annotated with fine-grained span-level annotations. We label factual statements (whether true or false) in each sentence, as well as any epistemic appeals that ground the statement to a source of epistemic authority. We additionally include fine-grained annotations for features of epistemic factual appeals, such as the type of source invoked (e.g. Expert, Witness, Active Participant), the method of appeal (direct or indirect quotation) and more. Our annotations span a wide range of appeal types, such as official statements, reports, and testimonies, offering fine-grained insights into how claims are constructed and backed by different types of epistemic authority.

2 Related Work

Understanding how factual claims are supported has been the focus of several research strands in natural language processing. In this section, we review the literature on claim detection and verification, epistemic modality and argumentation, and source attribution. We then explain how our work extends these efforts by jointly modeling factuality and detailed epistemic appeals.

2.1 Verifiability and Claim Verification

Early work on factuality detection aimed at determining whether statements describe verifiable events (Saurí and Pustejovsky, 2009, 2012; Hassan et al., 2017). More recent lines of research have emphasized claim verification, exemplified by large-scale benchmarks such as FEVER (Thorne et al., 2018), which require systems to determine if a claim is *supported* or *refuted* based on evidence. Other datasets have focused on specific domains or subtasks, such as SciFact (Wadden et al., 2020) for verifying scientific claims or FactRel (Mor-Lan and Levi, 2024) for factual entailment in news. While these resources have substantially advanced fact-checking methods, they focus primarily on detecting claims and modeling relations between claims, rather than providing a complete epistemic schema describing how a claim itself is constructed and supported.

2.2 Epistemic Modality and Argumentation

Research on epistemic modality seeks to capture linguistic markers of certainty and belief (Rubin,

2010; Saurí and Pustejovsky, 2012), while argumentation mining explores how claims are constructed and supported within discourse (Feng and Hirst, 2011). A related task is epistemic stance detection, which models a source’s degree of commitment—such as certainty or doubt—toward a claim (Gupta et al., 2022). Such analyses are crucial in persuasive language, where models have been trained to identify broad strategies like appeals to authority (Da San Martino et al., 2019). While these works focus on belief, commitment, or general persuasive tactics, our approach offers a more granular view. By contrast, we pinpoint the *concrete sources* invoked (e.g., a named expert, a witness) and classify the structural nature of the appeal itself, such as the source type and method of quotation. This allows for more precise modeling of how claims receive or signal credibility.

2.3 Source Attribution and Quotation Analysis

Prior studies have addressed the task of detecting quotations and attributing them to entities (Pareti et al., 2013; O’Keefe et al., 2012), which is crucial for scientific, journalistic and legal texts. However, these methods do not typically classify sources by *type* (e.g., expert vs. witness) or capture whether appeals are invoked through direct speech or paraphrasing. We build on these works by jointly modeling factuality and source-based epistemic appeals, thereby revealing how news articles invoke or display a source’s authority to support a factual claim.

3 Annotation Scheme

3.1 Overview

We propose a span-level annotation scheme for detecting epistemic appeals in news media, labeling each relevant textual span alongside its associated features. The tags are provided both as character indices and as XML-style tags. Span-level tags are a key advantage of FACTAPPEAL, allowing differentiating factual appeals, facts without appeals and non-factual components in a single text, as well as identifying multiple epistemic sources. Tags of different types may also be nested. The tags are:

- **Fact Without Appeal** — factual claim made without epistemic appeal to a source.
- **Fact With Appeal** — factual claim made with an epistemic appeal to a source. This tag has one modifier, an additional tag for whether

the identified fact reproduces the source’s speech verbatim or paraphrases and processes it. It is always annotated with respect to Fact With Appeal spans, with two possible values:

- Direct quote
 - Indirect quote
- **Source** — epistemic source to which a claim is attributed. This tag has two additional modifiers annotated with respect to all identified source spans.

First, whether the source is mentioned by name or not:

- Named
- Unnamed

Second, the type of epistemic source:

- Active Participant
 - Witness
 - Direct Evidence
 - Official
 - Expert
 - Expert Document
 - News Report
 - null (cannot be determined)
- **Source Attribute** — marking relevant epistemic attributes of the sources, such as a title, office or status held by the epistemic source, or any information about the source cited as epistemic credentials.
 - **Recipient** — recipient receiving the information from the appeal source.
 - **Appeal Time** — time in which appeal was made.
 - **Appeal Location** — physical, virtual or symbolic location in which appeal was made.

The primary tags are further explained below.

