

Politics of Questions in News: A Mixed-Methods Study of Interrogative Stances as Markers of Voice and Power

Victor Bros^{1, 2}, Matilde Barbini², Patrick Gerard³, Daniel Gatica-Perez^{1, 2}

¹Idiap Research Institute, Switzerland

²EPFL, Switzerland

³University of Southern California, USA

victor.bros@idiap.ch, matilde.barbini@epfl.ch, pgerard@isi.edu, gatica@idiap.ch

Abstract

Interrogatives in news discourse have been examined in linguistics and conversation analysis, but mostly in broadcast interviews and relatively small, often English-language corpora, while large-scale computational studies of news rarely distinguish interrogatives from declaratives or differentiate their functions. This paper brings these strands together through a mixed-methods study of the "Politics of Questions" in contemporary French-language digital news. Using over one million articles published between January 2023 and June 2024, we automatically detect interrogative stances, approximate their functional types, and locate textual answers when present, linking these quantitative measures to a qualitatively annotated subcorpus grounded in semantic and pragmatic theories of questions. Interrogatives are sparse but systematically patterned: they mainly introduce or organize issues, with most remaining cases being information-seeking or echo-like, while explicitly leading or tag questions are rare. Although their density and mix vary across outlets and topics, our heuristic suggests that questions are overwhelmingly taken up within the same article and usually linked to a subsequent answer-like span, most often in the journalist's narrative voice and less often through quoted speech. Interrogative contexts are densely populated with named individuals, organizations, and places, whereas publics and broad social groups are mentioned much less frequently, suggesting that interrogative discourse tends to foreground already prominent actors and places and thus exhibits strong personalization. We show how interrogative stance, textual uptake, and voice can be operationalized at corpus scale, and argue that combining computational methods with pragmatic and sociological perspectives can help account for how questioning practices structure contemporary news discourse.

1 Introduction

The digital transition has reshaped how news is produced, financed, and encountered. Audiences increasingly consume news online and via platforms, while legacy outlets face intensified competition for attention and pressure on business models (Hindman 2018; Udris, Fürst, and Eisenegger 2024). In many countries this has contributed to "news deserts", where local reporting capacity erodes and communities lose access to robust, professional journalism, with documented

democratic harms (Abernathy 2023; Vogler, Weston, and Udris 2023; Shaker 2014; Barclay et al. 2025). Against this backdrop, questions about who and what is made visible in news, whose problems are problematized, and how publics are addressed acquire renewed democratic significance (Pickard 2020).

These democratic roles are enacted not only through topic selection and framing, but also through interactional practices such as questioning and challenging. In broadcast settings, conversation-analytic work has shown how question design ranges from deferential information-seeking to highly adversarial, accountability-pressuring formats (Greatbatch 1988; Clayman and Heritage 2002; Heritage 2012). Linguistics and pragmatics emphasize that interrogatives are context-dependent acts that can request information, highlight issues, embed presuppositions, or function rhetorically (Athanasidou 1991; Ilie 1994; van Rooy 2003; Ciardelli 2016). Most of this work relies on close analysis of relatively small, often English-language corpora. Computational studies of rhetorical questions and stance in online debate (Ranganath et al. 2016; Oraby et al. 2017) show that some pragmatic distinctions can be modeled automatically, but remain limited in genre and scale. Large-scale NLP studies of news, in turn, focus mainly on topics, sentiment, or frames and usually treat all sentences alike, without distinguishing interrogatives from declaratives or separating different interrogative functions.

French-language news provides a particularly interesting setting for bridging these strands. It spans multiple media systems and political contexts, including France, Switzerland, Belgium, francophone Canada, Senegal, and pan-African outlets, with distinct histories of press freedom and journalistic cultures (Hallin and Mancini 2004; Straubhaar 2021). Yet French functions as a shared linguistic resource, enabling partly common repertoires of questioning and address while leaving room for local variation in how power and publics are constructed. This combination makes francophone news a useful testbed for our approach: it moves beyond the Anglophone focus of much prior work while keeping a single language and covering both Global North and Global South media systems. We bring a mixed-methods, computational lens to what we call the *Politics of Questions* in this fragmented but interconnected media space. Using more than 1.2 million French-language digital news

articles (January 2023-June 2024), we automatically identify interrogative stances in context, approximate their functional types, locate textual answers when they are present, and characterize the actors and places that populate questions and answers. These quantitative components are linked to a qualitatively annotated subcorpus, grounded in semantic and pragmatic theories of interrogatives, which we use for both evaluation and discourse analysis.

Focusing on *interrogative stances* allows us to make the politics of visibility and accountability empirically tractable. Questions are a particularly salient subset of journalistic moves: they mark issues as open, allocate epistemic authority, and can be detected and typed from surface cues in ways that most other interactional practices cannot. At the same time, they are normatively charged, because asking or not asking, and how a question is framed, can shape which actors are foregrounded, which problems are rendered discussable, and how accountability is textually staged.

This leads to the following research questions:

RQ1: How prevalent are interrogative stances in French-language news, and how do their types vary across outlets, countries, and topics?

RQ2: When journalists pose interrogative stances, to what extent are those questions textually answered, and how often do answers involve quoted external voices versus self-answers by the journalist?

RQ3: Which kinds of actors and social groups are most often foregrounded in interrogative contexts, and how frequently do such contexts include quoted external voices rather than only journalistic narration?

RQ4: How can a triangulated methodology, combining quantitative and qualitative analyses, enhance the interpretation and validity of the quantitative measures?

Contributions. Our contributions are threefold: (i) conceptually, we extend semantic and pragmatic work on questions to written news by operationalizing *interrogative stances*, answerhood, and dialogicity at corpus scale; (ii) empirically, we provide, to our knowledge, the first large-scale, cross-regional description of interrogative practices in French-language digital news, showing how interrogatives structure coverage, how often they are resolved in-text, and which actors and places they center; and (iii) methodologically, we propose and validate a triangulated pipeline in which large language models and fine-tuned neural classifiers are combined with qualitative annotation to yield interpretable proxy measures relevant to debates on gatekeeping, voice, and agenda-setting.

2 Related Work

Theories of questions: semantics and pragmatics Linguistics and philosophy of language treat questions not just as syntactic configurations but as discourse moves whose interpretation depends on context. Formal semantics represents questions as abstract objects that partition the space of possible worlds and determine which answers resolve an issue (Groenendijk and Stokhof 1984; Ciardelli, Groenendijk, and Roelofsen 2018), while contextual models such as Questions Under Discussion and Table-based accounts (Farkas and Bruce 2010; Ginzburg 2015) embed

this view in a dynamic picture of interaction where questions guide relevance and shape how common ground is updated. Pragmatic studies emphasize that interrogatives must be classified according to speaker intentions and social situations: they can request information, highlight issues, regulate knowledge display, or indirectly request action (Hudson 1975; Athanasiadou 1991). Classic distinctions between information-seeking, rhetorical, leading, and echo questions, and between more or less “conductive” designs, show how presuppositions and turn design can steer respondents toward particular answers or implicatures (van Rooy 2003; Ilie 1994; Frank 1990; Braun 2011). These insights underpin our use of *interrogative stance* as a contextual property of utterances and motivate the stance types we distinguish in the empirical analysis.

Questions, accountability, framing, and gatekeeping in news discourse Beyond their semantic content, questions enact social relations and allocate accountability. Athanasiadou (Athanasiadou 1991) notes that questioning carries a “dominance function” tied to roles: the same interrogative can be heard as a neutral request among peers or as a command in hierarchical settings. Stivers and Rossano (Stivers and Rossano 2010) show that interrogative turns mobilize responses with varying strength depending on action, design, and sequential position. In news interviews, Greatbatch (Greatbatch 1988) and Clayman and Heritage (Clayman and Heritage 2002) document how journalists control topical progression and use question design to negotiate norms of neutrality and adversarialness. Diachronic work on U.S. presidential press conferences finds a rise in more direct and negative interrogatives, interpreted as evidence of increasing pressure on presidents to respond in particular ways (Clayman and Heritage 2023). Media sociology and communication research add a broader perspective on how news institutions filter and present public issues. Classic gatekeeping work shows how newsroom routines and professional values prioritize certain topics, locations, and source types, leading to the overrepresentation of political and economic elites (Gans 2004; Shoemaker and Vos 2009). Framing theory conceptualizes how media selectively emphasize certain aspects of reality to promote particular problem definitions, causal stories, or remedies (Entman 1993). Our focus on interrogative stances, their answerability, and their actor populations connects these interactional and sociological perspectives by asking how questions help to distribute accountability and voice in news discourse.

Computational approaches to questions and news content Computational work has begun to model pragmatic functions of questions, but mostly in social media and with coarse distinctions. Ranganath et al. (Ranganath et al. 2016) and Oraby et al. (Oraby et al. 2017) compile corpora of rhetorical questions from online debate forums and Twitter and train classifiers that distinguish rhetorical from non-rhetorical and sarcastic from non-sarcastic questions with solid performance. These studies demonstrate that some pragmatic categories can be learned automatically, but focus on English user-generated dialogue and do not consider how questions are taken up or answered in subsequent discourse.

In journalism and communication studies, most computational analyses of text have focused on topics, frames, and temporal dynamics rather than interrogatives. Large-scale frame analyses apply supervised or semi-supervised models to millions of news articles to trace how frames fluctuate with events and evolve over longer periods (Kwak, An, and Ahn 2020; Guo 2011; Dehler-Holland, Schumacher, and Fichtner 2021). Other work combines topic modeling with time-series methods to study intermedia agenda-setting, for example in coverage of vaccine scandals or policy crises on different platforms (Shi and Wang 2023). Recent studies use topic models and named-entity recognition to assess how well COVID-19 coverage fulfills normative functions such as informing, monitoring, and providing a platform for debate (Nguyen and van Es 2024). Across this literature, sentences are typically treated uniformly for modeling topics, frames, or sentiment, and questions are not analyzed as distinct interactional moves. Moreover, even when headlines or rhetorical devices are examined qualitatively, there is little systematic attention to the micro-pragmatics of questioning, such as conduciveness, presupposition, or dialogicity, across large corpora. Our work complements these approaches by foregrounding interrogative acts themselves and by linking question design, answerhood, and actor mentions to agenda-setting and gatekeeping.

Mixed-methods and triangulated computational media research A parallel methodological literature argues for combining computational text analysis with qualitative inquiry. Triangulation is understood not only as cross-validation but as bringing multiple kinds of evidence and perspectives to bear on the same phenomenon (Olsen 2004). Nelson (Nelson 2020) formulates this as “computational grounded theory”, in which unsupervised or weakly supervised methods surface patterns, qualitative analysis refines concepts, and further computational analyses test their scope. Large language models are increasingly integrated into such pipelines as annotation or coding tools. Gilardi et al. (Gilardi, Alizadeh, and Kubli 2023) compare ChatGPT to crowdworkers and trained annotators on multiple text-coding tasks and show that zero-shot LLM classifications can rival crowdworkers in accuracy and intercoder agreement at much lower cost, though human oversight remains essential. We build on this line of work by using an LLM to generate high-confidence pseudo-labels of interrogative stance, distilling these into fine-tuned neural classifiers, and then interpreting their outputs through a qualitatively annotated subcorpus. This triangulated design allows us to operationalize pragmatic concepts such as interrogative stance, answerhood, and voice at scale while retaining a close connection to discourse-level analysis.

