

ExTax: Explainable Disinformation Detection via Persuasion, Emotion, and Narrative Role Taxonomies

Shang Luo^{1*}, Yingguang Yang^{2*}, Zhenchen Sun^{3*}, Yang Liu¹, Bin Chong^{1†},
Jingru Chen⁴, Yancheng Chen⁵, Jiayu Liang⁶, Kefu Xu¹, Hao Peng⁷, Philip S. Yu⁸,

¹Peking University, ²University of Science and Technology of China,

³North China University of Science and Technology,

⁴Nanjing University of Aeronautics and Astronautics,

⁵University of Chinese Academy of Sciences,

⁶Soochow University, ⁷Beihang University, ⁸University of Illinois Chicago

Correspondence: chongbin@pku.edu.cn

Abstract

The democratization of LLMs has accelerated the generation and circulation of highly fluent disinformation, making traditional syntax-semantic verification increasingly insufficient. Such deception often combines persuasive rhetoric, emotional manipulation, and narrative role construction to shape readers' interpretations. However, existing detectors typically emphasize isolated signals—such as syntax, external knowledge, persuasion, or affective cues—and thus struggle to capture multifaceted manipulation or provide structured, taxonomy-grounded evidence. To address this gap, we present **ExTax**, a taxonomy-aligned framework for interpretable disinformation detection. ExTax unifies persuasive rhetoric, emotional manipulation, and narrative roles into a 17-dimensional taxonomic space comprising 6 persuasive-rhetoric strategies, 5 emotional-manipulation methods, and 6 narrative-role categories. It elicits attributes from multiple frontier LLMs, reconciles their disagreements through Entropy-driven Dynamic Label Smoothing, and fuses the resulting taxonomic representations with contextual encodings via Heterogeneous Multi-Head Attention. Across five cross-domain and cross-genre benchmarks, ExTax achieves an overall Macro F_1 of 0.8456, outperforming strong deep learning and LLM-based baselines. It also remains robust under severe genre imbalance, where the strongest deep baseline drops from 0.9454 to 0.6194 across genres. Our code is available at <https://github.com/owlstan49/ExTax>.

1 Introduction

Disinformation detection has become increasingly challenging as social media platforms and generative AI systems accelerate the production and circulation of deceptive yet highly fluent content.

Such content rarely operates through surface wording or outright falsity alone; instead, it often combines rhetorical strategies, emotional manipulation, and narrative framing to attack or simplify targets, evoke high-arousal emotions, and assign actors to morally charged roles. In this way, disinformation shapes public interpretation and sharing behavior through multiple cognitive routes (Saeed et al., 2022; Loth et al., 2026). Meanwhile, alarmist discourse around disinformation may itself reshape public risk perception, making reliable and interpretable detection increasingly important (Jungherr and Rauchfleisch, 2024).

Recent studies improve robustness by incorporating richer textual, contextual, or external signals. Syntax-aware models capture dependency-level distortions (Xiao et al., 2024), dynamic representation methods adapt feature views to instance-level difficulty (Farhangian et al., 2025), and knowledge- or retrieval-augmented LLM frameworks incorporate external evidence or affective cues to enable more timely detection (Jin et al., 2025; Liu et al., 2025). Intent-aware detection additionally suggests that communicative purposes provide stable signals beyond surface wording (Wang et al., 2025). However, most existing methods still model syntax, knowledge, emotion, persuasion, or pattern in isolation, leaving underexplored how multiple human-interpretable dimensions jointly reveal the manipulative structure of disinformation.

We argue that robust disinformation detection should reason over three complementary facets of manipulation: persuasive rhetoric, emotional manipulation, and narrative roles. These facets characterize which rhetorical strategy is employed, which emotional-manipulation method is activated, and which morally charged role is assigned. Grounded in persuasion-oriented taxonomies, emotional manipulation studies, as well as narrative framing research (Piskorski et al., 2023b; Jamieson et al., 2025; Zhao et al., 2024; Mahmoud et al., 2025),

*Equal contribution.

†Corresponding author.

these facets provide human-interpretable manipulation evidence that extends beyond single-cue reasoning.

Realizing such a unified taxonomy introduces two challenges: no existing corpus jointly annotates these facets, and attributes elicited from a single LLM may inherit model-specific biases. To address these issues, we propose **ExTax**, an **Explainable Taxonomy**-aligned framework for disinformation detection. ExTax elicits taxonomy-guided attributes from multiple frontier LLMs, regularizes their disagreements through entropy-driven dynamic label smoothing, and constructs a unified 17-dimensional taxonomic representation consisting of 6 persuasive-rhetoric strategies, 5 emotional-manipulation methods, and 6 narrative-role categories. A taxonomy-aligned heterogeneous attention module then integrates this representation with contextual encodings, grounding each prediction in a per-sample manipulation profile.

Across five cross-domain, cross-genre benchmarks, ExTax achieves a Macro F_1 of 0.8456, outperforming the strongest LLM and deep-learning baselines by 2.01 and 3.13 percentage points, respectively. It also remains robust under severe genre imbalance, where the strongest deep baseline degrades from 0.9454 to 0.6194 across genres.

Overall, ExTax makes three key contributions: (1) **Multifaceted taxonomic modeling**. We reframe disinformation detection as reasoning over three complementary facets of manipulation—persuasive rhetoric, emotional manipulation, and narrative roles—and instantiate these facets as a unified 17-dimensional taxonomic representation. (2) **Noise-robust attribute synthesis**. We aggregate taxonomy-guided outputs from multiple frontier LLMs and regularize noisy hard decisions through entropy-driven dynamic label smoothing, thereby transforming inter-LLM disagreement into reliable synthetic supervision. (3) **Taxonomy-aligned detection and interpretability**. We integrate learned taxonomic signals with contextual encodings through heterogeneous attention, enabling robust disinformation detection while exposing rhetorical, affective-manipulative, and narrative evidence underlying each prediction.

2 Related Work

Disinformation detection. Disinformation detection has been studied under both supervised

learning and LLM-based prompting paradigms. Supervised approaches typically rely on lexical, semantic, engagement, or contextual features to learn task-specific detectors (Aslam et al., 2021), while recent zero- and few-shot LLM methods address data scarcity and cross-domain generalization (Peltine et al., 2023; Bang et al., 2023). Prompting strategies such as zero-shot Chain-of-Thought (Lucas et al., 2023) substantially improve detection performance. However, content-level methods may still overfit to dataset artifacts and generalize poorly across genres and topics, motivating supervised detectors that integrate structured, human-interpretable stylistic, rhetorical, or contextual evidence (Modzelewski et al., 2025).

Emotion, persuasion, and narrative framing in disinformation. Communication research has established emotion as central to how audiences process misleading content. The Appraisal-Tendency Framework (Han et al., 2007) explains how discrete emotions such as anger and sadness produce distinct cognitive appraisals. Zhao et al. (2024) extended this framework to narrative disinformation, showing that anger prompts heuristic action while sadness fosters deliberation, and Paletz et al. (2023) found that discrete emotions better predict post virality than valence or arousal. Bago et al. (2022) further showed that emotional reliance increases susceptibility to fake news. In parallel, NLP research has produced persuasive-technique taxonomies (Piskorski et al., 2023a), anchoring shared tasks on propaganda and manipulation detection. Entity framing and role portrayal provide a complementary lens for analyzing how actors are positioned within news narratives (Mahmoud et al., 2025).