3.2 Factual Claims

We first examine the factuality of a sentence. **Factual claims** are sentences that primarily make a statement about the state of the external world, which could be either true or false. They correspond to what Jakobson describes as the referential function of language, which is concerned with conveying information about the external world and

is “oriented toward the context” (Jakobson, 1960), as well as to the assertive speech act described by Searle, in which the speaker commits to the truth of what is asserted (Searle, 2013). Thus, statements that primarily convey a personal experience or subjective feeling are non-factual, and receive a null annotation:

- (1) **“Even so, when I visited Chennai, I felt okay about the media future we’re heading into.”**

Note that the use of quotation marks does not necessitate that a cited statement is an epistemic appeal or even factual, as these categorizations depend on the dominant function of the statement.

Normative statements that primarily express a value judgment are considered non-factual within this annotation scheme:

- (2) **They shouldn’t have had anything to do with this investigation, with this case, any submission to the FISA court.**

Similarly, questions, pleas, commands, calls to action and similar speech acts fall outside the scope of factual statements:

- (3)
 - a. **What exactly are you going to do?**
 - b. **Add your name to millions demanding Congress take action on the President’s crimes.**

Factual appeals are factual claims accompanied by a reference to a purported source of knowledge. Appeals are generally performed via some form of reference or citation,² which could take the form of direct quotation reproducing speech verbatim, or indirect reference including any forms of paraphrasing or knowledge mediation.

Thus, a **brute factual statement** is a factual claim that lacks any epistemic appeal, and is annotated as follows:

- (4) **<Fact_No_Appeal> Sometimes called street cameras, the portable P.D.Q. (Photography Done Quickly) model**

²Including unattributed quotes, in which the existence of a source is implied but its identity is not determined.

could produce pocket-size photographs directly onto paper, eliminating the need for negatives. </Fact_No_Appeal>

A challenging aspect of FACTAPPEAL is distinguishing cases where an entity is mentioned merely as the subject of a report from instances where the source is cited to bolster a factual claim through its authority. For example:

- (5) <Fact_No_Appeal> **After the successful test hop, Mr Musk said: “One day Starship will land on the rusty sands of Mars.”** </Fact_No_Appeal>

Here, although Elon Musk is quoted, his authority is not invoked as evidence for a verifiable fact; instead, the statement primarily *reports on* Musk making this comment. Consequently, this is annotated as Fact_No_Appeal rather than an epistemic appeal.

3.3 Types of Epistemic Appeals

We develop a structured typology of appeal sources grounded in the nature of the evidence that supports each factual claim. This framework is essential for distinguishing among various forms of authority and for clarifying how these authorities function within epistemic appeals.

As shown in Figure 1, our typology classifies sources according to two fundamental dimensions: (1) *proximity to the event (internal vs. external)* and (2) *whether the source is human or non-human*. An internal source has a direct, firsthand connection to the event, whereas an external source provides more generalized expertise. **Internal appeals** thus involve a factual grounding via an epistemic source with immediate or sensory contact to the events. They comprise the following types:

Active participants are actors taking active roles in the events related to the fact.

- (6) <Source:Named:Participant> **Emily** </Source> **told the** <Recipient> **Buffalo News** </Recipient> <Fact_Appeal:Indirect> **she had received a text from her mother that read: “Well, I am done with you.”** </Fact_Appeal>

Witnesses are observers who provide firsthand testimony of events but are not active participants.

- (7) **Another** <Source_Attribute> **witness to the shooting,** </Source_Attribute> <Source:Named:Witness> **Megan Chadwick,** </Source> **said** <Fact_Appeal:Indirect> **her husband saw the civilian take down the shooter.** </Fact_Appeal>

Officials are active participants who also have extra non-epistemic authority on events or on facts — e.g., legal, political, bureaucratic authority. Officials, such as government authorities, often provide statements that carry legal or formal weight. Importantly, officials wield power that can alter states of affairs related to the factual claim.

- (8) <Source:Named:Official> **Doug Erickson,** </Source> <Source_Attribute> **the EPA’s communications director for the transition,** </Source_Attribute> **told** <Recipient> **National Public Radio** </Recipient> **that** <Fact_Appeal:Direct> **“we’ll take a look at what’s happening so that the voice coming from the EPA is one that’s going to reflect the new administration.”** </Fact_Appeal>

Direct evidence is an appeal to a piece of evidence found “at the scene” and bearing on the facts related to the factual claim.

- (9) <Source:Unnamed:Direct_Evidence> **This 2013 photo** </Appeal_Source> **provided to** <Recipient> **The Associated Press** </Recipient> **shows** <Fact_Appeal:Indirect> **now-defrocked Catholic priest Richard Daschbach leading a service at a church in Kutet, East Timor.** </Fact_Appeal>

External appeals on the other hand involve appeals to a source without a firsthand connection to events, whose epistemic credentials are grounded in general expertise. These sources possess epistemic expertise which bears on the factual claim:

Experts, such as scientists or specialists, offer appeals rooted in professional expertise and specialized knowledge.