3 Methodology

3.1 Data

We assemble a French-language news corpus from two complementary sources. The first component is derived from CCNews, a large-scale news collection built from Common Crawl (CommonCrawl 2024). Starting from the most active French-language CCNews domains in our period (top 40 by

article volume), we manually retained outlets that (i) publish in French, (ii) remain active in both 2023 and 2024, (iii) can be identified as established news organizations via their domains and websites, and (iv) yield sufficiently clean article extractions, excluding domains with substantial scraping noise. To avoid overconcentration in a single national market, we also prioritized coverage across distinct francophone regions. This yields just over 1.2 million articles spanning global, national, and regional coverage from French-speaking Europe, America, and Africa.

Hyper-local actors and issues are underrepresented in this CCNews-derived subset. To better cover local journalism, we integrate a second open dataset of Swiss local news (Bros and Gatica-Perez 2025), which provides high-quality reporting from cantonal and hyper-local outlets in Romandy. From this collection we include all French-language articles published between 2023 and mid-2024. Combined, the final corpus comprises 24 outlets. For each outlet we manually code an *editorial scale* (hyper-local, regional, national, transnational, thematic) and a primary *country or region*, used in later analyses of cross-country and cross-scale variation. This corpus should therefore be read as a curated comparative sample rather than a statistically representative census of French-language digital news. All country- and scale-level aggregates should therefore be read as descriptive comparisons within this curated sample rather than representative estimates of francophone news systems, because outlet volumes are highly unequal, larger outlets contribute more to article-level group means. Full outlet lists and counts are provided in Appendix A.

Ethical considerations All data come from publicly accessible news websites or established web corpora. All external models and corpora we use are under permissive research licenses. We analyze text for scientific purposes only, report results in aggregate primarily, use a small number of short, translated and lightly paraphrased excerpts, and do not attempt to deanonymize individuals beyond what is already present in published news. We do not redistribute full-text articles. Any shared data will be limited to derived annotations and metadata, subject to the terms of use of the original sources.

3.2 Definition of interrogative stance

Following semantic and pragmatic work on questions (Hudson 1975; Ciardelli 2016; Ginzburg 2015; Athanasiadou 1991), we treat *interrogative stance* as a contextual property of utterances rather than a purely syntactic category. At the sentence level, a journalistic sentence expresses an interrogative stance if it does at least one of the following:

1. Explicitly presents a question or problem to be answered (direct questions with “?” or indirect questions).
2. Implicitly raises an issue or unknown whose resolution is projected in the surrounding discourse.
3. Highlights that some aspect of the situation remains open, unresolved, or unanswered, thereby framing it as an issue rather than a settled fact.

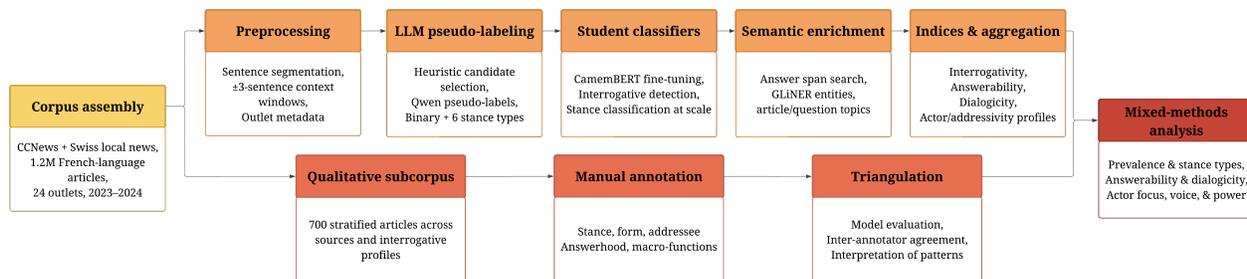


Figure 1: **Overview of the mixed-methods pipeline for interrogative stance analysis.** Starting from a corpus of 1.2M French-language news articles from 24 outlets (2023-2024), the upper branch applies a teacher-student NLP pipeline to detect interrogative stances and derive measures of interrogativity, answerability, dialogicity, and actor/addressivity profiles. In parallel, a stratified qualitative subcorpus is manually annotated for stance, form, addressee, answerhood, and macro-functions. The two branches are then triangulated to evaluate the models and interpret corpus-level patterns of prevalence, answerability, voice, and power.

Sentences that merely mention a "question" or "issue" as a topic, without presenting a concrete unknown or projecting its resolution, are treated as non-interrogative. Likewise, purely evaluative statements of uncertainty are not classified as interrogative unless they are clearly framed as an open question to be addressed.

Building on pragmatic typologies (Athanasidou 1991; Ilie 1994; Frank 1990; van Rooy 2003), we distinguish six interrogative stance types:

- **Framing-procedural:** interrogative or issue-raising sentences whose primary role is to structure the article (introduce a problem, announce a section) or to mark that a question remains open, rather than to request an answer from a specific interlocutor.
- **Information-seeking:** sentences that express a genuine request for factual information or clarification; the writer or quoted speaker is presented as lacking the information and expecting an answer.
- **Rhetorical:** formally interrogative sentences that do not genuinely seek an answer, but instead express evaluation, criticism, or emphasis; the answer is treated as obvious or irrelevant.
- **Leading:** questions that guide the reader or addressee toward a particular answer or interpretation, typically by embedding presuppositions or value judgments. Strongly oriented rhetorical questions are subsumed here.
- **Tag:** declarative statements followed by a short interrogative tag seeking confirmation or alignment.
- **Echo-clarification:** interrogative fragments or sentences that repeat or closely mirror prior wording in order to check understanding, register surprise, or distance oneself from a quoted formulation.

3.3 Automatic detection of interrogative stances

Our goal is to automatically detect interrogative stances and assign stance types to all sentences in more than one million French-language news articles. We adopt a teacher-student pipeline in which a large language model provides pseudo-labels on a subset of sentences, and a French transformer

model, CamemBERT-large (Antoun et al. 2024), is then fine-tuned on these pseudo-labels and applied at scale.

Articles are segmented into sentences, and during model training and inference each sentence is represented together with a short local context window (neighboring sentences from the same article). This preserves local discourse cues that are often crucial for interpreting interrogative force, while still allowing sentence-level labels. Full preprocessing details are given in Appendix B.

To obtain high-quality pseudo-labels, we apply a French-capable LLM, Qwen3-30B-A3B-Instruct (Team 2025), to a subset of sentences detected by a high-recall heuristic based on question marks, French interrogative patterns, and common question-raising constructions. All candidates, plus a random subsample of non-candidates for calibration, are passed to the LLM together with their local context. The model is prompted to perform a two-stage classification: (i) decide whether the target sentence expresses an interrogative stance (binary), and (ii) for positives, assign one of the six stance types defined above. We retain only high-confidence pseudo-labels for training. Implementation is documented in Appendix B.

Using these pseudo-labels, we fine-tune two CamemBERT-large classifiers: a **binary interrogative detector** (interrogative vs. non-interrogative) and a **six-way stance classifier** for interrogative sentences. High-confidence LLM labels are split into training and validation sets with articles kept disjoint across splits, and models are fine-tuned with weighted cross-entropy loss to compensate for label imbalance. Training hyperparameters and architecture details are reported in Appendix B. On held-out pseudo-labeled validation sets, the binary detector achieves high accuracy (around 0.97), while the stance classifier reaches reasonable six-way performance (macro-F1 around 0.69), providing a usable approximation of pragmatic types for large-scale analysis. When evaluated against our human gold-standard annotations (Section 3.7), the binary detector reaches an F1 of 0.78 for the interrogative class and the six-way stance classifier macro-F1 of 0.51, which we treat as sufficient for corpus-level analyses rather than

sentence-level claims.

At corpus scale, we apply the binary and stance models sequentially to every sentence. We store both predictions and confidence scores for downstream analyses in Section 4.

3.4 Answer identification

To approximate whether interrogative stances are textually resolved within an article, we implement an embedding-based answer search. For each article, we consider as *questions of interest* all sentences predicted as interrogative with sufficient classifier confidence, and group immediately consecutive questions into a single local issue. We then compute sentence embeddings for all sentences in the article using CamemBERT and, for each group of questions, search only among subsequent sentences for the most similar contiguous span of limited length. If the best-scoring span exceeds a similarity threshold, we treat it as the textual answer and record its location and text. Otherwise, the question group is marked as unanswered in that article. This procedure yields an approximate but systematic measure of textual uptake at scale, namely whether an interrogative is followed by a semantically matched answer-like span within the same article. The resulting rates should be read as indicative upper bounds rather than exact estimates of fully resolved answerhood. Exact implementation is detailed in Appendix B.

3.5 Entity and topic annotations

To characterize which actors and groups are mentioned in questions and answers, we run named-entity recognition (NER) on the subset of sentences involved in interrogative stances and their detected answer spans. We apply the multilingual GLiNER model (Zaratiana et al. 2023) with a coarse, robust label set: PERSON, ORGANIZATION, LOCATION, NATIONALITY OR RELIGIOUS OR POLITICAL GROUP, GENERIC SOCIAL GROUP, PUBLIC OR AUDIENCE, and EVENT. For each interrogative sentence, we run GLiNER on a short local context, and for each detected answer span we run it on the answer text. Entities with model scores below 0.5 are discarded. These annotations form the basis for our actor-level analyses of which actors, groups, and places are foregrounded in interrogative and answer contexts. They capture mention patterns rather than direct accountability relations: an entity named near a question is not necessarily the actor being addressed, challenged, or held responsible by that question.

We also derive topics at both article and question level. Using CamemBERT embeddings, we apply BERTopic (Grootendorst 2022) to the combined corpus to obtain unsupervised article topics, and in parallel run BERTopic on interrogative sentences (with their local context) to obtain question-level topics. For interpretability, we manually group the 100 most frequent BERTopic topics into eight broader *meta-topics* by inspecting top keywords and representative articles. Smaller or mixed clusters are assigned to the closest meta-topic on this basis. Algorithmic details are given in Appendix B.

3.6 Indices for interrogativity, answerability, dialogicity, and actor profiles

To relate micro-level interrogative practices to broader debates on agenda-setting, gatekeeping, and voice, we derive a set of interpretable indices from the sentence-level predictions, answer spans, and entity annotations. All indices are computed on the QA-enhanced sentence data and then aggregated by article, outlet, country, or topic.

At the article level, we first compute an *interrogative index* that captures how strongly an article structures its discourse through questions. For article a , let S_a be the total number of sentences and Q_a the number of sentences classified as interrogative (any of the six types) with sufficient confidence. The interrogative index is $ID_a = \frac{Q_a}{S_a}$, providing a simple measure of how much an article “problematizes” content rather than presenting it purely declaratively.