Integrating high-level communicative signals. Recent work integrates such signals into detection pipelines. RAEmoLLM (Liu et al., 2025) leverages emotional embeddings for cross-domain disinformation detection, while PCoT (Modzelewski et al., 2025) infuses persuasion knowledge into LLM reasoning. Kamali et al. (2024) further showed that persuasive writing strategies can serve as intermediate signals for health disinformation detection. Unlike prior work that treats emotion, persuasion, and narrative framing as separate evidence sources, ExTax jointly models them within a unified 17-dimensional taxonomic space and mitigates noise in LLM-elicited attributes through entropy-driven soft labeling.

3 Preliminaries

3.1 Problem Formulation

Disinformation detection is defined as the task of identifying false, inaccurate, or misleading information that is intentionally designed, presented, and promoted to cause public harm or generate profit (Modzelewski et al., 2026; Directorate-General for Communications Networks, Content and Technology (European Commission), 2018). Formally, this task is structured as a binary classification problem in which a model $\mathcal{F}$ evaluates an input text T , such as a news article or a social media post, to determine its credibility. Given a dataset $\mathcal{D} = \{(T_i, y_i)\}_{i=1}^n$ where T_i represents the content and $y_i \in \{0, 1\}$ is the ground-truth label, the objective is to learn a predictive function $f_{det} : \mathcal{T} \rightarrow \{0, 1\}$ such that:

$$\hat{y}_i = f_{det}(T_i), \quad (1)$$

where $\hat{y}_i = 1$ denotes that the text T_i is classified as disinformation, and $\hat{y}_i = 0$ indicates it is reliable information.

3.2 Taxonomy of Disinformation: Persuasion, Emotion, and Narrative Roles

ExTax operationalizes the three-route cognitive triangulation introduced in § 1 through three complementary taxonomic dimensions, each grounded in a distinct strand of prior literature. For **Persuasion**, we adopt the hierarchical taxonomy of Piskorski et al. (2023b), which organizes 23 fine-grained persuasive techniques into six high-level strategies: Attack on Reputation, Justification, Simplification, Distraction, Call, and Manipulative Wording. A subsequent extension by Piskorski et al. (2025a) expands the fine-grained inventory to 25 techniques. Since ExTax was developed and its taxonomy-aligned supervision was constructed based on the 2023 taxonomy, we retain the original 23-technique formulation for consistency, while noting the newer extension as a promising direction for future refinement. For **Emotion**, we consider five affective states associated with manipulative content—fear, anger, hope, anxiety, and sadness—motivated by prior findings on emotional manipulation and narrative misinformation (Jamieson et al., 2025; Zhao et al., 2024). For **Narrative Roles**, we build upon the entity-framing taxonomy of Mahmoud et al. (2025) and adopt six role categories: Ethical Stabilizers, Altruistic Catalysts, Overt Aggressors, Deceptive Subversives, Institutional Toxins, and

Marginalized Sufferers. The 6+5+6 granularity is chosen to balance representational expressiveness and annotation reliability: overly coarse categories may fail to capture meaningful manipulation patterns, whereas excessively fine-grained categories can increase ambiguity and noise in LLM-based weak annotation. We therefore adopt these high-level categories as a practical taxonomy-aligned representation, rather than claiming that this decomposition is exhaustive or uniquely optimal. Together, these three dimensions constitute a unified 17-dimensional taxonomic space. These three dimensions should not be interpreted as a complete definition of disinformation. Rather, they provide complementary taxonomy-based evidence for disinformation detection from rhetorical, affective, and entity-framing perspectives. In particular, Narrative Roles complement Persuasion and Emotion by modeling how actors are functionally and morally positioned within a narrative. We do not directly model author intent, which is difficult to infer reliably from text alone without external evidence. Detailed category definitions are provided in Appendix A.

4 Methods

The architecture of ExTax is illustrated in Figure 1, and comprises three core components: (1) Taxonomic Attribute Elicitation, which employs multiple LLMs to extract Persuasion, Emotion, and Narrative Role attributes and applies entropy-driven dynamic label smoothing to mitigate model noise; (2) Taxonomy-aligned Representation Learning, which aligns posts or news articles with a unified 17-dimensional taxonomic space; and (3) Taxonomy-based Disinformation Detection, which integrates taxonomic attributes with contextual text encodings for final prediction. Additional architectural details are provided in Appendix C.

4.1 Taxonomic Attribute Elicitation

Since no large-scale corpus jointly labels the three facets of manipulation, we construct soft taxonomic supervision by querying $N_{LLM} = 4$ top-tier LLMs and reconciling their disagreement. The four LLMs are used only offline as weak annotators to construct taxonomy-level supervision and are not queried during ExTax inference. We use deterministic decoding with temperature = 0. Detailed LLM and implementation settings are provided in Appendix D.4. Specifically, for each input




Figure 1: Overall framework of ExTax.

$T_i \in \mathcal{D}$, the Taxonomic Attribute Elicitation process extracts attributes associated with Persuasive Rhetoric, Emotional Manipulation, and Narrative Roles. This process yields an augmented dataset $\mathcal{D}_{\text{tax}} = \{(T_i, y_i, \mathbf{y}_{\text{tax}}^{(i)})\}_{i=1}^n$, where $\mathbf{y}_{\text{tax}}^{(i)} \in \mathcal{Y}_{\text{smooth}}$ is a 17-dimensional dynamically smoothed soft target vector encoding the refined presence probability of each taxonomic category in T_i .

Divergent attribute synthesis. For each text T_i , we query four frontier LLMs $\{M_k\}_{k=1}^4$ with taxonomy-specific prompts (Appendix B) to obtain three binary feature vectors over the persuasive-rhetoric ($|\mathcal{C}_P| = 6$), emotional-manipulation ($|\mathcal{C}_E| = 5$), and narrative-role ($|\mathcal{C}_N| = 6$) category sets: $\mathbf{p}_{i,k} \in \{0, 1\}^6$, $\mathbf{e}_{i,k} \in \{0, 1\}^5$, $\mathbf{n}_{i,k} \in \{0, 1\}^6$. A component is assigned a value of 1 if the LLM identifies the corresponding category in T_i . The raw outputs collected from all four models form the raw taxonomic dataset $\mathcal{R}$, on which we subsequently apply the entropy-driven smoothing procedure described below.

Entropy-driven dynamic label smoothing. To reduce noise in LLM-elicited attributes, we transform the hard decisions produced by four LLMs into entropy-aware soft labels. For notational simplicity, we use $z_{i,k}^{(c)}$ to denote the binary output of the k -th LLM for any category $c \in \mathcal{C}_P \cup \mathcal{C}_E \cup \mathcal{C}_N$.

For each sample T_i and category c , we compute the positive vote proportion:

$$p_{i,c} = \frac{1}{4} \sum_{k=1}^4 z_{i,k}^{(c)}, \quad (2)$$

where $p_{i,c}$ represents the fraction of LLMs predicting the presence of category c . We then estimate the inter-model uncertainty using binary entropy:

$$H_{i,c} = -[p_{i,c} \log_2 p_{i,c} + (1 - p_{i,c}) \log_2 (1 - p_{i,c})]. \quad (3)$$

with $0 \log_2 0 = 0$. The final smoothed taxonomic target is defined as:

$$y_{\text{tax}}^{(i,c)} = p_{i,c}(1 - \alpha H_{i,c}) + 0.5(\alpha H_{i,c}), \quad (4)$$

where α controls the smoothing strength. Thus, high-disagreement labels are pulled toward the neutral value 0.5, while high-agreement labels remain close to the original vote proportion. This procedure is applied to all categories in $\mathcal{C}_P \cup \mathcal{C}_E \cup \mathcal{C}_N$, yielding the smoothed taxonomic targets $\mathcal{Y}_{\text{smooth}}$.