- (10) <Fact_Appeal:Direct> **“The dolphins of Sarasota Bay are really good indicators of the health of our ecosystem,”** </Fact_Appeal> **said** <Source:Named:Expert> **Dr. Wells.** </Source>

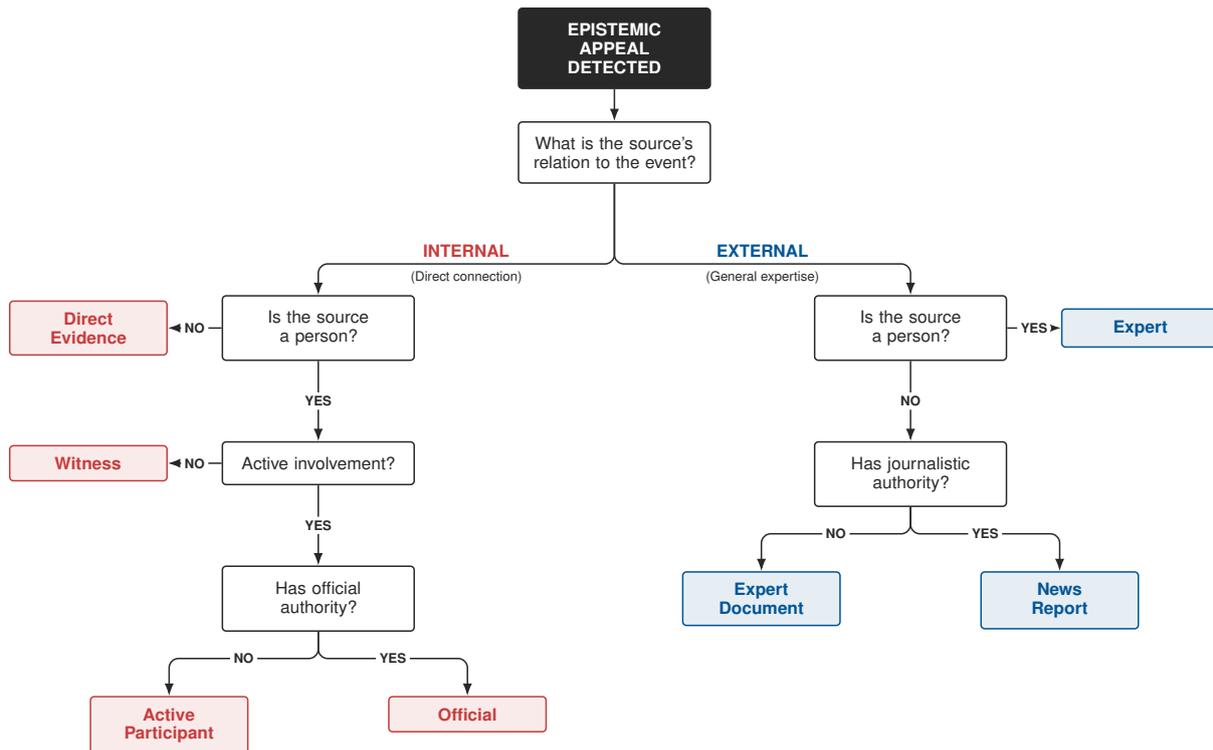


Figure 1: Typology of Epistemic Appeal Sources

Expert Document refers to expert knowledge embodied in non-human objects, such as research documents, scientific and institutional reports.

- (11) A 2013 <Source:Unnamed:Expert_Doc> study </Source> found that <Fact_Appeal:Indirect> peppermint oil has potent antiseptic properties which are useful against oral pathogens. </Fact_Appeal>

News Report refers to citation of previous news reports.

- (12) <Fact_Appeal:Direct> "We are praying to God and asking that all Argentines help us to pray that they keep navigating and that they can be found," </Fact_Appeal> <Appeal_Source:Named:Witness> Claudio Rodriguez, </Appeal_Source> <Source_Attribute> the brother of one of the crew members </Source_Attribute> told <Recipient> the local Todo Noticias TV channel, </Recipient> according to <Appeal_Source:Named:News_Report> the AP. </Appeal_Source>

The distinction between internal and external sources also reflects two modes or logics of epistemology—a common-wisdom logic preferring those with direct relations to the matter at hand, as opposed to an expertise-based logic preferring “detached” experts. Whereas internal sources have epistemic credentials in virtue of their specific history and contact with the situation at hand, external sources possess epistemic credibility due to their attained expertise (Pierson, 1994; Collins and Evans, 2002).

4 Dataset

The dataset contains 3,226 sentences sampled from diverse English-language news articles published between 2020 and 2022. Each sentence was annotated by one of two annotators: one of the authors and a student research assistant (see appendices A and B). The dataset has been randomly split into a training set (70%), development set (15%) and test set (15%).