We then define an *answerability index*: $Ans_a = \frac{A_a}{Q_a}$, where A_a is the number of interrogative sentences in a for which the answer-identification procedure finds a sufficiently similar answer span in the subsequent text. Low Ans_a values indicate orphan questions that are raised but not taken up by the matching procedure in the article, while high values signal a tendency for interrogatives to be followed by semantically matched answer-like spans in-text, linking them more tightly to explanatory or justificatory work.

To distinguish monologic from more dialogic patterns, we compute a *dialogicity* profile. For each answered question, we detect whether the answer span contains markers of direct speech. At the article level, we track the shares of questions that are (i) unanswered; (ii) answered internally (only narrative text); and (iii) answered externally (through direct quotes). The proportion of externally answered questions in an article functions as a proxy for how far answers are delegated to sources rather than supplied solely by the journalist.

Using the NER annotations, we derive actor-level indices. Each interrogative sentence is categorized as *actor-focused* (mentioning one specific person, organization, or location), *group-focused* (mentioning only broad social or public categories such as national/religious/political groups, generic social groups, or audiences), or *issue-focused* (no detected human or collective entities). The relative frequencies of these types provide an *addressivity* profile, indicating whether interrogative contexts are centered on identifiable actors, diffuse publics, or abstract issues.

3.7 Qualitative analysis

We complement the automatic pipeline with a qualitatively annotated subcorpus used for evaluation and for detailed discourse analysis.

Sampling. We draw a stratified sample of 700 articles from the combined corpus, balancing data sources and interrogative profiles. Using the LLM-based pseudo-labels from the teacher model, we first mark articles as *question-containing* or *no-question*. Among question-containing articles, we compute a dominant pseudo-labeled interrogative stance per article to ensure coverage of the full range of interrogative functions. We then sample equally from

the CCNews-derived corpus and the Swiss local corpus, so that half of the manually annotated articles in each subset come from each source. In total, 400 articles (with both question-containing and no-question items) form the main evaluation set for the classification, 100 additional question-containing articles are double-coded by both annotators for inter-annotator agreement, and 200 further question-containing articles (100 per annotator) broaden the qualitative coverage of outlets and interrogative stances.

Coding scheme. The unit of analysis is a single *interrogative stance*, which may correspond to an explicit interrogative sentence or to an implicit question-introducing construction that pragmatically opens a Question Under Discussion in the sense of Roberts (Roberts 2012). For each unit, annotators first code the interactional context (interview vs. non-interview) and the addressee type (individual, collective actor, audience, self), then the grammatical form (wh-question, polar, alternative, tag, declarative-question, elliptic, indirect), and finally a primary pragmatic function label that maps onto the six stance types used in the automatic classifier (information-seeking, rhetorical, leading, tag, echo-clarification, framing-procedural).

Beyond these micro-pragmatic labels, the scheme captures how questions enact power and voice through four layers: (i) an interpersonal layer (register, inferred status relation, face-threat); (ii) an information layer (inquisitive vs. informative vs. rhetorical, presuppositions, conduciveness, directive force); (iii) a stance/style layer (evaluative stance, irony markers, performative display); and (iv) an uptake layer (answerhood expectation, resolution status, and whether an answer is realized in the text). On this basis, each interrogative stance is mapped to one or two macro-functional axes: *Authority positioning*, *Framing/agenda-setting*, *Stance/alignment*, *Legitimation*, or *Discursive strategy*. The detailed codebook and decision tree are available at <https://gitlab.idiap.ch/socialcomputing/politics-of-questions>. Appendix D reports inter-annotator agreement metrics.

4 Results

4.1 RQ1: Prevalence and distribution of interrogative stances

Across the corpus, interrogative stances are relatively rare but not exceptional. On average, about 2.5% of sentences in an article are classified as interrogative (mean interrogative index $ID_a = 0.025$, SD 0.054), and the median article contains no interrogative stance at all ($P_{50} = 0$, $P_{90} \approx 0.09$). Among interrogative sentences, framing-procedural questions are the dominant type, accounting for just over half of all interrogatives. Taken together, these distributions indicate that in written news interrogatives are used far more often to structure exposition and highlight issues for readers than to pose open information requests to specific actors. Information-seeking questions account for about one fifth, rhetorical questions for roughly one seventh, and echo-clarification questions for just under one tenth, while explicitly leading and tag questions each make up only about

Table 1: Global distribution of predicted interrogative stance types in the corpus.

Stance	N	% of interrogatives
echo-clarification	70,281	9.2
framing-procedural	396,298	52.1
information-seeking	153,404	20.2
leading	16,221	2.1
rhetorical	107,862	14.2
tag	16,116	2.1

2% of interrogatives (Table 1). Given the moderate reliability of the six-way classifier, we interpret these stance-type shares as coarse corpus-level approximations and avoid strong claims about small differences among rare or pragmatically adjacent categories. This prevalence estimate is robust to reasonable decision thresholds: varying the binary/stance confidence cutoff from 0.6 to 0.8 changes the absolute number of detected interrogatives but keeps mean article-level interrogative density in a narrow band ($ID_a = 0.0236$ – 0.0261 , Appendix Table 5).

Interrogative density varies systematically across outlets and scales. Grouping by outlet country, Swiss francophone outlets have the highest mean interrogative index ($\bar{I} \approx 0.033$), followed by French outlets ($\bar{I} \approx 0.028$), with Canadian outlets somewhat lower ($\bar{I} \approx 0.022$). Belgian, Senegalese, and pan-African outlets cluster around 0.017–0.018, while outlets coded as broader "Europe" or "International" have the lowest interrogative indices ($\bar{I} \approx 0.011$ – 0.014). When grouping by editorial scale, thematic outlets show the highest interrogative index ($\bar{I} \approx 0.036$), national outlets follow ($\bar{I} \approx 0.024$), regional outlets are slightly lower ($\bar{I} \approx 0.022$), and transnational outlets lowest ($\bar{I} \approx 0.020$). Figure 2 summarizes these patterns. Overall, interrogative stances are used more intensively in context-rich, domestically anchored coverage and in thematic verticals, and less in transnational or wire-type reporting.

To examine topical variation, we cluster articles with BERTopic and group frequent topics into eight meta-topics (professional sports, national/local politics, geopolitics, local news, lifestyle/entertainment, *faits divers*, business/economy, technology). Three domains stand out as particularly "question-saturated": professional sports (mean $ID_a \approx 0.029$ over $\sim 181,000$ articles), local news (≈ 0.030 over $\sim 39,600$ articles), and lifestyle/entertainment (≈ 0.027 over $\sim 44,500$ articles). Business/economy, geopolitics, and technology are less interrogative (means around 0.018–0.022), with national/local politics and *faits divers* in the middle (around 0.021–0.022, see Appendix Table 3). Stance distributions are broadly similar across topics: framing-procedural questions account for around half to two-thirds of interrogatives everywhere, information-seeking questions are relatively more prominent in *faits divers* and business/economy, while rhetorical and leading questions are somewhat more frequent in national/local politics and geopolitics.

These aggregate patterns match how interrogatives are

used in individual articles. In a Canadian explainer on procedures for a one-off family allowance, for instance, the piece is explicitly structured around a sequence of questions (*Who is eligible? How do you apply? What documents are needed?*), each immediately followed by a short answer. Here, interrogatives function as framing-procedural devices that break down a complex policy into manageable steps, illustrating why local and service-oriented topics have relatively high interrogative indices. By contrast, a Swiss commentary on a proposed pension payment is almost entirely declarative, but closes its presentation of the proposal with a single question, *What should we make of this argument?*, which serves as a rhetorical hinge to orient readers toward the ensuing evaluative discussion. Even in low-interrogative articles, a lone framing question can thus perform macro-structuring work.

4.2 RQ2: Answerability and dialogicity

Answerability. Once interrogative stances appear in news text, they are overwhelmingly taken up in subsequent discourse. Across the corpus, 760,182 sentences are classified as interrogative, and for 726,911 of them our heuristic identifies a subsequent answer-like span in the same article, corresponding to a 95.6% rate of heuristic local textual uptake, which we interpret as an upper bound on answerhood rather than as verified resolution in the strict pragmatic sense. At the article level, the median answerability index Ans_a is 1.0 in all outlet groups: in a typical article that contains questions, almost all are followed by some form of semantically matched textual uptake. Answerability varies only modestly across stance types: information-seeking and echo-clarification questions have the highest match rates (97.2% and 97.4%), framing-procedural questions the lowest (94.8%), and rhetorical, leading, and tag questions fall in between (around 95%, see Appendix Table 4). Taken together with our qualitative coding, this suggests that even when interrogatives function evaluatively or rhetorically, they are usually embedded in local question-answer sequences rather than left entirely hanging. Because the answer-matching step captures semantic relatedness rather than manually verified resolution, we interpret the 95.6% figure as a heuristic upper bound on local textual uptake rather than a literal estimate of fully resolved answerhood. Still, this result is not driven by a narrow parameter choice: answerability is unchanged for cosine thresholds from 0.05 to 0.80 (Appendix Table 6), and a manual audit of 50 question groups found clear or partial local answers in 34/40 predicted answered cases, an answer somewhere in the article in 37/40, and confirmed all 10 predicted unanswered cases (Appendix Table 7).

A minority of "orphan" questions nonetheless remain. For example, in a French local story on new street ashtrays installed to curb cigarette litter, the journalist reports municipal justifications and local associations' concerns before asking whether more bollard-ashtrays should be added at a specific location. The article ends without answering this question, leaving readers to infer that current measures may be insufficient or inappropriate. Such unresolved, suggestion-like interrogatives correspond closely to

the small set of unanswered questions flagged by our automated method.

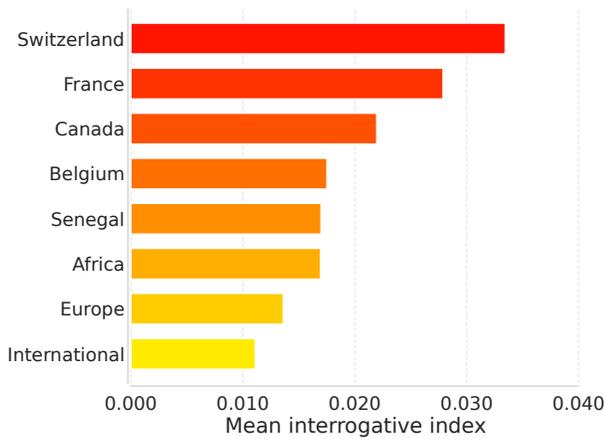
Dialogicity. Although most interrogatives are textually taken up, the form of that uptake differs. Overall, 4.4% of interrogative sentences remain unanswered, 80.3% are answered internally by the journalist in narrative text, and 15.4% are answered through direct quotes or clearly marked reported speech (Appendix Table 4, bottom panel). In other words, questions in French-language news are overwhelmingly taken up in-text, but in roughly one out of six cases this uptake takes the form of letting a quoted source "speak back" rather than the journalist alone filling the gap.