4.2 Taxonomy-aligned Representation Learning

The objective of Taxonomy-aligned Representation Learning is to map input texts, including posts and news articles, into a multi-faceted taxonomic

space defined by Persuasion, Emotion, and Narrative Roles. Specifically, given the textual dataset $\mathcal{D}$ and the dynamically smoothed taxonomic labels $\mathcal{Y}_{\text{smooth}}$, the task requires the model to learn a predictive mapping from the textual domain to the corresponding taxonomic label space:

$$f_{\text{tax}} : \mathcal{T} \rightarrow \mathcal{Y}_{\text{smooth}}. \quad (5)$$

Namely, for any input news text $T_i \in \mathcal{T}$, the model yields a prediction vector $\hat{\mathbf{y}}_{\text{tax}}^{(i)}$ whose dimensionality aligns with that of $\mathcal{Y}_{\text{smooth}}$ (i.e., a 17-dimensional vector representing the combined taxonomic facets).

Adaptive gated pooling. For each input text T_i , we obtain token-level hidden states from a pre-trained LLM:

$$\mathbf{S}_i = \text{LLM}(T_i) \in \mathbb{R}^{L_i \times D}, \quad (6)$$

where L_i is the sequence length and D is the hidden dimensionality. Instead of relying on mean or first-token pooling, we introduce a learnable gating vector that adaptively assigns importance weights to individual tokens:

$$\mathbf{v}_i = \sum_{h=1}^{L_i} g_{i,h} \mathbf{S}_{i,h} \in \mathbb{R}^D, \quad (7)$$

where $\mathbf{g}_i = \text{softmax}(\mathbf{w}_{1:L_i})$ denotes the normalized token-level gating weights. This mechanism allows the model to learn a weighted aggregation of token features before taxonomic projection.

Taxonomy-aligned representation learning training. Given the pooled representation $\mathbf{v}_i \in \mathbb{R}^D$, we employ three facet-specific MLP heads, corresponding to Persuasion, Emotion, and Narrative Roles, to project the text into their respective taxonomic spaces:

$$\hat{\mathbf{p}}_i = f_P(\mathbf{v}_i), \quad \hat{\mathbf{e}}_i = f_E(\mathbf{v}_i), \quad \hat{\mathbf{n}}_i = f_N(\mathbf{v}_i), \quad (8)$$

where each head outputs sigmoid-normalized probabilities over its corresponding taxonomy categories. The model is trained using Binary Cross-Entropy (BCE) supervision against the smoothed taxonomic targets:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_P + \mathcal{L}_E + \mathcal{L}_N. \quad (9)$$

Each facet-specific loss is computed independently over its corresponding category space using the smoothed labels derived from $\mathcal{Y}_{\text{smooth}}$.

4.3 Taxonomy-based Disinformation Detection

In the Taxonomy-based Disinformation Detection phase, we integrate the taxonomic feature vector $\mathbf{t}_i = [\hat{\mathbf{p}}_i \parallel \hat{\mathbf{e}}_i \parallel \hat{\mathbf{n}}_i]$, formed by concatenating the attributes elicited from the taxonomic alignment tasks, with the text feature sequence $\mathbf{S}_i$. This fusion mechanism enables the model to leverage structured rhetorical, emotional, and narrative evidence when predicting whether a given news article constitutes disinformation.

Text feature transformation. To align contextual token features with the taxonomic space, we project each token representation into the same 17-dimensional space as the taxonomic attributes:

$$\tilde{\mathbf{S}}_{i,h} = \text{TransMLP}(\mathbf{S}_{i,h}) \in \mathbb{R}^{17}, \quad h = 1, \dots, L_i. \quad (10)$$

The resulting sequence $\tilde{\mathbf{S}}_i \in \mathbb{R}^{L_i \times 17}$ serves as the contextual feature input for heterogeneous multi-head attention.

Learnable taxonomic prompt. To enrich the semantic representations of the taxonomic dimensions, we introduce a Learnable Taxonomic Prompt mechanism. Specifically, the previously obtained taxonomic feature vector $\mathbf{t}_i \in \mathbb{R}^{17}$ is concatenated with N_{ppt} learnable semantic prompt vectors $\mathbf{q} \in \mathbb{R}^{N_{\text{ppt}} \times 17}$ of the same dimensionality, yielding an augmented representation:

$$\tilde{\mathbf{t}}_i = \text{Concat}(\mathbf{t}_i, \mathbf{q}) \in \mathbb{R}^{(1+N_{\text{ppt}}) \times 17}. \quad (11)$$

This concatenation operation enables the model to adaptively recalibrate feature importance based on both the original taxonomic attributes and the learnable prompts. Consequently, the framework can more effectively capture nuanced semantic cues.

Heterogeneous multi-head attention. To maintain a strict one-to-one correspondence with the pre-determined taxonomic spaces, we propose a Heterogeneous Multi-Head Attention mechanism. Unlike standard multi-head attention, where all heads share the same dimensionality, we configure three specialized attention heads whose dimensionalities correspond directly to the number of categories in the three taxonomic facets: $d_1 = |\mathcal{C}_P| = 6$, $d_2 = |\mathcal{C}_E| = 5$, and $d_3 = |\mathcal{C}_N| = 6$.

Formally, given the learnable taxonomic prompt sequence $\tilde{\mathbf{t}}_i$ and the transformed token feature sequence $\tilde{\mathbf{S}}_i$, we project them into the query, key, and

Method	Overall		Post		Article	
	MacroF1	MacroRecall	MacroF1	MacroRecall	MacroF1	MacroRecall
RAEmoLLM	0.8255 $\pm$.002	0.7981 $\pm$.002	0.8625 $\pm$.001	0.8528 $\pm$.001	0.7978 $\pm$.002	0.7704 $\pm$.002
FACTUAL	0.7563 $\pm$.004	0.7377 $\pm$.004	0.6428 $\pm$.006	0.6474 $\pm$.005	<u>0.8027</u> $\pm$.004	0.7780 $\pm$.004
DYNAMO	0.6680 $\pm$.007	0.6972 $\pm$.009	0.5737 $\pm$.020	0.6200 $\pm$.019	0.6936 $\pm$.002	0.7177 $\pm$.015
PCoT	0.8218 $\pm$.020	0.8053 $\pm$.019	0.7675 $\pm$.007	0.7606 $\pm$.007	0.8439 $\pm$.027	0.8226 $\pm$.025
RoBERTa-large	0.8143 $\pm$.028	0.8085 $\pm$.038	0.9270 $\pm$.016	0.9300 $\pm$.014	0.7583 $\pm$.042	0.7517 $\pm$.053
DRES	0.7294 $\pm$.001	0.7170 $\pm$.001	<u>0.9454</u> $\pm$.002	<u>0.9418</u> $\pm$.002	0.6194 $\pm$.002	0.6145 $\pm$.002
MSynFD	0.7337 $\pm$.026	0.7206 $\pm$.026	0.8887 $\pm$.006	0.8924 $\pm$.003	0.6375 $\pm$.051	0.6356 $\pm$.039
ExTax (ours)	0.8456 $\pm$.011	0.8423 $\pm$.015	0.9548 $\pm$.003	0.9565 $\pm$.004	0.7965 $\pm$.018	<u>0.7913</u> $\pm$.022