4.1 Inter-Annotator Agreement Analysis

We conducted an inter-annotator agreement (IAA) analysis on a subset of data labeled by both annotators. To facilitate the comparison, each span annotation was converted into binary word-level labels. Using these labels, we computed several met-

rics—namely the union and intersection counts, the number of words where neither annotator marked the tag, the Intersection over Union (IoU), and Cohen’s Kappa. Table 1 summarizes the IAA statistics for each tag. The overall IoU of 0.74 and a Cohen’s Kappa of 0.82 indicate substantial agreement between the annotators. However, some span annotations are relatively rare and have few instances.

Tag	Union	Intersection	Unlabeled	IoU	Cohen’s κ
Fact w/o Appeal	511	372	1194	0.73	0.79
Fact with Appeal	986	732	719	0.74	0.70
Appeal Time	15	11	1690	0.73	0.85
Appeal Location	27	17	1678	0.63	0.77
Recipient	14	14	1691	1.00	1.00
Source	131	104	1574	0.79	0.88
Source Attribute	90	83	1615	0.92	0.96
Indirect Quote	669	420	1036	0.63	0.67
Direct Quote	368	261	1337	0.71	0.79
Active Participant	21	12	1684	0.57	0.73
Witness	22	19	1683	0.86	0.93
Direct Evidence	20	13	1685	0.65	0.79
Official	23	15	1682	0.65	0.79
Expert	14	12	1691	0.86	0.92
External Document	27	15	1678	0.56	0.71

Table 1: Word-level Inter-annotator Agreement Metrics³

4.2 Descriptive Statistics

We examine the share of sentences containing any factual claim in Figure 2. More than 80% of statements are annotated as factual. While this may seem high, it corresponds well to the factual transmitting nature of news reports.

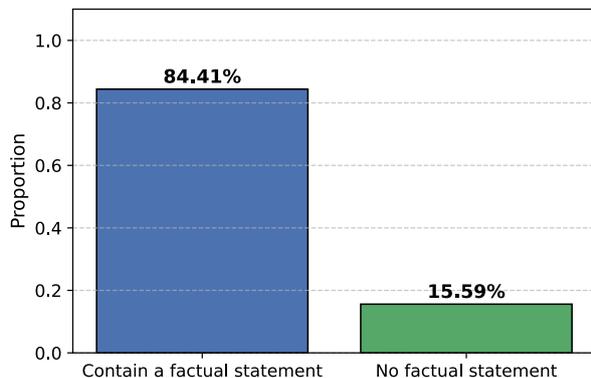


Figure 2: Proportion of Factual Sentences

In Figure 3, we present the distribution of span annotations. We first observe that statements without epistemic appeals appear nearly twice as frequently as those containing appeals. Moreover, we observe that most factual appeals utilize paraphrasing (66%) rather than direct quotation (34%).

³Named/Unnamed are excluded as they were added later. External Document includes both Expert Document and News Report, which were later split from this category.

When an appeal source is mentioned, it is usually mentioned by name (64%). For named sources, the most popular types are active participants (20%), news reports and expert documents (20%), officials (19%) and experts (19%), and are thus almost equally prevalent. Witnesses and direct evidence account for a smaller share and thus appear substantially less common as sources of knowledge.

For appeal sources that are unnamed (35%), news reports and expert documents are most common (24%), followed by the null category for indeterminate types (19%) and officials (17%). Witnesses, experts, active participants and direct evidence thus appear less frequently as unnamed sources.

5 Experiments

We compare two modeling strategies for Epistemic Appeal Identification (see Appendix C):

Token-level multi-label classification with encoder models. Since different tag types may overlap, in this setting we represent the tags as token-level multi-label binary annotations, with 18 labels corresponding to each of the tags and possible modifier values. We fine-tune pre-trained transformer encoder models, using the base model versions of RoBERTa (Liu et al., 2019), DeBERTa v3 (He et al., 2021) and ModernBERT (Warner et al., 2024). The encoder models are trained for up to 12 epochs with focal loss (Lin et al., 2017).

Generative decoder models. In this setting, annotations are represented as XML-style tags (similar to the presentation in Section 3). Models are trained to produce the annotated sentence given the raw sentence. We fine-tune several smaller pre-trained LLMs such as Gemma 2 (2B and 9B) (Gemma Team et al., 2024), Llama 3.1 8B (Grattafiori et al., 2024) and Mistral v0.3 7B (Jiang et al., 2023). The models are trained with QLORA (Detmeters et al., 2023) with 4-bit quantization for 3 epochs, with $r = 256$ and $\alpha = 256$. We mask the loss of the input prompt and train on completions.

In order to allow for partial matches and robustness across tokenizers, we calculate word-level scores for each of the 18 tag categories. Table 2 reports macro-average word-level precision, recall and F_1 scores on the test set. While differences between top encoder models and top decoder models are small, the largest decoder, Gemma 2 9B, achieves the best macro- F_1 score of 0.73.