These shares are relatively stable across countries and outlet scales in terms of answerability, but external dialogicity varies more. Swiss and Belgian outlets show the highest mean fraction of questions answered by quoted speech (around 18%), followed by French, pan-African, and Senegalese outlets (around 15–16%), while Canadian and international outlets are somewhat lower (around 13%). By editorial scale, hyper-local outlets have the highest external dialogicity (20%), national, regional and transnational outlets follow closely, and thematic outlets (sports, lifestyle, verticals) show the lowest share (around 13%, Figure 3). Together, these results suggest that while questions are almost always resolved, national and regionally rooted outlets more frequently embed answers in dialogic exchanges with sources, whereas thematic verticals rely more heavily on internal narration.

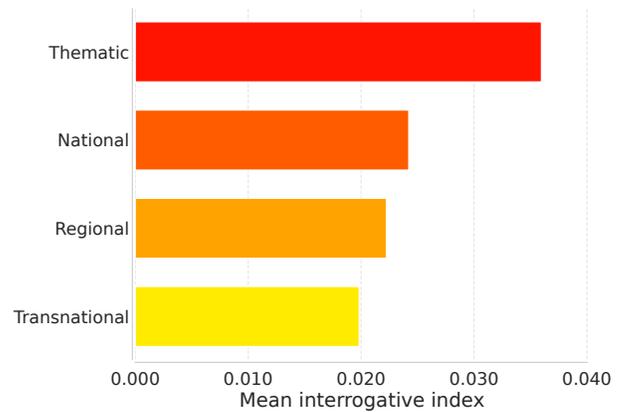
External dialogicity is visible in concrete articles. In a piece on doctored political images and AI-generated "fake news", for instance, the journalist raises the issue of where the misleading pictures originated and immediately attributes responsibility in a quoted exchange with a radio host who admits to having created them. Here, an information-seeking question is resolved through an external answer, matching our finding that a substantial share of questions are answered dialogically via quoted sources.

4.3 RQ3: Who is questioned, and how are interrogatives populated with actors and places?

Where questions are voiced. Interrogative stances occur predominantly in journalistic narration rather than as part of direct quotations. Globally, only 7.8% of interrogative sentences contain quotation markers in the question sentence itself, and 22.3% if we consider the concatenation of the question and its immediate answer span (figures derived from the same counts as Appendix Table 4). The remaining three quarters of questions are thus framed in the journalist's own voice, even when followed by quoted answers. This asymmetry suggests that, within the article text, journalists more often than sources control which questions are explicitly articulated. Sources more often appear as respondents within question-led structures than as initiators of the interrogative agenda. This pattern holds across stance types, but rhetorical, leading, and echo-clarification questions are noticeably more likely to appear inside quoted speech, reflecting their frequent use in talk shows, political interventions,



(a) Mean interrogative index by region.



(b) Mean interrogative index by editorial scale.

Figure 2: Variation in interrogative density across (a) regions and (b) editorial scales in French-language news (2023–2024).

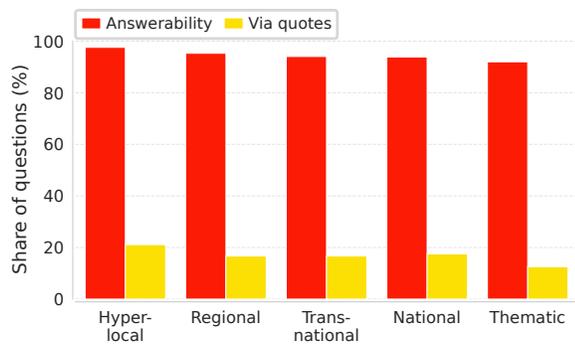


Figure 3: Mean answerability and external dialogicity by editorial scale. Values are shares of all interrogative sentences in each scale.

or opinionated statements by sources. Information-seeking and framing-procedural questions, by contrast, are typically voiced by journalists.

Personalization and collectivization. Named-entity annotations show that interrogative space appears densely populated with concrete actors and sites. A majority of interrogative sentences mention at least one PERSON (58.5%), and nearly half mention an ORGANIZATION (43.7%), almost one third contain a LOCATION (31.8%), and one quarter co-occur with an EVENT entity such as a match, election, or crisis. By contrast, explicit collectivizing references are comparatively rare: only 2.7% of questions mention a PUBLIC OR AUDIENCE, 1.3% a NATIONALITY OR RELIGIOUS OR POLITICAL GROUP, and 0.3% a GENERIC GROUP. Interrogative space in this corpus is therefore strongly personalized and institutionalized, with broad social groups and publics appearing much less frequently as explicit addressees. In most cases, questions thus function as organizing devices in discourse *about* actors and places rather than as contributions to an ongoing dialogic exchange with them. Part of this pattern likely reflects standard news sourcing norms, which privi-

lege identifiable elites and institutions over diffuse publics, and the fact that appeals to “the public” are often left implicit rather than being named as such.

This personalization is high across all countries but varies in degree. Senegalese outlets have the highest personalization index (68.1% of questions mentioning at least one person), followed by Canadian (60.9%) and French (59.8%) outlets. Swiss outlets and Europe-wide or transnational feeds display somewhat lower levels (around 47–50%), suggesting a slightly less person-centered interrogative style. When grouping by editorial scale, thematic verticals (especially sports and people/lifestyle) show the highest personalization (72.8% of interrogatives mentioning a person), national outlets are intermediate (55.9%), and regional and hyper-local outlets are lower (50.1% and 44.0%), consistent with a stronger emphasis on institutions, places, and issues than on a small set of named individuals. Collectivization remains a minority phenomenon everywhere (roughly 3–6% of questions mention a collective), slightly more common in pan-African, Senegalese, and transnational outlets. For brevity, full breakdowns by country and scale are omitted.

Institutional and geographic anchoring by topic. Entity annotations also reveal systematic differences in how questions are anchored in institutions, places, and events across topics (Appendix Table 3). In business and economy coverage, interrogatives are especially institution-centered: nearly 58% of questions mention an organization, compared to 30% mentioning a location and 25% an event. Professional sports also shows high organizational anchoring (52%) and strong event orientation (33%), with questions frequently tied to specific clubs, competitions, or fixtures. In geopolitics, geographic anchoring is pronounced: almost half of interrogatives mention at least one location and about one third an event, reflecting the centrality of countries, regions, and crises. National/local politics and local news likewise exhibit substantial organizational and locational anchoring, indicating that interrogative contexts in these domains are frequently organized around institutions and localities. By

contrast, lifestyle/entertainment and technology topics have lower institutional and geographic anchoring, with interrogatives more often revolving around individuals or generalized situations than around specific organizations or places.

Frequency lists of named entities provide a more concrete picture of who and what dominates interrogative space. Among EVENT labels, major sports competitions and high-salience political or social events are most frequent (*Mercato, élections, grève*). The LOCATION list is headed by national and regional nodes of attention, including *France, Paris, Ukraine, Québec, Gaza, Suisse, and Marseille*. For ORGANIZATION, football clubs and media brands dominate (*Real Madrid, Ligue 1*), alongside generic institutional labels (*gouvernement, police, Union européenne*). Among PERSON entities with proper names, high-profile athletes (Kylian Mbappé, Lionel Messi) and national leaders (Emmanuel Macron, Donald Trump) appear most often, together with a smaller set of recurring political and media figures. An example illustrates how interrogatives can simultaneously personalize, moralize, and anchor issues geographically. In a French report on protests against Israel's military actions, a quoted left-wing MP asks, *What is he waiting for before denouncing the Netanyahu government's war crimes?*, referring to the French president. The question is embedded in direct speech, targets named elite actors (the president and the Netanyahu government), presupposes culpable delay, and implicitly addresses a wider French public. It exemplifies the pattern whereby rhetorical and leading questions are often voiced by politicians or activists, concentrate on prominent individuals and governments, and combine interrogative form with strong evaluative force.

4.4 RQ4: Triangulated validation and qualitative insights

Model evaluation and reliability. We assess the reliability of our automatic pipeline by comparing model predictions to the manually annotated gold-standard sample (Section 3.7). On 8,399 gold sentences (both corpora, both annotators), the binary interrogative detector achieves 0.97 accuracy, with precision 0.76, recall 0.80, and F1-score 0.78 for the positive (interrogative) class. On the subset of 516 gold-positive sentences, the six-way stance classifier attains a macro-averaged F1 of 0.51 (micro-F1 0.51). Appendix Table 8 summarizes these aggregate figures. These scores indicate that while fine-grained pragmatic distinctions remain challenging, the models reliably capture interrogative vs. non-interrogative status and provide a usable, if imperfect, approximation of stance types for large-scale analysis. Because the student models are trained on high-confidence LLM pseudo-labels, they may also inherit systematic biases from the teacher model, especially for rare or pragmatically adjacent stance types. Appendix Table 9 and the conditional confusion matrix in Figure 4 further show that the main ambiguities are concentrated among framing-procedural, information-seeking, rhetorical, and leading questions: information-seeking and rhetorical cases are often absorbed into framing-procedural, while leading questions are frequently mapped to rhetorical or framing-procedural. In other words, most residual errors

arise among pragmatically adjacent categories rather than between clearly dissimilar question types.

We also measure inter-annotator agreement on a separate set of 98 articles double-coded by both annotators. A fuzzy one-to-one alignment of annotated spans yields an approximate Jaccard overlap of 0.83 over the union of segments. On the 204 matched segments, the two annotators reach 0.84 accuracy on the six-way stance labels, with Cohen's $\kappa \approx 0.78$. This indicates substantial agreement and confirms that the stance taxonomy is interpretable and can be applied consistently. The gap between human-human and human-model performance is in line with expectations for a nuanced pragmatic task and underscores the value of combining automatic predictions with qualitative analysis. At the same time, the gold annotations should be understood as a carefully constructed reference standard rather than an error-free ground truth: because interrogative stance is context-sensitive and often multifunctional, some model-gold disagreements plausibly reflect borderline cases that admit more than one reasonable reading. We therefore treat stance distributions as approximate and use them primarily for corpus-level patterns, complemented by detailed qualitative analysis. This caution is reinforced by the threshold checks reported above: the main corpus-level patterns are stable even when the absolute number of detected interrogatives varies modestly with stricter confidence cutoffs. This mixed-method, teacher-student design offers a pragmatic compromise: a relatively low-cost way to obtain corpus-scale coverage that remains anchored in a reliable qualitative scheme.

Macro-functions and qualitative insights. The qualitative coding shows that interrogatives in French-language news serve a small set of recurrent macro-functions. Across the coded sample, framing and agenda-setting questions dominate (60% of coded units), followed by authority-positioning interrogatives (30%). The remaining macro-axes, stance/alignment, discursive strategy, and legitimation each account for only a few percent of questions, but are structurally distinctive. Annotators frequently noted plausible *secondary* axes within a single question, for example when framing questions also expressed stance or legitimation. This multi-functionality suggests that journalistic interrogatives often combine several interactional jobs at once.