Table 1: Performance comparison between baseline models and our framework. The methods above the horizontal dividing line represent LLM-based approaches, while those below are deep learning-based methods. The best and second-best results are highlighted in **bold** and underlined, respectively. We report the mean and standard deviation of the Macro F_1 -score and Macro recall over all random seeds. The improvement of ExTax over the strongest competing baseline, RAEmoLLM, is statistically significant under McNemar’s test ($p = 0.0193 < 0.05$).

value spaces of the k -th taxonomic head: $\mathbf{Q}_k \in \mathbb{R}^{(1+N_{ppt}) \times d_k}$, $\mathbf{K}_k \in \mathbb{R}^{L_i \times d_k}$, and $\mathbf{V}_k \in \mathbb{R}^{L_i \times d_k}$, where $d_k \in \{|\mathcal{C}_P|, |\mathcal{C}_E|, |\mathcal{C}_N|\}$. The heterogeneous attention output is computed as:

$$\mathbf{H}_{i,k} = \text{softmax} \left(\frac{\mathbf{Q}_k \mathbf{K}_k^\top}{\sqrt{d_k}} \right) \mathbf{V}_k, \quad (12)$$

$$\mathbf{H}_i^{\text{het}} = \text{Concat}(\mathbf{H}_{i,1}, \mathbf{H}_{i,2}, \mathbf{H}_{i,3}),$$

where $\mathbf{H}_{i,k}$ denotes the output of the k -th attention head for instance i . The concatenated representation $\mathbf{H}_i^{\text{het}} \in \mathbb{R}^{(1+N_{ppt}) \times 17}$ combines the three facet-specific outputs along the feature dimension, thereby preserving alignment with the unified 17-dimensional taxonomic space.

This design ensures that each attention head is dedicated to a single taxonomic dimension, preserving rigorous alignment with persuasion strategies, emotional manipulation, and narrative roles.

The heterogeneous multi-head attention module comprises a stack of N_{att} layers. Within each layer, the input representation first undergoes heterogeneous multi-head attention, followed by a residual connection and layer normalization.

Prediction. The classification head maps the final attention output $\mathbf{h}_i^{(N_{\text{att}})} \in \mathbb{R}^{(N_{\text{ppt}}+1) \times 17}$ into the target prediction space. Specifically, a Weighted Sum Queries layer first aggregates the representation across the query dimension to produce a unified 17-dimensional taxonomic vector, which is subsequently projected by a compact MLP layer to

generate the final classification:

$$\mathbf{h}_i^{(\text{sum})} = \sum_{j=1}^{N_{\text{ppt}}+1} \beta_j \cdot \mathbf{h}_{i,j}^{(N_{\text{att}})}, \quad (13)$$

$$\hat{\mathbf{y}}_i = \text{Linear}(\text{LN}(\text{GELU}(\text{Linear}(\mathbf{h}_i^{(\text{sum})})))), \quad (14)$$

where β_j denotes the learnable query weights, and $\hat{\mathbf{y}}$ represents the final prediction vector.

5 Experiments

5.1 Setup

Dataset. We adopt the five datasets used in PCoT (Modzelewski et al., 2025): **CoAID** (Cui and Lee, 2020) and **ECTF** (Paka et al., 2021; Bansal et al., 2021) cover COVID-19 health disinformation in articles and tweets, respectively, and provide train/validation/test splits; **ISOT Fake News** (Ahmed et al., 2018, 2017), **MultiDis** (Modzelewski et al., 2025), and **EUDisinfo** (Modzelewski et al., 2025) are test-only article corpora spanning mainstream news, multi-topic European political/social disinformation, and pro-Kremlin geopolitical narratives. Full specifications are provided in Appendix D.1.

Baseline. We compare against seven representative baselines, including three deep learning-based methods, **RoBERTa** (Liu et al., 2019), **DRES** (Farhangian et al., 2025), and **MSynFD** (Xiao et al., 2024), and four LLM-based methods, **RAEmoLLM** (Liu et al., 2025), **FACTUAL** (Li et al., 2025), **DYNAMO** (Jin et al., 2025), and **PCoT** (Modzelewski et al., 2025). Their introductions and implementation details are provided in Appendix D.2 and Appendix D.5, respectively.

Metrics. To quantitatively evaluate our framework, we employ three standard metrics: **Accuracy**, **Macro Recall** and **Macro F_1 -score**. For a more detailed definition, please refer to Appendix D.

Implementation. ExTax uses a frozen roberta-large backbone with parameter-efficient trainable components (≈ 267 K parameters), optimized in a two-phase architecture across three random seeds. All deep learning baselines and ExTax are trained and validated on the CoAID and ECTF training/validation partitions, then evaluated across all five test sets. Detailed implementation and hyperparameter configurations are provided in Appendix D.4.

5.2 Main Results

We categorize all five datasets into *Post* and *Article* genres based on text typology. Table 1 reports the main results, from which we draw three key observations. (i) **Best Overall performance.** ExTax achieves the highest Overall MacroF1 (**0.8456**) and MacroRecall (**0.8423**), improving by +2.01 and +4.42 percentage points over the strongest LLM-based RAG baseline (RAEmoLLM), and by +3.13 and +3.38 percentage points over the strongest deep-learning baseline (RoBERTa-large). As noted by Modzelewski et al. (2025), prompting-based baselines may further benefit from potential pre-training exposure, an advantage that ExTax does not require. (ii) **Genre trade-off rather than uniform dominance.** The Overall margin is mainly driven by short-form *Posts*, where ExTax reaches **0.9548** MacroF1 and outperforms PCoT by +18.7 percentage points. On long-form *Articles*, ExTax is the runner-up with **0.7965** MacroF1, behind PCoT (**0.8439**), which may be attributable to the 30% *Article* minority in training and potential pre-training exposure favoring article-style text for prompting baselines. Crucially, PCoT pays for this *Article* advantage by collapsing on *Posts* (**0.7675**), whereas ExTax does not exhibit a genre-collapse failure mode. (iii) **Robustness under severe genre imbalance.** Despite the 70/30 genre imbalance, ExTax generalizes consistently across both genres, while conventional deep-learning baselines exhibit substantial performance degradation. For example, DRES’s MacroF1 drops from **0.9454** on *Posts* to **0.6194** on *Articles* (−32.6 percentage points), suggesting that taxonomy-aligned features capture more genre-invariant manipulation patterns,

LLM	stage 1	stage 2
	Avg. F1	Overall
GPT-5 mini	0.5026	0.8260
DeepSeek-V4-Flash	0.3234	0.8216
Grok 4.1 Fast Reasoning	0.2571	0.8384
Qwen3.6 Flash	0.4286	<u>0.8387</u>
ours (hybrid)	<u>0.4489</u>	0.8608

Table 2: Performance comparison (Macro F_1) of our hybrid approach (Divergent Attribute Synthesis and Entropy-driven Dynamic Label Smoothing) against single LLM attribute synthesis across Stage 1 (Taxonomy-aligned Representation Learning) and Stage 2 (Taxonomy-based Disinformation Detection). Stage 1 reports the average Macro F_1 across the three taxonomic dimensions.