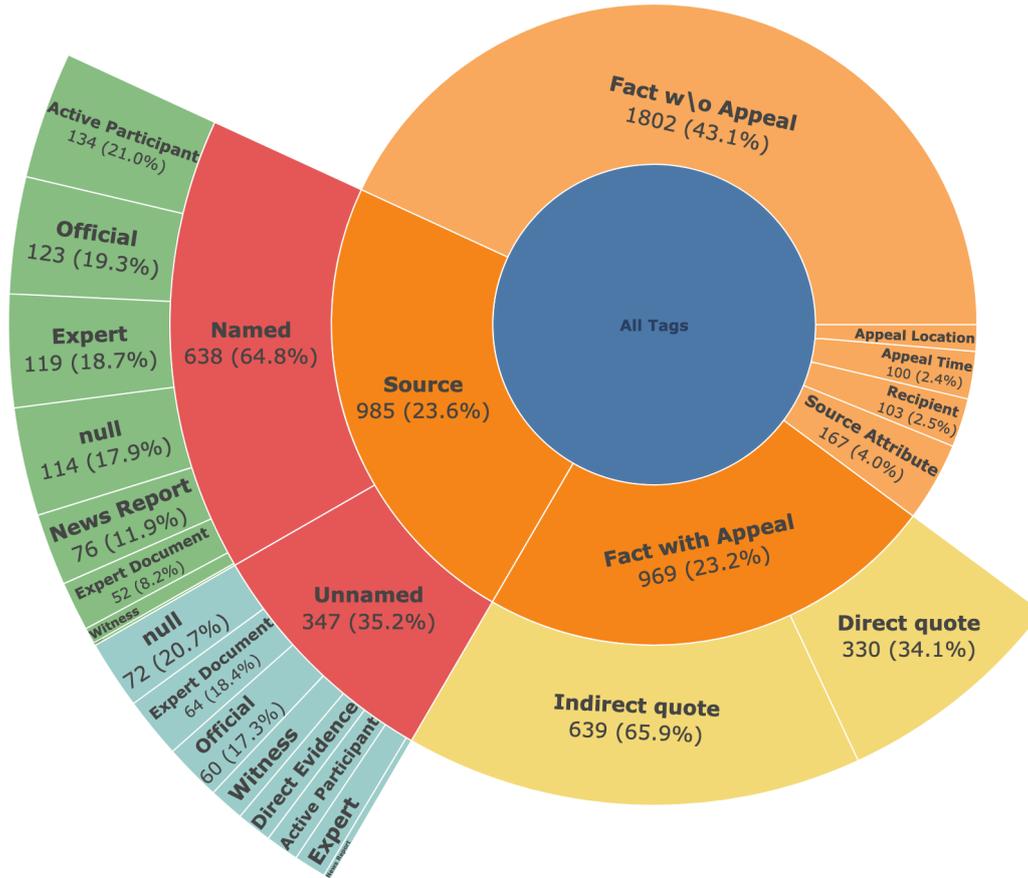


Figure 3: Distribution of Span Tags⁴

In Table 3 we take a look at the F_1 scores per tag in the test set. While some tags are learned well by encoder models, encoder models show higher variation in performance across categories. Moreover, encoder models show a stronger correlation between tag counts and test F_1 scores ($\rho_s = 0.72$) than decoders ($\rho_s = 0.66$).

For source type annotations specifically, performance is less correlated with the prevalence of the tags, as the four more prevalent tags *Active Participant*, *Expert Document*, *Official* and *Expert* are not necessarily better detected than the less prevalent *Direct Evidence* and *Witness*. Here again, tag prevalence is more strongly correlated with F_1 scores for encoder models ($\rho_s = 0.18$) and less so for decoders ($\rho_s = 0.07$).

Overall, these results indicate that *Epistemic Appeal Identification* remains challenging for encoder-only models and smaller LLMs, highlighting significant room for improvement.

Model	Precision	Recall	F_1
RoBERTa (base)	0.75	0.67	0.7
DeBERTa v3 (base)	0.73	0.67	0.69
ModernBERT (base)	0.73	0.47	0.54
Gemma 2 9B	0.76	0.73	0.73
Mistral v0.3 7B	0.73	0.68	0.7
Llama 3.1 8B	0.75	0.65	0.68
Gemma 2 2B	0.65	0.58	0.6

Table 2: Global Macro Metrics, Test Set

6 Conclusion

In this work, we introduce FACTAPPEAL, a novel dataset and task formulation aimed at identifying epistemic appeals in news media factual claims. Our dataset captures both the factuality of claims and the underlying epistemic structures that lend these claims credibility. The experiments comparing token-level predictions using encoder models with generative LLMs underscore the challenges of modeling nuanced epistemic appeals, as well as the strength of generative models and the feasibility of