Three cross-cutting findings help interpret the large-scale patterns reported above. First, interview contexts act as a powerful normalizing force. Questions addressed to live interlocutors are overwhelmingly information-seeking, non-conductive, and low in face-threat, enacting what Clayman and Heritage term "soft accountability": journalists invite explanations or justifications from epistemically privileged sources and almost always receive answers. Outside interviews, by contrast, journalists use interrogatives more freely in a narrative-authority mode, posing questions that they then answer themselves. These questions often include presuppositions and section-opening roles, allowing journalists to display inquiry while retaining interpretive control.

Second, presuppositions are strongly associated with conduciveness, but the qualitative data clarify that it is their

content that matters. Many framing questions rely on factual presuppositions (e.g. that a statistical decline has occurred) that anchor shared context without steering evaluation, supporting relatively neutral framing. A smaller but important subset combines evaluative presuppositions with mild conduciveness, gently guiding readers toward particular interpretations, especially in hyper-local commentary. This refines classical claims that presuppositions are inherently ideological: in practice they are a flexible resource that can be deployed for both neutral and engaged framing.

Third, hyper-local outlets exhibit an “engaged but largely neutral” profile. They use more mildly leading or evaluative questions than national outlets, particularly when highlighting local institutional strains or community concerns, yet most of these interrogatives remain low in face-threat and leave space for competing views. This helps explain why, in the aggregate, hyper-local and regional outlets show higher interrogative indices and higher personalization of questions than transnational sources, while still maintaining very high answerability and relatively soft forms of accountability. Together, these qualitative insights support the validity of our computational indices and show how micro-pragmatic features, presupposition, conduciveness, addressee design, and context, combine to produce the macro-level distributions of interrogativity, answerability, dialogicity, and actor presence observed in French-language news.

5 Discussion

Our analysis shows that interrogative stances are sparse in French-language news but functionally rich. Although only a small share of sentences are interrogative, they are recruited for a limited and recurrent set of macro-functions: framing and agenda-setting dominate, authority-positioning forms a substantial secondary cluster, and stance/alignment, discursive strategy, and legitimation remain rarer. This pattern is broadly stable across countries, outlet types, and topics, while the frequent presence of secondary macro-axes confirms that individual interrogatives often perform more than one interactional job at once.

In terms of authority and the textual staging of accountability, our findings resonate with and nuance prior work on news interviews and epistemics (Clayman and Heritage 2002; Heritage 2012). Authority-positioning interrogatives in our corpus are overwhelmingly “soft” in Clayman and Heritage’s sense: in interview contexts, journalists typically adopt an epistemically subordinate stance, invite explanations from K^+ sources, and almost always receive textual uptake, a pattern more consistent with platform-based accountability than with adversarial challenge. Outside interviews, journalists use interrogatives in a narrative-authority mode, posing questions they then take up themselves and often answer in the surrounding narration, thereby consolidating interpretive control while sustaining a rhetorical appearance of inquiry. High answerability in this sense likely reflects not only accountability dynamics but also the conventions of explanatory and service-oriented digital journalism, where question-led structures are routinely used to organize exposition. Compared to Anglo-American traditions that have documented rising face-threat and negative

interrogatives in high-stakes press conferences (Clayman and Heritage 2023), French-language news in our corpus remains largely consensus-oriented: face-threatening questions are rare, and even hyper-local outlets with somewhat higher conduciveness and leading forms seldom resort to overtly confrontational designs.

Framing interrogatives suggest that questions can serve as a central resource for agenda-setting and issue construction (Entman 1993; Van Gorp 2007), but they also complicate simple claims about presuppositions and bias. The large majority of framing questions in our qualitative sample either introduce topics without presuppositions or employ factual presuppositions that anchor shared context without pushing a particular evaluation. A smaller but important subset uses evaluative presuppositions and mild conduciveness to gently steer interpretation, particularly in hyper-local journalism where community-oriented guidance appears more acceptable. Alternative (binary) questions, though rare, illustrate how interrogative form can discretely narrow the interpretive space by restricting readers to two options. These patterns refine Fowler’s (Fowler 2013) view of presuppositions as inherently ideological: in our data, presuppositions are a flexible pragmatic resource that can support neutral or engaged framing depending on what is taken for granted. Our findings on voice and actor salience can likewise be read in relation to classic concerns in media sociology (Shoemaker and Vos 2009; McCombs and Shaw 1993), while adding a pragmatic lens. At scale, interrogatives are heavily populated with named persons, organizations, and places, and a relatively small set of actors and sites, national leaders, key institutions, emblematic locations, recur disproportionately in questions, especially in sports and political coverage. Publics and broad social groups, by contrast, are rarely addressed directly in interrogative form. These mention patterns should be read as indicators of interrogative foregrounding and salience, not as direct measurements of who is actually being held accountable. Qualitatively, we observe two main voice-allocation mechanisms: in interviews, actors are explicitly given the floor to “speak back” under soft accountability. In non-interview contexts, questions more often organize discourse *about* actors and issues without creating interactional rights to respond. This helps to explain why our dialogicity measures suggest that most questions are followed by textual uptake, often by the journalist, and only a minority of interrogatives are resolved through quoted voices.

Methodologically, the study illustrates how a triangulated, mixed-methods approach can render abstract pragmatic and sociological concepts empirically tractable at scale. Large language model pseudo-labels and fine-tuned CamemBERT classifiers allowed us to detect interrogative stances and approximate their types over 1.2m articles, an embedding-based heuristic provided usable answer spans, and multilingual NER, combined with simple indices, gave a coarse yet informative picture of actor and place salience. The qualitative coding then validated and nuanced these automatic outputs: it confirmed robust patterns (such as interview neutrality constraints and the link between presuppositions and conduciveness) and revealed more subtle distinctions (such

as neutral vs. evaluative irony, or neutral vs. engaged framing) that would be invisible from model outputs alone.

At the same time, our approach has clear limitations. More specifically, the teacher-student setup may propagate some of the pseudo-labeler’s preferences into the final classifiers. Human evaluation mitigates this risk, but does not eliminate it, particularly for infrequent and borderline categories. Our actor-level measures rely on named-entity mentions and quoted-speech markers, so they capture textual foregrounding and voice allocation only imperfectly and should not be read as direct measures of accountability or agenda-setting effects. It is restricted to written French news over a two-year period, stance and NER models are imperfect, especially for rare pragmatic types and marginal actor categories, and the qualitative sample, while theoretically informed, remains relatively small. Future work could extend this framework to broadcast interviews and parliamentary debates, refine actor typing (for instance, distinguishing state, corporate, NGO, and citizen voices more systematically), and replicate the analysis in other languages and media systems. More broadly, the combination of large-scale NLP and detailed pragmatic coding offers a promising route to studying how questions, answers, and the distribution of voice structure contemporary public discourse.

6 Conclusion

This paper examined the *politics of questions* in contemporary French-language news through a mixed-methods analysis of more than 1.2 million articles (2023-2024), combining LLM-based pseudo-labeling, fine-tuned neural classifiers, heuristic answer detection, multilingual NER, and a qualitatively coded subcorpus grounded in pragmatic theory.

For **RQ1**, we find that interrogative stances are quantitatively sparse but highly structured: only a small share of sentences per article are interrogative, yet framing-procedural questions dominate wherever they appear, with information-seeking and echo-clarification questions forming most of the remainder and explicitly leading and tag questions rare. Interrogative density varies systematically across media systems, with Swiss and French outlets and thematic verticals, especially sports, local news, and lifestyle, relying more on question-based structuring than business, geopolitics, or technology. For **RQ2**, our heuristic identifies an answer-like textual continuation for the vast majority of questions, but this uptake is mostly monologic: questions are usually followed by the journalist’s own narrative continuation, and only a minority are resolved through quoted speech. These rates should be read as upper-bound indicators of textual uptake rather than fully verified answerhood. National and some regional outlets are somewhat more dialogic than thematic or transnational sources, but the overall pattern of high uptake and limited external dialogicity is robust.

Turning to **RQ3**, our NER-based analysis reveals that interrogative space is densely populated with named persons, organizations, locations, and emblematic events, while explicit publics and generic social groups are rarely addressed. Personalization is especially strong in Senegalese and Canadian outlets and in thematic verticals, and somewhat lower in Swiss and hyper-local outlets, which more often foreground

institutions and places. Across topics, interrogatives in business and sports are strongly anchored in organizations and events, while geopolitical and local-political questions foreground locations. Elites and key institutions thus occupy the center of both questioning and response. Finally, for **RQ4**, our teacher–student pipeline yields high reliability for the binary distinction between interrogative and non-interrogative sentences and usable, though imperfect, performance for six-way stance types. Qualitative coding organized around macro-functional axes is essential for interpreting these outputs, distinguishing neutral from engaged framing and clarifying how presuppositions, conduciveness, and face-threat shape the interactional work of questions.

These findings speak to debates in journalism studies, pragmatics, and computational social science. In our corpus, interrogative stances appear intermittently and are most often mobilized to structure explanation or mark issues for consideration rather than to stage overt confrontation. Authority-positioning interrogatives tend to take “soft” forms: in interviews, journalists invite explanations from epistemically privileged sources, while in non-interview contexts they frequently pose questions they themselves go on to answer, combining displays of inquiry with a strong degree of narrative control. Framing questions, especially those with evaluative presuppositions or binary alternatives, show how interrogatives can subtly steer interpretation, yet many presuppositions function in a relatively neutral way, anchoring shared context without straightforwardly imposing ideological commitments. From a gatekeeping perspective, the strong personalization and institutionalization of interrogative discourse, combined with relatively modest direct address to publics, are consistent with the idea that questions tend to foreground actors already central to the news agenda and to allocate voice unevenly across article texts. Methodologically, the study demonstrates that integrating pragmatic concepts such as interrogative stance, presupposition, answerhood, and dialogicity into large-scale NLP can enrich computational accounts of agenda-setting and voice.

The study also has limitations. It is restricted to written French-language news over a two-year period, and several automatic components remain approximate: stance classification is imperfect, the NER schema is coarse, and the analysis relies on textual traces alone rather than the prosodic and interactional cues available in broadcast or face-to-face settings. Our indices of interrogativity, answerability, dialogicity, and voice should therefore be read as interpretable proxies rather than precise reconstructions of interactional practice. Future work could extend this framework to other media genres, languages, and historical periods while refining actor typing and dialogic attribution.

Acknowledgments

This work was supported by the EU Horizon Europe program through the ELIAS project (No. 101120237) and by the EU Horizon Europe program, Marie Skłodowska-Curie Actions (MSCA), alignAI project (No. 101169473). We thank Isabelle Segarini and Frédéric Keller (ESH Médias - <https://www.eshmedias.ch/>) for discussions and support with data news provision.