Conf.	Acc _{fake}	Acc _{real}	MacroF1
Baseline	0.5952	0.9313	0.7789
Hetero. Attention	0.6041	0.9415	0.7898
Ada. Gated Pooling	0.6789	0.8973	0.7966
hybrid	0.7132	0.9001	0.8137
$N_{\text{ppt}} = 1$	0.5165	0.9551	0.7541
$N_{\text{ppt}} = 2$	<u>0.7512</u>	<u>0.9517</u>	0.8644
$N_{\text{ppt}} = 3(\text{Full})$	0.8058	0.8823	<u>0.8608</u>

Table 3: Ablation study of ExTax components. The baseline configuration utilizes RoBERTa paired with an MLP. Random seed is 445. Components are incrementally integrated from top to bottom until the full ExTax framework is achieved.

whereas surface-pattern detectors may overfit to genre-specific artifacts.

5.3 Ablation

Effectiveness of hybrid taxonomic supervision.

As shown in Table 2, our hybrid Taxonomic Attribute Elicitation framework does not achieve the best Stage-1 performance in taxonomy-aligned representation learning, but yields the strongest downstream Macro F_1 in Stage 2. This is by design: Stage 1 favors hard-label fidelity to synthesized annotations, whereas entropy-driven dynamic label smoothing deliberately trades such fidelity for calibrated uncertainty. When multiple LLMs disagree, supervision is softened to prevent overfitting to uncertain attributes. As a result, Stage 2 receives better-calibrated taxonomic supervision, improving robustness against synthesized attribute noise and leading to stronger disinformation detection.




Figure 2: Distribution of persuasion, emotion, and narrative role taxonomic attributes across all datasets, where distributions for disinformation and real news are highlighted in red and green, respectively.




Figure 3: Sankey diagram illustrating the co-occurrence patterns across the three taxonomic dimensions (Persuasion, Emotion, and Narrative Roles) and the final ground-truth labels. The flows reveal which manipulative attributes tend to appear together. "None" denotes instances where the LLMs did not detect any of the predefined categories within a specific taxonomic facet.

Ablation of model design choices. Table 3 details the incremental contributions of ExTax’s core components over the baseline (0.7789 macF1). Introducing heterogeneous multi-head attention effectively integrates taxonomic representations with contextual text features, improving performance to 0.7898 macF1. Adding adaptive gated pooling dynamically captures key semantic tokens, boosting fake news detection accuracy (Acc_{fake}) to 0.6789. Further integration of the hybrid module (divergent attribute synthesis and entropy-driven label smoothing) filters synthetic feature noise, increasing the macF1 score to 0.8137.

Regarding prompt scaling, although $N_{\text{ppt}} = 2$ achieves the highest overall macF1 (0.8644), it suffers from severe classification bias toward the real label ($\text{Acc}_{\text{real}} = 0.9517$ vs. $\text{Acc}_{\text{fake}} = 0.7512$).

Consequently, ExTax adopts $N_{\text{ppt}} = 3$ to achieve more balanced diagnostic performance between fake ($\text{Acc}_{\text{fake}} = 0.8058$) and real ($\text{Acc}_{\text{real}} = 0.8823$) news, while maintaining a competitive Macro F1 of 0.8608.

Tables 4 and 5 further validate the main design choices. Adaptive Gated Pooling achieves the best Overall and Article Macro F1, while combining textual and taxonomic representations outperforms either source alone. Removing any taxonomy facet degrades performance, with Narrative Roles producing the largest drop. Entropy Smoothing also outperforms both Single-LLM supervision and Multi-LLM Hard Voting, supporting its effectiveness in handling inter-LLM disagreement.

External validation on human annotations. On independent human-annotated benchmarks, the Per-

Method	Overall	Post	Article
<i>Pooling Strategy</i>			
CLS Pooling	81.97	95.40	76.05
Mean Pooling	83.54	96.37	78.14
Adaptive Gated Pooling	86.08	95.22	82.01
<i>Representation Source</i>			
Text-only (RoBERTa-large)	81.43	92.70	75.83
Taxonomy-only	78.05	76.89	78.49
Text + Taxonomy (Full)	86.08	95.22	82.01

Table 4: Architecture ablation of ExTax in terms of pooling strategy and representation source. Results are reported in Macro F1 (%). The text-only variant corresponds to the RoBERTa-large baseline.

Method	Overall	Post	Article
<i>Taxonomy Facet Ablation</i>			
w/o Persuasion	79.44	94.57	71.17
w/o Narrative Role	77.88	94.71	68.00
w/o Emotion	85.20	95.22	80.70
All Features	86.08	95.22	82.01
<i>Weak-supervision Strategy</i>			
Single LLM	84.18	96.20	78.20
Multi-LLM Hard Vote	81.86	94.67	75.24
Entropy Smoothing	86.08	95.22	82.01

Table 5: Ablation of taxonomy facets and weak-supervision strategies. Results are reported in Macro F1 (%).

suation and Narrative Role heads achieve Macro AUCs of 0.720 and 0.613 on SemEval-2023 Task 3 (Piskorski et al., 2023a) and SemEval-2025 Task 10 (Piskorski et al., 2025b), respectively, without retraining. These results provide external evidence that the learned representations retain meaningful alignment with human annotations. Detailed results are provided in Appendix D.6.

5.4 Interpretability

Figure 2 visualizes the global distribution of ExTax’s three taxonomic dimensions across all datasets. Disinformation and real news exhibit clearly distinct patterns within the 17-dimensional taxonomic space: fake news is more strongly associated with “Attack on Reputation”, “Simplification”, and “Manipulative Wording”, high-arousal emotions such as “Anger” and “Fear”, and negative narrative roles such as “Deceptive Subversives” and “Institutional Toxins”. These observations are consistent with prior findings on persuasion and emotional manipulation in disinformation (Modzelewski et al., 2025; Jamieson et al., 2025; Zhao et al., 2024).

Figure 3 further illustrates the co-occurrence

structure among persuasion, emotion, narrative roles, and final labels. Aggressive persuasion frequently co-occurs with negative emotions, whereas non-manipulative attributes tend to propagate across dimensions and are more strongly associated with real news. These results suggest that ExTax captures interpretable manipulative patterns underlying deceptive texts. The full interpretability analysis and detailed case studies are provided in Appendix D.7.

6 Conclusion

We introduced ExTax, a taxonomy-aligned framework for disinformation detection that jointly models persuasion, emotion, and narrative roles in a unified 17-dimensional space. ExTax combines multi-LLM weak supervision, entropy-driven smoothing, and heterogeneous attention, achieving strong cross-domain and cross-genre performance. Its structured taxonomy-grounded profiles additionally make model predictions more inspectable than binary fake/real outputs.

Limitations

Several further limitations are worth noting. (1) **Reliance on commercial frontier LLMs.** The divergent attribute elicitation stage relies on four commercial frontier LLMs, which incurs non-trivial inference cost and raises reproducibility concerns as those APIs evolve. Although entropy-driven smoothing mitigates inconsistency across models, it cannot correct systematic biases shared by all of them. (2) **English- and Western-centric corpora.** Our five evaluation datasets are predominantly English and focused on Western political and health discourse; thus, we do not validate ExTax on multilingual or non-Western disinformation, where persuasion conventions and narrative archetypes may differ. (3) **Explanations are not yet evaluated with humans.** While ExTax outputs structured per-sample 17-dimensional manipulation profiles, we have not conducted a user study with fact-checkers or moderators to verify whether these profiles are more useful in practice than opaque verdicts. We leave controlled human evaluation to future work.