Tag	RoBERTa	ModernBERT	DeBERTa v3	Gemma 2 2B	Llama 3.1 8B	Mistral v0.3 7B	Gemma 2 9B
Factuality							
Fact w/o Appeal	0.87	0.83	0.86	0.84	0.86	0.88	0.89
Fact with Appeal	0.84	0.84	0.86	0.78	0.83	0.84	0.85
Appeal Characteristics							
Appeal Time	0.75	0.69	0.65	0.44	0.60	0.68	0.77
Appeal Location	0.47	0.43	0.49	0.31	0.57	0.47	0.58
Recipient	0.77	0.51	0.90	0.62	0.71	0.81	0.89
Source	0.81	0.77	0.84	0.70	0.77	0.79	0.84
Source Attribute	0.71	0.53	0.80	0.65	0.66	0.58	0.79
Quotation Type							
Indirect Quote	0.83	0.83	0.87	0.77	0.82	0.81	0.83
Direct Quote	0.77	0.79	0.79	0.72	0.76	0.78	0.80
Source Named							
Named	0.75	0.76	0.80	0.68	0.71	0.72	0.77
Unnamed	0.69	0.53	0.67	0.49	0.69	0.64	0.65
Source Type							
Active Participant	0.43	0.16	0.26	0.40	0.45	0.44	0.54
Witness	0.66	0.10	0.55	0.57	0.57	0.81	0.57
Direct Evidence	0.58	0.19	0.51	0.30	0.73	0.57	0.71
Official	0.67	0.51	0.62	0.59	0.59	0.59	0.71
Expert	0.60	0.17	0.57	0.52	0.65	0.63	0.62
Expert Document	0.61	0.43	0.62	0.63	0.67	0.73	0.68
News Report	0.76	0.68	0.75	0.79	0.68	0.78	0.68
<i>Standard Deviation</i>	0.12	0.25	0.17	0.16	0.10	0.13	0.11

Table 3: Per-Tag F_1 Scores

sequence-to-sequence representations for this task.

Beyond advancing the modeling of epistemic appeals, this work also contributes to the fields of factual detection and automated fact-checking, offering span-level annotations that capture how claims are justified in news media. By providing fine-grained annotations—differentiating factual from non-factual statements and detailing the types of epistemic appeals—our approach opens new avenues for more context-aware fact-checking. This dual focus on both factuality and the structure of supporting evidence addresses key limitations in current factuality detection frameworks and paves the way for more robust news factuality analysis.

Furthermore, FACTAPPEAL has important implications for social science research across political philosophy, social epistemology, and communication. Scholars such as Anderson (Anderson, 2021) and Lynch (Lynch, 2021) have highlighted that contrasting epistemic frameworks can lead to “deep disagreements” among political groups, and communication scholars have underscored the central role of media in shaping which facts gain prominence and how audiences interpret them (McCombs and Shaw, 1972; Entman, 1993). More recent studies demonstrate how the information environment influences factual beliefs, partisan divides, and public polarization (Jerit and Barabas, 2012; Aalberg et al., 2012; Garrett et al., 2016;

Djerf-Pierre and Shehata, 2017). By systematically identifying and modeling epistemic appeals, FACTAPPEAL offers a powerful tool for investigating how news media construct and validate factual claims—a process fundamental to understanding broader social dynamics and shifts in political discourse.

Future research can leverage these contributions in several ways. In factuality and fact-checking, our dataset may improve claim verification and evidence detection approaches by incorporating source-based credibility cues. Extending FACTAPPEAL to larger textual units, such as paragraphs or entire articles, could reveal more complex discourse structures and further enhance automated verification. In computational discourse analysis, FACTAPPEAL can facilitate deeper investigations of epistemic appeals in public discourse, shedding light on broader patterns of justification, knowledge transfer, and media polarization.

Secondly, the community can utilize FACTAPPEAL to refine factual appeal modeling even further—by exploring appeals in larger contexts, linking multiple sources and claims, and identifying additional attributes of factual epistemic appeals. Future expansions could also include social media content and other distinct types of discourse.

Limitations

While FACTAPPEAL marks an important step forward in capturing epistemic structures in news media, our work has several limitations. First, the dataset employs only sentence-level annotations, which restricts the amount of contextual information that can be captured. Future studies might extend annotations to paragraphs or entire articles, where relationships among claims, sources, and evidence can be modeled more comprehensively.

Second, although multiple sources or factual claims can appear in a single sentence, the current annotations do not explicitly link each source to its corresponding claim. Such explicit linkage could improve the granularity of epistemic appeal analyses and enable more precise modeling of how diverse sources relate to one or more claims within the same sentence.

Finally, FACTAPPEAL comprises English-language news articles from a particular time frame (2020–2022). This narrow focus may limit the generalizability of our findings to other languages, domains, or historical periods. Future research could address these limitations by applying the annotation scheme to broader contexts and by leveraging multilingual corpora.