References

- Abernathy, P. M. 2023. News Deserts: A Research Agenda for Addressing Disparities in the United States. *Media and Communication*, 11(3): 290–292.
- Antoun, W.; Kulumba, F.; Touchent, R.; Éric de la Clergerie; Sagot, B.; and Seddah, D. 2024. CamemBERT 2.0: A Smarter French Language Model Aged to Perfection.
- Athanasiadou, A. 1991. The Discourse Function Of Questions. *Pragmatics; Vol 1, No 1 (1991)*, 1.
- Barclay, S.; Barnett, S.; Moore, M.; and Townend, J. 2025. Local news as political institution and the repercussions of ‘news deserts’: A qualitative study of seven UK local areas. *Journalism*, 26(9): 1803–1821. Publisher: SAGE Publications.
- Braun, D. 2011. Implicating Questions. *Mind & Language*.
- Bros, V.; and Gatica-Perez, D. 2025. The Suisse Romande Local News Dataset. *Proceedings of the International AAAI Conference on Web and Social Media*, 19(1): 2396–2401.
- Ciardelli, I.; Groenendijk, J.; and Roelofsen, F. 2018. *Inquisitive Semantics*. Oxford University Press. ISBN 978-0-19-881478-8.
- Ciardelli, I. A. 2016. *Questions in logic*. S.l: s.n. ISBN 978-90-6464-977-6. OCLC: 945552181.
- Clayman, S.; and Heritage, J. 2002. *The News Interview: Journalists and Public Figures on the Air*. Studies in Interactional Sociolinguistics. Cambridge University Press.
- Clayman, S. E.; and Heritage, J. 2023. Pressuring the President: Changing language practices and the growth of political accountability. *Journal of Pragmatics*, 207: 62–74.
- CommonCrawl. 2024. Common Crawl Corpus.
- Dehler-Holland, J.; Schumacher, K.; and Fichtner, W. 2021. Topic Modeling Uncovers Shifts in Media Framing of the German Renewable Energy Act. *Patterns*, 2(1): 100169.
- Entman, R. M. 1993. Framing: Toward Clarification of a Fractured Paradigm. *Journal of Communication*, 43(4): 51–58.
- Farkas, D. F.; and Bruce, K. B. 2010. On Reacting to Assertions and Polar Questions. *Journal of Semantics*, 27(1): 81–118.
- Fowler, R. 2013. *Language in the News: Discourse and Ideology in the Press*. Routledge.
- Frank, J. 1990. You call that a rhetorical question?: Forms and functions of rhetorical questions in conversation. *Journal of Pragmatics*, 14(5): 723–738.
- Gans, H. J. 2004. *Deciding What’s News: A Study of CBS Evening News, NBC Nightly News, Newsweek, and Time*. Northwestern University Press. ISBN 978-0-8101-2237-6.
- Gilardi, F.; Alizadeh, M.; and Kubli, M. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30): e2305016120. Publisher: Proceedings of the National Academy of Sciences.
- Ginzburg, J. 2015. *The Interactive Stance: Meaning for Conversation*. Oxford, GB: Oxford University Press.
- Greatbatch, D. 1988. A Turn-Taking System for British News Interviews. *Language in Society*.
- Groenendijk, J.; and Stokhof, M. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure.
- Guo, S. 2011. Framing distance: local vs. non-local news in Hong Kong press. *Chinese Journal of Communication*.
- Hallin, D. C.; and Mancini, P. 2004. *Comparing Media Systems: Three Models of Media and Politics*.
- Heritage, J. 2012. Epistemics in Conversation. In *The Handbook of Conversation Analysis*, 370–394. John Wiley & Sons, Ltd.
- Hindman, M. 2018. *The Internet Trap: How the Digital Economy Builds Monopolies and Undermines Democracy*. Princeton University Press. ISBN 978-0-691-15926-3.
- Hudson, R. A. 1975. The Meaning of Questions. *Language*.
- Ilie, C. 1994. *What Else Can I Tell You?: A Pragmatic Study of English Rhetorical Questions as Discursive and Argumentative Acts*. Almqvist & Wiksell International.
- Kwak, H.; An, J.; and Ahn, Y.-Y. 2020. A Systematic Media Frame Analysis of 1.5 Million New York Times Articles from 2000 to 2017. In *Proceedings of the 12th ACM Conference on Web Science*. Association for Computing Machinery.
- McCombs, M. E.; and Shaw, D. L. 1993. The Evolution of Agenda-Setting Research: Twenty-Five Years in the Marketplace of Ideas. *Journal of Communication*, 43(2): 58–67.
- Nelson, L. K. 2020. Computational Grounded Theory: A Methodological Framework. *Sociological Methods & Research*, 49(1): 3–42. Publisher: SAGE Publications Inc.
- Nguyen, D.; and van Es, K. 2024. Exploring the Value of Computational Methods for Metajournalistic Discourse: The Example of COVID-19 Reporting in Dutch Newspapers. *Journalism Studies*, 25(10): 1160–1181. Publisher: Routledge .eprint: <https://doi.org/10.1080/1461670X.2024.2358118>.
- Olsen, W. 2004. Triangulation in Social Research: Qualitative and Quantitative Methods Can Really Be Mixed.
- Oraby, S.; Harrison, V.; Misra, A.; Riloff, E.; and Walker, M. 2017. Are you serious?: Rhetorical Questions and Sarcasm in Social Media Dialog. In *Proceedings of the 18th SIGdial Meeting on Discourse and Dialogue*. Association for Computational Linguistics.
- Pickard, V. 2020. *Democracy without Journalism?: Confronting the Misinformation Society*.
- Qi, P.; Zhang, Y.; Zhang, Y.; Bolton, J.; and Manning, C. D. 2020. Stanza: A Python Natural Language Processing Toolkit for Many Human Languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics.
- Ranganath, S.; Hu, X.; Tang, J.; Wang, S.; and Liu, H. 2016. Identifying Rhetorical Questions in Social Media. *Proceedings of the International AAAI Conference on Web and Social Media*, 10.
- Roberts, C. 2012. Information structure: Towards an integrated formal theory of pragmatics. *Semantics and pragmatics*, 5: 6–1.
- Shaker, L. 2014. Dead Newspapers and Citizens’ Civic Engagement. *Political Communication*, 31(1): 131–148. Publisher: Routledge .eprint: <https://doi.org/10.1080/10584609.2012.762817>.
- Shi, J.; and Wang, H. 2023. Examining the Intermedia Agenda Setting Effects amid the Changsheng Vaccine Crisis: A Computational Approach. *International Journal of Environmental Research and Public Health*, 4052.
- Shoemaker, P. J.; and Vos, T. 2009. *Gatekeeping Theory*. New York: Routledge. ISBN 978-0-203-93165-3.
- Stivers, T.; and Rossano, F. 2010. Mobilizing Response. *Research on Language and Social Interaction*, 43(1).
- Straubhaar, J. 2021. Cultural Proximity. In *The Routledge Handbook of Digital Media and Globalization*. Routledge.
- Team, Q. 2025. Qwen3 Technical Report. arXiv:2505.09388.
- Udris, L.; Fürst, S.; and Eisenegger, M. 2024. Are news from public service media crowding out private news media? Usage and willingness to pay in Switzerland.
- Van Gorp, B. 2007. The constructionist approach to framing: Bringing culture back in. *Journal of communication*.
- van Rooy, R. 2003. Questioning to Resolve Decision Problems. *Linguistics and Philosophy*, 26(6): 727–763.
- Vogler, D.; Weston, M.; and Udris, L. 2023. Investigating News Deserts on the Content Level: Geographical Diversity in Swiss News Media. *Media and Communication*.
- Zaratiana, U.; Tomeh, N.; Holat, P.; and Charnois, T. 2023. GLiNER: Generalist Model for Named Entity Recognition using Bidirectional Transformer.

G Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes, and the work analyzes publicly available news articles and reports only aggregate statistics about outlets, topics, and actors, with no collection of private or user-generated data.**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Yes, and the Abstract and Introduction state the mixed-methods design, the corpus scope, and the contributions without overclaiming.**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes, and we explain why combining large-scale computational analysis with qualitative coding is suitable for our research questions in the Methodology and Results sections.**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes, and we describe corpus composition (countries, outlet types, editorial scales, time window) and note coverage and representativeness limitations, for example the focus on French-language digital news and on specific outlet sets.**
 - (e) Did you describe the limitations of your work? **Yes, and we discuss limitations in the Discussion and Conclusion.**
 - (f) Did you discuss any potential negative societal impacts of your work? **No, because the paper presents very limited risks in this regard. The contributions and the results are presented in aggregations.**
 - (g) Did you discuss any potential misuse of your work? **No, because the paper presents very limited risks in this regard.**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes, and the code for reproducibility are available at <https://gitlab.idiap.ch/socialcomputing/politics-of-questions>.**
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes, and the study has been designed to conform to ethics guidance as we understand it.**
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? **NA**
 - (b) Have you provided justifications for all theoretical results? **NA**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? **NA**
 - (b) Did you include complete proofs of all theoretical results? **NA**
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Due to news copyright constraints, we do not redistribute full-text articles. Code and derived non-copyright-restricted artifacts are available at <https://gitlab.idiap.ch/socialcomputing/politics-of-questions>**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes, and we describe model architectures, training/validation splits, and hyperparameters in the Methodology section.**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **No, because we report single-run metrics for each configuration and did not repeat training with multiple random seeds.**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **No, because the current draft does not yet specify hardware details or total compute. Time constraints for the publication prevented a reliable estimation of the compute resources.**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes, and we justify our evaluation via a manually annotated gold-standard subset, inter-annotator agreement analysis, and model-human comparisons, and we argue that the resulting performance is adequate for aggregate, corpus-level analyses.**
 - (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? **Yes, and we explain that stance and NER errors limit fine-grained interpretation, so our indices should be treated as approximate proxies rather than precise reconstructions, and we discuss these limitations in the Methodology, Results, and Discussion.**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity...**

- (a) If your work uses existing assets, did you cite the creators? **Yes, and we cite the creators of the news corpora and the main NLP models and toolkits we use (CCNews, the Suisse Romande corpus, CamemBERT, BERTopic, GLiNER, stanza, and the LLM teacher model).**
 - (b) Did you mention the license of the assets? **No, because we currently describe data sources and tools as they recommend to be cited but do not yet explicitly state licenses.**
 - (c) Did you include any new assets in the supplemental material or as a URL? **We release code and derived artifacts at <https://gitlab.idiap.ch/socialcomputing/politics-of-questions>, but not the reconstructed full-text corpus due to copyright restrictions.**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **No, because we exclusively use publicly available news articles.**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes, and we have not encountered such content during the development of the project. Though the corpus may contain personal names and potentially sensitive content as typical in news reporting. We analyze and report only in aggregate, do not redistribute full text, and avoid exposing individuals beyond what is already public.**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? **NA**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? **NA**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**
- (a) Did you include the full text of instructions given to participants and screenshots? **NA**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **NA**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **NA**
 - (d) Did you discuss how data is stored, shared, and de-identified? **NA**

A Outlets with their ontologies

Table 2 lists all outlets included in the corpus, with article counts for 2023–2024 and the manually coded outlet ontologies used in the analyses (primary country/region and editorial scale/type).