Although ExTax is designed to support explainable disinformation detection, it may still produce false positives or false negatives, especially on long-form articles or domains underrepresented in the training data. False positives could unfairly flag legitimate content as disinformation, while false

negatives could allow harmful misleading content to remain undetected. In addition, the LLM-elicited taxonomic attributes may reflect biases or inconsistencies from the underlying LLMs, even though our entropy-driven smoothing mechanism is designed to mitigate synthetic labeling noise. Therefore, Ex-Tax should be used as a decision-support tool rather than as an automated content moderation system, and its outputs should be interpreted with human oversight.

Acknowledgements

This work was supported in part by the Science Research Project of Hebei Higher Education Institutions under Grant CYZD2026005, the Shijiazhuang Science and Technology Plan Project under Grant 2511301807A, and the National Natural Science Foundation of China under Grant 7240041225.

References

- Hadeer Ahmed, Issa Traore, and Sherif Saad. 2017. [Detection of online fake news using n-gram analysis and machine learning techniques](#). In *Intelligent, Secure, and Dependable Systems in Distributed and Cloud Environments*, pages 127–138, Cham. Springer International Publishing.
- Hadeer Ahmed, Issa Traore, and Sherif Saad. 2018. [Detecting opinion spams and fake news using text classification](#). *SECURITY AND PRIVACY*, 1(1):e9.
- Nida Aslam, Irfan Ullah Khan, Farah Salem Alotaibi, Lama Abdulaziz Aldaej, and Asma Khaled Al-dubaikil. 2021. [Fake Detect: A Deep Learning Ensemble Model for Fake News Detection](#). *Complexity*, 2021:5557784.
- Bence Bago, Leah R. Rosenzweig, Adam J. Berinsky, and David G. Rand. 2022. [Emotion may predict susceptibility to fake news but emotion regulation does not seem to help](#). *Cognition and Emotion*, 36(6):1166–1180.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- Rachit Bansal, William Scott Paka, Nidhi, Shubhashis Sengupta, and Tanmoy Chakraborty. 2021. [Combining exogenous and endogenous signals with a semi-supervised co-attention network for early detection of covid-19 fake tweets](#). In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*.
- Limeng Cui and Dongwon Lee. 2020. [Coaid: Covid-19 healthcare misinformation dataset](#). *Preprint*, arXiv:2006.00885.
- Directorate-General for Communications Networks, Content and Technology (European Commission). 2018. *A Multi-Dimensional Approach to Disinformation: Report of the Independent High Level Group on Fake News and Online Disinformation*. Publications Office of the European Union.
- Faramarz Farhangian, Leandro Augusto Ensina, George D C Cavalcanti, and Rafael M. O. Cruz. 2025. [DRES: Fake news detection by dynamic representation and ensemble selection](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20023–20041, Suzhou, China. Association for Computational Linguistics.
- Seunghye Han, Jennifer S. Lerner, and Dacher Keltner. 2007. [Feelings and Consumer Decision Making: The Appraisal-Tendency Framework](#). *Journal of Consumer Psychology*, 17(3):158–168.
- Jack Jamieson, Toru Hara, and Mitsuaki Akiyama. 2025. [Flagging Emotional Manipulation: Impacts of Manipulative Content Warnings on Sharing Intentions and Perceptions of Health-Related Social Media Posts](#). In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '25, pages 1–9, New York, NY, USA. Association for Computing Machinery.
- Di Jin, Jun Yang, Xiaobao Wang, Junwei Zhang, Shuqi Li, and Dongxiao He. 2025. [A dynamic knowledge update-driven model with large language models for fake news detection](#). In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 3000–3008. International Joint Conferences on Artificial Intelligence Organization.
- Andreas Jungherr and Adrian Rauchfleisch. 2024. [Negative Downstream Effects of Alarmist Disinformation Discourse: Evidence from the United States](#). *Political Behavior*, 46(4):2123–2143.
- Danial Kamali, Joseph D. Romain, Huiyi Liu, Wei Peng, Jingbo Meng, and Parisa Kordjamshidi. 2024. [Using Persuasive Writing Strategies to Explain and Detect Health Misinformation](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 17285–17309, Torino, Italia. ELRA and ICCL.
- Ang Li, Jingqian Zhao, Bin Liang, Lin Gui, Hui Wang, Xi Zeng, Xingwei Liang, Kam-Fai Wong, and Ruifeng Xu. 2025. [Mitigating Biases of Large Language Models in Stance Detection with Counterfactual Augmented Calibration](#). In *Proceedings of*

- the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), pages 7075–7092, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Zhiwei Liu, Kailai Yang, Qianqian Xie, Christine de Kock, Sophia Ananiadou, and Eduard Hovy. 2025. [RAEmoLLM: Retrieval Augmented LLMs for Cross-Domain Misinformation Detection Using In-Context Learning Based on Emotional Information](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16508–16523, Vienna, Austria. Association for Computational Linguistics.
- Alexander Loth, Martin Kappes, and Marc-Oliver Pahl. 2026. [Industrialized deception: The collateral effects of llm-generated misinformation on digital ecosystems](#). In *Companion Proceedings of the ACM Web Conference 2026, WWW Companion '26*, page 295–302, New York, NY, USA. Association for Computing Machinery.
- Jason Lucas, Adaku Uchendu, Michiharu Yamashita, Jooyoung Lee, Shaurya Rohatgi, and Dongwon Lee. 2023. [Fighting Fire with Fire: The Dual Role of LLMs in Crafting and Detecting Elusive Disinformation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14279–14305, Singapore. Association for Computational Linguistics.
- Tarek Mahmoud, Zhuohan Xie, Dimitar Iliyanov Dimitrov, Nikolaos Nikolaidis, Purificação Silvano, Roman Yangarber, Shivam Sharma, Elisa Sartori, Nicolas Stefanovitch, Giovanni Da San Martino, Jakub Piskorski, and Preslav Nakov. 2025. [Entity framing and role portrayal in the news](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 302–326, Vienna, Austria. Association for Computational Linguistics.
- Arkadiusz Modzelewski, Witold Sosnowski, Tiziano Labruna, Adam Wierzbicki, and Giovanni Da San Martino. 2025. [PCoT: Persuasion-Augmented Chain of Thought for Detecting Fake News and Social Media Disinformation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 24959–24983, Vienna, Austria. Association for Computational Linguistics.
- Arkadiusz Modzelewski, Witold Sosnowski, Eleni Papadopulos, Elisa Sartori, Tiziano Labruna, Giovanni Da San Martino, and Adam Wierzbicki. 2026. [MA-Licious INTent dataset and inoculating LLMs for enhanced disinformation detection](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3125–3148, Rabat, Morocco. Association for Computational Linguistics.
- William Scott Paka, Rachit Bansal, Abhay Kaushik, Shubhashis Sengupta, and Tanmoy Chakraborty. 2021. [Cross-SEAN: A cross-stitch semi-supervised neural attention model for covid-19 fake news detection](#). *Applied Soft Computing*, 107:107393 – 107393.
- Susannah B. F. Paletz, Michael A. Johns, Egle E. Murauskaite, Ewa M. Golonka, Nick B. Pandža, C. Anton Rytting, Cody Buntain, and Devin Ellis. 2023. [Emotional content and sharing on Facebook: A theory cage match](#). *Science Advances*, 9(39):eade9231.
- Kellin Pelrine, Anne Imouza, Camille Thibault, Meilina Reksoprodjo, Caleb Gupta, Joel Christoph, Jean-François Godbout, and Reihaneh Rabbany. 2023. [Towards reliable misinformation mitigation: Generalization, uncertainty, and GPT-4](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6399–6429, Singapore. Association for Computational Linguistics.
- Jakub Piskorski, Dimitar Dimitrov, Filip Dobranić, Marina Ernst, Jacek Haneczok, Ivan Koychev, Nikola Ljubešić, Michal Marcinczuk, Arkadiusz Modzelewski, Ivo Moravski, and Roman Yangarber. 2025a. [SlavicNLP 2025 shared task: Detection and classification of persuasion techniques in parliamentary debates and social media](#). In *Proceedings of the 10th Workshop on Slavic Natural Language Processing (Slavic NLP 2025)*, pages 254–275, Vienna, Austria. Association for Computational Linguistics.
- Jakub Piskorski, Tarek Mahmoud, Nikolaos Nikolaidis, Ricardo Campos, Alipio Mario Jorge, Dimitar Dimitrov, Purificação Silvano, Roman Yangarber, Shivam Sharma, Tanmoy Chakraborty, Nuno Guimaraes, Elisa Sartori, Nicolas Stefanovitch, Zhuohan Xie, Preslav Nakov, and Giovanni Da San Martino. 2025b. [SemEval 2025 task 10: Multilingual characterization and extraction of narratives from online news](#). In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*, pages 2610–2643, Vienna, Austria. Association for Computational Linguistics.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023a. [SemEval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multi-lingual setup](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361, Toronto, Canada. Association for Computational Linguistics.
- Jakub Piskorski, Nicolas Stefanovitch, Nikolaos Nikolaidis, Giovanni Da San Martino, and Preslav Nakov. 2023b. [Multilingual Multifaceted Understanding of Online News in Terms of Genre, Framing, and Persuasion Techniques](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational*