Acknowledgments

This work was supported by the Israel Science Foundation (Grant no. 2501/22). We thank Hagar Kaminer for diligent research assistance, the members of the DeepStory lab, and especially Effi Levi for insightful comments and suggestions that improved the final version of this manuscript.

References

- Toril Aalberg, Peter Van Aelst, and James Curran. 2012. Media systems and the political information environment: A cross-national comparison. In *How Media Inform Democracy*, pages 33–49. Routledge.
- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. The fact extraction and verification over unstructured and structured information (feverous) shared task. In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 1–13.
- Elizabeth Anderson. 2021. Epistemic bubbles and authoritarian politics. In Elizabeth Edenberg and Michael Hannon, editors, *Political Epistemology*, pages 11–30. Oxford University Press.
- Harry M Collins and Robert Evans. 2002. The third wave of science studies: Studies of expertise and experience. *Social studies of science*, 32(2):235–296.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news articles. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5636–5646, Hong Kong, China. Association for Computational Linguistics.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. *Qlora: Efficient finetuning of quantized llms*.
- Monika Djerf-Pierre and Adam Shehata. 2017. Still an agenda setter: Traditional news media and public opinion during the transition from low to high choice media environments. *Journal of Communication*, 67(5):733–757.
- Robert M Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of communication*, 43(4):51–58.
- Vanessa Wei Feng and Graeme Hirst. 2011. Classifying arguments by scheme. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1*, pages 987–996.
- R Kelly Garrett, Brian E Weeks, and Rachel L Neo. 2016. Driving a wedge between evidence and beliefs: How online ideological news exposure promotes political misperceptions. *Journal of Computer-Mediated Communication*, 21(5):331–348.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Ankita Gupta, Su Lin Blodgett, Justin Gross, and Brendan O’connor. 2022. *Examining political rhetoric with epistemic stance detection*. In *Proceedings of the Fifth Workshop on Natural Language Processing and Computational Social Science (NLP+CSS)*, pages 89–104, Abu Dhabi, UAE. Association for Computational Linguistics.
- Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by claim-buster. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1803–1812.

- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#).
- Roman Jakobson. 1960. Closing statement: Linguistics and poetics. *Style in language*, pages 350–377.
- Jennifer Jerit and Jason Barabas. 2012. Partisan perceptual bias and the information environment. *The Journal of Politics*, 74(3):672–684.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#).
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Doll  r. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Michael P. Lynch. 2021. Political disagreement, arrogance, and the pursuit of truth. In Elizabeth Edenberg and Michael Hannon, editors, *Political Epistemology*, pages 244–258. Oxford University Press.
- Maxwell E McCombs and Donald L Shaw. 1972. The agenda-setting function of mass media. *Public opinion quarterly*, 36(2):176–187.
- Guy Mor-Lan and Effi Levi. 2024. [Exploring factual entailment with NLI: A news media study](#). In *Proceedings of the 13th Joint Conference on Lexical and Computational Semantics (*SEM 2024)*, pages 190–199, Mexico City, Mexico. Association for Computational Linguistics.
- Timothy O’Keefe, Silvia Pareti, James R. Curran, Irena Koprinska, and Matthew Honnibal. 2012. [A sequence labelling approach to quote attribution](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 790–799, Jeju Island, Korea. Association for Computational Linguistics.
- Silvia Pareti, Tim O’Keefe, Ioannis Konstas, James R. Curran, and Irena Koprinska. 2013. [Automatically detecting and attributing indirect quotations](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 989–999, Seattle, Washington, USA. Association for Computational Linguistics.
- Robert Pierson. 1994. The epistemic authority of expertise. In *PSA: Proceedings of the Biennial meeting of the Philosophy of Science Association*, volume 1994, pages 398–405. Cambridge University Press.
- Victoria L Rubin. 2010. Epistemic modality: From uncertainty to certainty in the context of information seeking as interactions with text. *Information Processing & Management*, 46(5):533–540.
- Roser Saur   and James Pustejovsky. 2009. Factbank: a corpus annotated with event factuality. *Language resources and evaluation*, 43(3):227–268.
- Roser Saur   and James Pustejovsky. 2012. [Are you sure that this happened? assessing the factuality degree of events in text](#). *Computational Linguistics*, 38(2):261–299.
- John R Searle. 2013. Illocutionary acts and the concept of truth. In *Truth and Speech Acts*, pages 31–40. Routledge.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. In *Proceedings of the 2018 conference of the north american chapter of the association for computational linguistics: Human language technologies, volume 1 (long papers)*, pages 809–819.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or fiction: Verifying scientific claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics.
- Benjamin Warner, Antoine Chaffin, Benjamin Clavi  , Orion Weller, Oskar Hallstr  m, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, et al. 2024. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *arXiv preprint arXiv:2412.13663*.