B Additional implementation details

For all the stochastic runs described below, we used a fixed random seed.

Preprocessing and sentence representation

All experiments were implemented in Python using stanza (Qi et al. 2020):

- Articles were segmented into sentences with the French stanza pipeline and each sentence was assigned an `article_id` and a sentence index (`sent_id`).
- We constructed local context windows by sorting sentences by `article_id`, `sent_id` and, for each sentence, concatenating up to *three* preceding and *three* following sentences from the same article into a single `context_text`.
- In `context_text`, the target sentence was wrapped in `<tgt> ... </tgt>` and neighboring sentences were separated by the delimiter " `</s>` ".

LLM pseudo-labeling with Qwen

All Qwen3-30B calls were run locally on local GPU servers (no prompts or texts were sent to external APIs). The exact system and user prompts used to query Qwen for binary and six-way stance decisions are released with the code at <https://gitlab.idiap.ch/socialcomputing/politics-of-questions>.

Candidate detection. To reduce LLM calls, we first flagged *candidate* interrogative sentences with a high-recall heuristic:

- Presence of a question mark "?"
- Sentence-initial interrogative patterns (case-insensitive), e.g. *comment*, *pourquoi*, *combien*, *est-ce que*, *y a-t-il*, *faut-il*, *peut-on*, *doit-on*, *que faire*, *que va-t-il*.
- Question-raising verb patterns, e.g. *on peut se demander*, *on est en droit de se demander*, *(se) demande(r)*, *(s') interroger*.
- Question-raising noun patterns, e.g. *pose la question*, *soulève la question*, *remet en question*, *la question se pose/reste/demeure*, *question en suspens*, *questions sans réponse*, *reste à savoir*.

In addition, a random subsample of non-candidate sentences (25% as many as candidates) was sent to Qwen for calibration.

Binary interrogative stance (teacher). For the binary decision (*interrogative stance* vs. *non-interrogative*), we used:

- Model: Qwen/Qwen3-30B-A3B-Instruct-2507.
- Input: a system prompt defining "interrogative stance" and a user prompt with up to three sentences of left context, the target sentence marked as "Phrase cible", and up to three sentences of right context.

- Decoding: greedy (temperature=0.0, top-p=1.0), max_new_tokens=64, batch size 8.
- Output: JSON with "is_interrogative" (boolean) and "confidence" $\in \{0.2, 0.5, 0.8, 0.95\}$.

Six-way stance classification (teacher). For sentences judged interrogative by the binary teacher, we used the same Qwen model with a different system prompt to assign one of six stance labels:

- Labels: "information-seeking", "rhetorical", "leading", "tag", "echo-clarification", "framing-procedural".
- Input: same context format as above; output: JSON with "label" in the six-way set and "confidence" $\in \{0.2, 0.5, 0.8, 0.95\}$.

For training student models, we retained only high-confidence LLM labels:

- Binary: positives (True) and negatives (False) with confidence ≥ 0.7 .
- Stance: six-way labels with confidence ≥ 0.7 .

Student models: CamemBERT classifiers

Binary interrogative detector. The binary student model is a CamemBERT-large sequence classifier:

- Base encoder: camembert/camembert-large.
- Labels: {"non-interrogative", "interrogative"}.
- Training data: high-confidence Qwen binary labels, excluding all articles in the human evaluation sets.
- Input: `context_text` with a ± 3 -sentence window and `<tgt>` markers, max sequence length 512 tokens.
- Split: Split on `article_id` with 10% held out for validation (no article overlap).
- Optimization: AdamW, learning rate 2×10^{-5} , weight decay 0.01, warmup ratio 0.06, batch sizes 16 (train) and 32 (validation), up to 20 epochs with early stopping (patience 3) on validation macro-F1. Weighted cross-entropy loss to balance classes.

Six-way stance classifier. The stance student model uses the same base encoder and input representation:

- Labels: the six stance types above.
- Training data: high-confidence Qwen stance labels (confidence ≥ 0.7) for sentences judged interrogative by the binary teacher, again excluding evaluation articles.
- Optimization: same schedule as the binary model, with class weights computed per stance label to compensate for imbalance (leading and tag are rare).

Both models were later used for two-step inference on the full sentence-level corpus, with a binary confidence threshold of 0.7 to gate sentences into the six-way stance classifier.

Table 2: Outlets in the corpus with total article counts (2023-mid-2024) and outlet ontologies.

Source	Articles	Country/region	Scale	Type
arcinfo.ch	11,204	Switzerland	Hyper-local	general
footmercato.net	35,579	France	Thematic	sports
fr.allafrica.com	116,680	Africa (pan-African)	Transnational	general
francetvinfo.fr	39,709	France	National	general
gala.fr	32,556	France	Thematic	celebrity
ici.radio-canada.ca	50,976	Canada	National	general
lacote.ch	6,145	Switzerland	Hyper-local	general
la-croix.com	19,486	France	National	general
ladepeche.fr	176,974	France	Regional	general
lapresse.ca	54,136	Canada	National	general
le10sport.com	61,404	France	Thematic	sports
lefigaro.fr	43,444	France	National	general
lelezard.com	23,594	International	Transnational	wire
lenouvelliste.ch	9,774	Switzerland	Hyper-local	general
lequipe.fr	51,389	France	Thematic	sports
lindependant.fr	60,960	France	Regional	general
midilibre.fr	89,704	France	Regional	general
news-24.fr	123,725	France	Transnational	wire
rtl.be	82,375	Belgium	National	general
seneweb.com	27,892	Senegal	National	general
tf1info.fr	42,118	France	National	general
voici.fr	32,335	France	Thematic	celebrity
watson.ch	16,127	Switzerland	National	general
zonebourse.com	30,120	Europe	Transnational	business

Answer span identification

Answer spans were identified with a simple embedding-based search using CamemBERT-large:

- Embedding model: `camembert/camembert-large`.
- Embedding input: for each sentence, an `embed_text` context consisting of the target sentence with a ± 1 -sentence window (same-article neighbors), with the target wrapped in `<tgt> ... </tgt>` and sentences separated by `" </s> "`.
- Tokenization: truncation to 256 tokens, padding to batch max length, batch size 64.
- Sentence embedding: mean pooling of the last hidden layer over non-padding tokens, followed by ℓ_2 normalization.

Within each article, we:

- Treated as *questions of interest* all sentences whose stance label was one of the six interrogative types with stance confidence ≥ 0.7 .
- Grouped consecutive question sentences (no gap in `sent_id`) into local question groups.

For each question group within an article:

1. Compute a group embedding as the average of the normalized sentence embeddings in the group, re-normalized to unit length.
2. Precompute cumulative sums over the article’s sentence embeddings to allow fast average embeddings over any contiguous window.

3. Search only among subsequent sentences up to 15 sentences ahead of the last question sentence in the group.
4. For each candidate window length $L \in \{1, 2, 3, 4, 5\}$ and each possible start position, compute the mean embedding and its cosine similarity with the group embedding.
5. If the best window has cosine similarity ≥ 0.40 , treat it as the answer span. Sensitivity checks reported in Appendix C show that the stored similarity scores are sharply bimodal, so the main answerability estimates are effectively invariant across a broad range of thresholds. Otherwise mark the group as unanswered.

For each interrogative sentence we stored whether an answer was found, the similarity score, and the sentence indices and concatenated text of the answer span. These values underlie the article-level answerability and dialogicity indices.

BERTopic configuration

Article-level topics were derived with BERTopic on CamemBERT-large article embeddings, using GPU-accelerated UMAP and HDBSCAN:

- **model** `~1M` national articles: `UMAP` `n_components=5,` `n_neighbors=50,` `min_dist=0.0,` `metric="cosine";` `HDBSCAN` `min_cluster_size=200,` `min_samples=20.`

For each article we recorded its topic ID and the maximum topic probability. Frequently occurring topics were

manually grouped into eight meta-topics used in the main analysis. For transparency, the 100 most frequent BERTopic clusters were manually assigned to the eight meta-topics by inspecting each cluster’s top keywords and representative articles. Clusters with mixed semantics were assigned according to their dominant content after joint author inspection.

Named-entity recognition with GLiNER

We applied NER to questions and their detected answer spans using the multilingual GLiNER model:

- Model: `urchade/gliner_multi-v2.1`.
- Label set (coarse and robust): PERSON, ORGANIZATION, LOCATION, NATIONALITY OR RELIGIOUS OR POLITICAL GROUP, GENERIC SOCIAL GROUP, PUBLIC OR AUDIENCE, EVENT.
- Inference settings: batch size 16 and score threshold 0.5.

NER was run on the QA-enhanced sentence files produced by the answer-identification step. We proceeded as follows:

- **Selecting questions.** We selected as *questions of interest* the same sentences used in the answer-identification step: those whose CamemBERT stance label was one of the six interrogative types and whose stance confidence was at least 0.7.
- **Question-side NER.** For each selected question, we built a short context string:
previous sentence (if any) + question sentence + next sentence (if any),
- **Answer-side NER.** For questions marked as `has_answer = True`, we passed the concatenated `answer_span_text` (without extra context) to GLiNER.
- We stored the resulting entities in columns as lists of entity records.

C Additional quantitative summaries

Table 3 provides a detailed overview of meta-topics, showing how interrogative density and the presence of organizations, locations, and events vary across article domains. Table 4 summarizes answerability and dialogicity at the stance level, including the overall shares of unanswered questions and of questions answered internally versus via quoted speech.

Sensitivity to confidence and answer-similarity thresholds

To assess the robustness of the main corpus-level findings to threshold choices, we conducted two simple sensitivity checks. First, we varied the confidence threshold used to retain interrogative predictions from the binary and six-way stance classifiers. As shown in Table 5, stricter thresholds reduce the absolute number of detected interrogatives, as expected, but the overall prevalence estimates remain in the same range: interrogatives remain sparse, with mean article-level interrogative indices between 0.024 and 0.026.

Second, we varied the cosine-similarity threshold used in answer matching. The stored similarity scores were sharply separated rather than clustered around the baseline cutoff: all unmatched cases fell below 0.35, whereas matched spans were overwhelmingly assigned much higher scores. Consequently, answerability estimates were identical for thresholds from 0.05 to 0.80 and declined only marginally at stricter cutoffs (Table 6). This indicates that the high answerability rate is not driven by a narrow threshold choice. At the same time, because the answer-matching procedure captures semantic relatedness rather than manually verified resolution, these values should still be interpreted as heuristic proxies for textual uptake rather than literal estimates of fully resolved answerhood.

Manual spot check of answer spans

To complement the threshold sensitivity analysis, we manually audited 50 randomly sampled question groups from 50 different articles, including 40 predicted answered and 10 predicted unanswered cases split across the local and national corpora. Table 7 summarizes the results.

Overall, 34 of the 40 predicted answered cases (85.0%) corresponded to clear or partial answers, and in 37 of 40 cases (92.5%) the article contained an answer somewhere, while all 10 predicted unanswered cases were confirmed as genuinely unanswered.