Linguistics (Volume 1: Long Papers), pages 3001–3022, Toronto, Canada. Association for Computational Linguistics.

Mohammed Saeed, Nicolas Traub, Maelle Nicolas, Gianluca Demartini, and Paolo Papotti. 2022. [Crowd-sourced fact-checking at twitter: How does the crowd compare with experts?](#) *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*.

Zhengjia Wang, Qiang Sheng, Danding Wang, Beizhe Hu, and Juan Cao. 2025. [Bridging thoughts and words: Graph-based intent-semantic joint learning for fake news detection](#). *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*.

Liang Xiao, Qi Zhang, Chongyang Shi, Shoujin Wang, Usman Naseem, and Liang Hu. 2024. [Msynfd: Multi-hop syntax aware fake news detection](#). *Proceedings of the ACM Web Conference 2024*.

Xinyan Zhao, Jessica Shaw, and Zexin Ma. 2024. [How individuals cope with anger- and sadness-induced narrative misinformation on social media: Roles of transportation and correction](#). *Online Media and Global Communication*, 3(3):418–446.

Appendix

A Taxonomies

This appendix provides the detailed definitions of the three taxonomic dimensions used in ExTax: Persuasion, Emotion, and Narrative Roles. These taxonomies are used during the Taxonomic Attribute Elicitation stage to guide LLMs in extracting interpretable textual attributes from each input text. To further clarify the semantic scope of the 17-dimensional taxonomy, we additionally provide structured definitions and representative examples for all categories in Tables 6–8. The definitions are derived from the prompt configuration used for attribute elicitation, while prompt-formatting instructions and JSON output templates are reported separately in Appendix B.

For Persuasion, ExTax follows the hierarchical taxonomy of Piskorski et al. (2023b), which organizes 23 fine-grained persuasion techniques into six high-level categories. We retain this formulation because it is the taxonomy used to construct the weak supervision and train the Persuasion head throughout our experiments. We note that a subsequent extension by Piskorski et al. (2025a) expands the fine-grained inventory to 25 techniques. Incorporating this updated inventory would require reconstructing the corresponding taxonomic supervision, and we therefore leave it for future work.

A.1 Persuasion Taxonomy Used in ExTax

Persuasion refers to the use of language intended to influence readers. We distinguish six high-level persuasion approaches: Attack on Reputation, Justification, Simplification, Distraction, Call, and Manipulative Wording.

A.1.1 Attack on Reputation

Attack on Reputation refers to arguments that do not directly address the topic itself, but instead target a participant, group, organization, object, or activity in order to question or undermine its credibility.

- **Name Calling or Labelling:** The use of loaded labels directed at an individual, group, object, or activity, typically in an insulting or demeaning way, although sometimes using labels that the target audience may find desirable.
- **Guilt by Association:** Attacking an opponent, activity, or concept by associating it with another group, activity, or concept that carries strong negative connotations for the target audience.
- **Casting Doubt:** Questioning the character or personal attributes of a person or entity in order to undermine its general credibility or perceived quality.
- **Appeal to Hypocrisy:** Attacking a target’s reputation by accusing it of hypocrisy or inconsistency.
- **Questioning the Reputation:** Making strong negative claims about a target, with emphasis on undermining its character or moral standing rather than addressing the substantive topic.

A.1.2 Justification

Justification refers to arguments composed of a statement and an accompanying explanation or appeal, where the latter is used to justify or support the former.

- **Flag Waving:** Justifying an idea by invoking pride in a group or highlighting benefits for that group.
- **Appeal to Authority:** Giving weight to an argument, idea, or piece of information by presenting an entity regarded as authoritative as its source.

Category	Sub-techniques	Definition	Example
Attack on Reputation	Name Calling/Labelling, Guilt by Association, Casting Doubt, Appeal to Hypocrisy, Questioning the Reputation	The argument targets the participant (person, group, or organization) rather than the topic itself, in order to undermine their credibility.	<i>"But then, Trump is notorious for appointing like-minded camp-followers such as Johnson to important ambassadorial jobs."</i>
Justification	Flag Waving, Appeal to Authority/Popularity/Values/Fear-Prejudice	A statement is supported by an explanation or appeal (to authority, values, fear, etc.) that justifies it.	<i>"Post-Brexit Britain should use its power to defend democracy around the world, Jeremy Hunt will say today."</i>
Simplification	Causal/Consequential Oversimplification, False Dilemma	The argument excessively simplifies a problem's cause, consequence, or the existence of choices.	<i>"There is no U.K. majority for any of these options, but if there is no agreement, then the default will be a no-deal Brexit."</i>
Distraction	Strawman, Red Herring, Whataboutism	The argument shifts focus away from the main topic to distract the reader.	<i>"I often use the example of an iPhone ... which is manufactured in China, uses American technology and these are two countries we deal with on WTO terms..." (diverting a Brexit-deal debate to an unrelated anecdote)</i>
Call	Slogans, Conversation Killer, Appeal to Time	Not an argument, but an encouragement to act or think in a particular way.	<i>"The words 'London is open' rang in the new year as the capital welcomed 2019 with a dazzling riverside fireworks display."</i>
Manipulative Wording	Loaded Language, Repetition, Exaggeration/Minimisation, Obfuscation	Non-neutral, confusing, exaggerating, or emotionally loaded wording used to influence the reader.	<i>"...her only chance is if Labour decides that they want to dishonour democracy and effectively keep us in the EU."</i>

Table 6: Definitions, fine-grained sub-techniques, and representative examples of the six Persuasion categories used by ExTax.