A Annotators

The dataset has been annotated by two annotators, one of the authors and a paid student research assistant. The annotators are a man and a woman in their 20s–30s from the EMEA region.

B Annotation Guidelines

These guidelines describe what constitutes a factual statement, how to detect whether it appeals to an external source, and how to label the source and its attributes. They also detail how to mark the relevant spans in the text.

B.1 Determining Factuality

Definition. A sentence is **factual** if it primarily makes a statement about the external world that can be objectively true or false. Statements focusing on

subjective feelings, judgments, calls to action, or questions generally do **not** count as factual for this annotation scheme.

Label.

- Fact_No_Appeal (“Fact Without Appeal”) for factual statements that do not cite an external source.
- Fact_Appeal (“Fact With Appeal”) for factual statements that explicitly reference an external source or evidence to support their claim.

Non-Factual Content. If a sentence is *primarily* non-factual (for instance, it is dominated by a personal opinion or call to action), it receives no fact-related annotation.

B.2 Identifying Epistemic Appeals

Definition. An **epistemic appeal** is a factual claim that is accompanied by a reference to an external source or evidence. The reference can be direct (quoted verbatim) or indirect (paraphrased or summarized).

Distinguishing Reporting from Appeals. When a statement merely *covers* someone’s words or remarks without using the speaker’s position or information as evidence for a factual claim, it is annotated as Fact_No_Appeal. By contrast, if the statement explicitly *invokes* external authority or specialized knowledge as the reason to accept the factual claim, it is Fact_Appeal.

Method of Appeal. For each Fact_Appeal span, annotate the manner in which the claim references its source:

- Direct (quoted verbatim)
- Indirect (paraphrased or mediated)

B.3 Source Annotations

When annotating a Fact_Appeal span, identify the Source span(s) explicitly referenced in that statement. The Source tag has two modifiers:

Source Name.

- Named: The text gives a proper name or explicit identity.
- Unnamed: The source is referenced, but not by name (e.g., “an official stated...”).

Source Type. Each source is labeled with one of the following:

- Active_Participant: Has a direct, primary role in the events in question.

- Witness: Observed the events but was not directly involved.
- Direct_Evidence: A non-human piece of evidence (e.g., footage, photograph) closely tied to the scene.
- Official: Holds a position of non-epistemic authority (legal, governmental, etc.).
- Expert: A person with specialized knowledge not derived from direct involvement (e.g., scientist, analyst).
- Expert_Document: A written or recorded source of expertise (e.g., a published paper).
- News_Report: A journalistic source.
- null: Source type cannot be determined.

B.4 Additional Attributes

If relevant information is present, you may also label the following:

- **Source_Attribute:** Any text specifying the authority, rank, credentials, or role of the source (e.g., an official title).
- **Recipient:** The entity or individual to whom the source directed the claim (if explicitly stated).
- **Appeal_Time:** The time when the appeal was made (if explicitly mentioned, e.g., “yesterday” or a date).
- **Appeal_Location:** The physical, virtual, or symbolic location (e.g., “during a press briefing at the White House”).

B.5 Marking the Spans

All annotations should be represented at the span level. Spans can overlap or nest. For instance, a Fact_Appeal span could contain one or more Source sub-spans. Make sure each factual statement is fully wrapped, and all relevant sources or attributes within it are separately tagged.

B.6 Edge Cases and Practical Tips

Multiple Sources or Claims. A single sentence may present more than one claim or more than one source. Tag each factual claim with or without appeal separately. If a sentence has multiple appeals or different source types, annotate each source individually.

Attribution Without Clear Source Type. If the text provides insufficient detail to determine the source type (e.g., just “sources say...” with no additional context), use the null label for Source_Type.

Unclear Factuality or Mixed Content. If the sentence intermixes factual and non-factual statements, identify which portion is factual, provided it constitutes a coherent factual claim. Non-factual segments do not receive tags.

C Experimental Setup

Models were trained on a single A100 GPU with 40GB VRAM, with the longest model run taking 4 hours to complete. A learning rate of $1e-5$ was used. The results in the paper correspond to a single run.

For decoder models, we used a batch size of 6. The AdamW optimizer was used with the default Huggingface settings.

We used the following prompt template for training and inference: {input_sentence}\n### Answer: {annotated_sentence}

The input sentence tokens were masked from loss calculation.

Table 4 documents the number of parameters in the models utilized in the experiments.

Model	Parameter Size
RoBERTa (base)	125M
ModernBERT (base)	150M
DeBERTa v3 (base)	184M
Gemma 2 2B	2.2B
Mistral v0.3 7B	7.0B
Llama 3.1 8B	8.0B
Gemma 2 9B	9.0B

Table 4: Parameter sizes of models used in experiments.