D Model and annotation quality

Table 8 complements the main-text summary of the performance assessment of the classification pipeline and the inter-annotator agreement.

Table 9 reports per-class precision, recall, F1, and support for the six-way stance classifier on the full gold interrogative evaluation set ($n = 516$). Figure 4 provides a complementary conditional confusion matrix computed only on gold interrogative sentences that were passed to the six-way stage by the binary detector. Because binary false negatives are excluded from this conditional view, the row-normalized diagonal values in Figure 4 should not be read as recalls and are therefore not directly comparable to the recall values reported in Table 9.

The conditional confusion matrix shows that most residual errors arise among pragmatically adjacent stance types. Information-seeking and rhetorical questions are often redistributed toward framing-procedural once a sentence has been recognized as interrogative, suggesting that framing-procedural partly functions as a broad residual category at the six-way stage. Leading questions are likewise frequently mapped to rhetorical or framing-procedural, consistent with the qualitative proximity between strongly oriented evaluation and more overtly leading design. By contrast, echo-clarification remains comparatively distinctive, while tag questions are moderately recovered but rely on a relatively small number of gold cases. More broadly, these disagreements should be interpreted in light of the intrinsic ambiguity of the task: many interrogatives plausibly support more than one stance reading, and post-hoc inspection of mismatches suggests that some model-gold discrepancies reflect

Table 3: Meta-topic overview: article volume, interrogative density, and institutional/geographic anchoring of questions.

Meta-topic	Articles	Mean interrogative index	% questions with ORG	% with LOC / EVENT
Local news	39,662	0.0298	38.3	36.3 / 28.4
Professional sports	181,056	0.0290	51.8	33.4 / 33.4
Lifestyle, entertainment & people	44,533	0.0271	27.8	25.8 / 26.1
<i>Faits divers</i>	29,932	0.0220	37.7	39.6 / 21.5
National / local politics	118,167	0.0214	47.1	35.7 / 23.8
Technology	12,104	0.0184	44.7	17.7 / 20.6
Business & economy	25,412	0.0182	57.9	29.7 / 24.8
Geopolitics	94,190	0.0180	42.9	46.7 / 31.0

Table 4: Answerability and dialogicity of interrogative stances.

Stance	<i>N</i> questions	% answered
information-seeking	153,404	97.2
echo-clarification	70,281	97.4
rhetorical	107,862	95.5
leading	16,221	95.4
tag	16,116	95.1
framing-procedural	396,298	94.8
All stances	760,182	95.6

Category	<i>N</i>	% of interrogatives
Unanswered	33,271	4.4
Answered (internal)	610,209	80.3
Answered (via quotes)	116,702	15.4

Table 5: Sensitivity of interrogative prevalence estimates to the confidence threshold used for the interrogative predictions.

Confidence	<i>N</i> questions	% of sentences	Mean ID_a
0.6	785,014	3.34	0.0261
0.7	760,182	3.23	0.0252
0.8	714,023	3.03	0.0236

borderline or multifunctional cases rather than unequivocal classification failures. We therefore use the six-way predictions as approximate corpus-level indicators and read fine-grained contrasts among adjacent stance types with appropriate caution.

E Additional qualitative observations and illustrative examples

For reasons of space, the main text can only briefly refer to individual examples from the qualitative subcorpus. This appendix presents a small set of additional excerpts that were cut from the main text but are analytically representative of broader patterns. Each example is accompanied by a short comment on why it is interesting in terms of interrogative

Table 6: Sensitivity of answerability and dialogicity estimates to the cosine-similarity threshold used for answer matching.

Similarity	answered	unanswered	internal	via quotes
0.05	95.6 %	4.4 %	80.3 %	15.4 %
0.40	95.6 %	4.4 %	80.3 %	15.4 %
0.80	95.6 %	4.4 %	80.3 %	15.4 %
0.95	94.9 %	5.1 %	79.6 %	15.3 %
0.975	82.5 %	17.5 %	69.1 %	13.5 %

Table 7: Manual spot check of the answer heuristic on 50 randomly sampled questions from 50 different articles.

Predicted answered	Local	National	Total
Clear answer	11 (55%)	14 (70%)	25 (62.5%)
Partial answer	5 (25%)	4 (20%)	9 (22.5%)
Answer elsewhere in article	1 (5%)	2 (10%)	3 (7.5%)
No genuine answer	3 (15%)	0 (0%)	3 (7.5%)

Predicted unanswered	Local	National	Total
Correctly unanswered	5 (100%)	5 (100%)	10 (100%)
Missed answer exists	0 (0%)	0 (0%)	0 (0%)

stance, macro-function, and the indices introduced in Section 3.

Framing local decline as a problem to be solved

Coverage of local socioeconomic trends sometimes uses interrogatives to recast descriptive statistics as collective problems. In a Valaisan article on the decline in apprenticeships, the lead introduces a factual decrease and then asks, in translation:

“Ten percent fewer apprentices: how can Valais fix this?”

The question is classified as *framing-procedural* but it also carries an evaluative presupposition: the decline is implicitly treated as undesirable and in need of remedy. It illustrates how factual presuppositions (*there has been a 10% drop*) can be combined with verbs like “fix” or “remedy” to move from neutral description to engaged framing, while

Table 8: Summary of model performance and inter-annotator agreement.

Binary interrogative detector	
Metric	Value
Evaluation sentences	8,399
Accuracy	0.97
Precision (interrogative)	0.76
Recall (interrogative)	0.80
F1 (interrogative)	0.78
Six-way stance classifier	
Metric	Value
Evaluation interrogatives	516
Macro-F1	0.51
Micro-F1	0.51
Inter-annotator agreement (stance)	
Metric	Value
Double-coded articles	98
Matched interrogative units	204
Jaccard overlap (spans)	0.83
Accuracy (stance labels)	0.84
Cohen's κ	0.78

Table 9: Per-class performance of the six-way stance classifier on the gold-standard evaluation set.

Stance	Precision	Recall	F1	Support
Framing-procedural	0.45	0.59	0.51	129
Information-seeking	0.80	0.37	0.51	171
Rhetorical	0.60	0.39	0.47	113
Leading	0.68	0.38	0.49	39
Tag	0.59	0.50	0.54	32
Echo-clarification	0.48	0.62	0.54	32

still remaining relatively low in face-threat and open-ended in terms of specific solutions.

Engaged framing and mild conduciveness in hyper-local politics

Hyper-local outlets sometimes use more pointed interrogatives to highlight institutional strain. In a feature on parliamentary workload in a Swiss canton, a journalist reports complaints about the number of motions and then asks:

"Are Valais deputies suffocating their own Parliament?"

"Is the parliamentary machine at risk of overload?"

These polar questions are annotated as *leading* and mapped to *Framing/agenda-setting*. They embed evaluative metaphors ("suffocating", "overload") and present a narrow choice between a problematic and a non-problematic reading of the situation. In our indices, such questions contribute to higher interrogative density and a slightly higher share

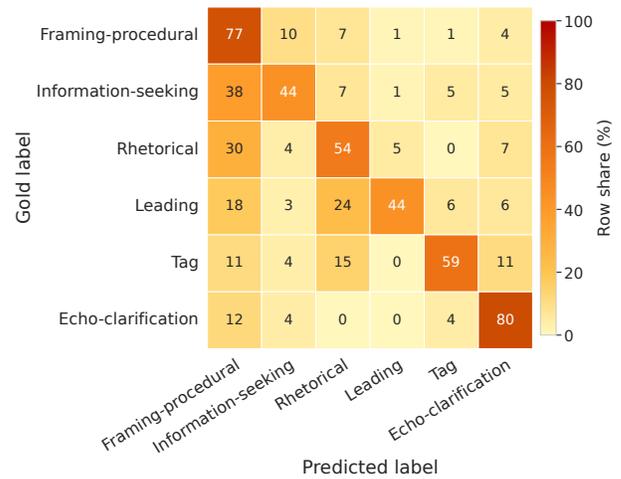


Figure 4: Row-normalized conditional confusion matrix for the six-way stance classifier.

of leading stances in hyper-local political coverage, while remaining substantially less confrontational than the adversarial formats documented in some broadcast traditions.

Authority positioning in expert interviews

In interviews with experts, interrogatives often enact an epistemic asymmetry in which the journalist adopts a knowledge-seeking role. In a piece on the Russian Orthodox Church's participation in ecumenical dialogue, after summarizing calls to sanction the Moscow Patriarchate, the interviewer turns to a scholar and asks:

"If the Russian Orthodox Church does not really want to take part in ecumenical debate, what interest does it have in participating at all?"

This is labeled *information-seeking* and mapped to the macro-axis of *Authority positioning*. The journalist's epistemic deficit is made explicit through the "what interest" formulation, while the presupposition about reluctance is attributed to prior sources rather than asserted in the journalist's own voice. The answer appears as an extended quoted response, contributing to high external dialogicity and illustrating the "epistemic extraction" mode in which questions are used to elicit explanations rather than to confront.

Personalizing cross-border relations

Some analytical pieces use interrogatives to personalize complex international relationships around a single political figure. In a Canadian analysis of U.S.–Canada relations under the prospect of a second Trump presidency, the closing paragraph asks:

"Did Donald Trump and his advisers leave Washington with a better understanding of Canada's position?"

Throughout the article, Trump is repeatedly named and foregrounded, and this final question recasts a structural diplomatic issue in terms of what one individual has or has

not learned. In our terminology, the interrogative is *framing-procedural* and contributes to a highly personalized interrogative space: a complex bilateral relationship is summarized as a question about a single actor's knowledge state.

Localization and collectivization of policy debates

Interrogatives can also localize national policy debates and link them explicitly to a specific public. In Quebec coverage of federal climate policy, after describing a proposed abolition of a federal levy, a journalist asks:

"Would abolishing this federal contribution cause Quebec's carbon market to collapse?"

This information-seeking question is also a framing move: it connects an abstract fiscal change to potential local consequences and names a specific place and collective concern (the provincial carbon market). In our NER-based indices, such cases combine LOCATION and PUBLIC/AUDIENCE mentions, illustrating the relatively small but significant share of interrogatives that directly implicate broad publics rather than only elites or institutions.

Everyday dialogicity in hyper-local features

Finally, some hyper-local feature stories include conversational question–answer pairs that foreground everyday interactions among non-elite actors. In a Swiss piece about a sewing class, the journalist recounts an exchange between a novice participant and his instructor:

*"Can I use my own sewing machine?" he asks.
"Of course, as long as it is in good working order,"
the teacher replies.*

These are straightforward *information-seeking* interrogatives realized in direct speech. They contribute to external dialogicity and to the presence of non-elite PERSON entities in both questions and answers. While numerically marginal in the corpus, such examples highlight how interrogatives can be used to stage ordinary voices and lived experience, especially in hyper-local reporting.

Taken together, these additional excerpts show how fine-grained qualitative distinctions, between neutral and engaged framing, epistemic extraction and narrative authority, personalization and localization, underlie the aggregate patterns reported in Section 4.