- **Appeal to Popularity:** Justifying an argument or idea by claiming that everyone, or the large majority, agrees with it, or that nobody disagrees with it.
- **Appeal to Values:** Supporting an idea by linking it to values that are viewed positively by the target audience.
- **Appeal to Fear or Prejudice:** Promoting or rejecting an idea by invoking fear, prejudice, or repulsion toward that idea.

A.1.3 Simplification

Simplification refers to arguments that excessively simplify a complex issue, usually by reducing its causes, consequences, or available choices.

- **Causal Oversimplification:** Assuming a single cause or reason for an issue when multiple causes are involved.
- **False Dilemma or No Choice:** Presenting only two options or sides when more alternatives exist, or explicitly instructing the audience to take one action while excluding other possible choices.
- **Consequential Oversimplification:** Claiming that an initial event or action will lead to a domino-like chain of increasingly serious consequences, where the causal chain is exaggerated, improbable, or insufficiently warranted.

A.1.4 Distraction

Distraction refers to arguments that divert attention away from the main issue or argument.

- **Strawman:** Giving the impression of refuting an opponent's argument while replacing the actual argument with a distorted or weaker version.
- **Red Herring:** Diverting the audience's attention from the main topic by introducing an irrelevant issue.
- **Whataboutism:** Attempting to discredit an opponent's position by accusing the opponent of hypocrisy without directly addressing or disproving the original argument.

A.1.5 Call

Call refers to language that is not primarily argumentative, but instead encourages the audience to act or think in a particular way.

- **Slogans:** Brief and striking phrases, often functioning as emotional appeals and potentially involving labelling or stereotyping.
- **Conversation Killer:** Words or phrases that discourage critical thinking or meaningful discussion about a given topic.
- **Appeal to Time:** Arguments centered on the claim that the time has come for a particular action.

Category	Sub-roles	Definition	Example
Ethical Stabilizers	Guardian, Peacemaker, Virtuous	Actors who use established values, moral authority, or mediation to maintain cohesion, safety, and justice.	<i>"Ukraine holds back the Russian troops due to robust defenses... the Ukrainian troops are 'capable and trained'." (Ukraine → Virtuous)</i>
Altruistic Catalysts	Rebel, Martyr, Underdog	Actors who challenge the status quo or endure hardship to drive change, often at personal cost.	<i>"Oxford residents are taking matters into their own hands and destroying the street zone sectioning barriers." (Oxford residents → Rebel)</i>
Overt Aggressors	Instigator, Tyrant, Terrorist, Bigot	Visible antagonists using direct force, systemic power, or violent ideology to dominate or divide.	<i>"Those 'Extinction Rebellion' activists everyone loves to hate are at it again. After multiple instances of vandalism against famous landmarks..." (Extinction Rebellion → Terrorist)</i>
Deceptive Subversives	Conspirator, Spy, Traitor, Deceiver, Saboteur, Foreign Adversary	Hidden actors operating through secrecy, betrayal, or manipulation to undermine trust.	<i>"According to top Kremlin propagandists, it's only a matter of time before the West admits its defeat and pays reparations to Moscow." (Kremlin propagandists → Deceiver)</i>
Institutional Toxins	Corrupt, Incompetent	Internal failures of an organization or system caused by human vice or incapacity, eroding it from within.	<i>"Washington 'cannot afford' to allow Russia to achieve victory in the Ukraine conflict as this would mean losing direct access to vast mineral assets." (Washington → Corrupt)</i>
Marginalized Sufferers	Forgotten, Exploited, Victim, Scapegoat	Entities enduring harm, exploitation, or unjust exclusion, serving as evidence of systemic failure.	<i>"While Garcetti was away, Los Angeles continued to struggle with the high crime and homelessness that have come to define his legacy." (Los Angeles → Forgotten)</i>

Table 7: Definitions, fine-grained sub-roles, and representative examples of the six Narrative Role categories used by ExTax.

Category	Definition	Example
Fear	Language that makes people feel scared or threatened, highlighting immediate serious risks to provoke a protective reaction.	<i>"ESSENTIAL OIL DIFFUSERS POISONOUS FOR CATS! ... I unknowingly have been poisoning my cat for years."</i>
Anger	Language provoking outrage or a sense of injustice, often targeting specific groups or institutions.	<i>"Corruption Exposed: Officials Withholding Cheap Life-Saving Medicine to Protect Corporate Profits!"</i>
Hope	Language eliciting positive emotion or promising miraculous breakthroughs, e.g., unverified cures.	<i>"Scientists Find Root That Kills 98% Of Cancer Cells In Only 48 Hours."</i>
Anxiety	Language creating uncertainty about the future or safety, e.g., disturbing medical cases or hidden daily dangers.	<i>"Five-foot long tapeworm came 'wiggling out' of man's body after he ate sushi."</i>
Sadness	Language focusing on irreversible loss or the suffering of helpless victims, triggering sympathy and moral guilt.	<i>"HE DIED WAITING: 82-year-old veteran found starving in a house full of medals."</i>

Table 8: Definitions and representative examples of the five Emotional Appeal categories used by ExTax.

A.1.6 Manipulative Wording

Manipulative Wording refers to language that is not necessarily an argument by itself, but uses non-neutral, confusing, exaggerated, or loaded expressions to influence the reader emotionally.

- **Loaded Language:** The use of words or phrases with strong emotional implications, either positive or negative, to influence the audience.
- **Obfuscation, Intentional Vagueness, or Confusion:** The use of deliberately unclear, vague, or ambiguous language that leaves room for audience-specific interpretations.
- **Exaggeration or Minimisation:** Representing something in an excessive manner or making

it appear less important or smaller than it actually is.

- **Repetition:** Repeated use of the same phrase or expression in order to increase persuasive impact.

A.2 Emotion

Emotion refers to language designed to bypass purely rational processing and influence readers by triggering affective responses. We distinguish five high-level emotional manipulation types: Fear, Anger, Hope, Anxiety, and Sadness.

- **Fear:** Language designed to make readers feel scared or threatened, often by emphasizing immediate risks, serious health concerns, or distressing outcomes in order to provoke a protective reaction.

- **Anger:** Language designed to provoke outrage, resentment, or a sense of injustice, often by targeting specific groups, institutions, or actors.
- **Hope:** Language designed to elicit positive expectations or promise breakthroughs, often by presenting unverified cures, revolutionary alternatives, or highly optimistic outcomes.
- **Anxiety:** Language designed to create uncertainty about future safety or well-being, often by describing disturbing cases or hidden dangers in everyday practices.
- **Sadness:** Language focusing on irreversible loss, heartbreaking misfortune, or the suffering of helpless victims, thereby triggering sympathy, sorrow, or moral guilt.

A.3 Narrative Roles

Narrative Roles describe how a person, group, or entity is functionally positioned within a narrative, independently of its actual moral status. We distinguish six high-level role categories: Ethical Stabilizers, Altruistic Catalysts, Overt Aggressors, Deceptive Subversives, Institutional Toxins, and Marginalized Sufferers.

A.3.1 Ethical Stabilizers

Ethical Stabilizers are actors portrayed as using established values, moral authority, or mediation to maintain social cohesion, safety, and justice.

- **Guardian:** Actors portrayed as protectors of values or communities, such as law enforcement officers, soldiers, or community leaders.
- **Peacemaker:** Actors portrayed as seeking harmony, conflict resolution, diplomacy, negotiation, or mediation.
- **Virtuous:** Actors portrayed as righteous, noble, fair, just, or morally exemplary.

A.3.2 Altruistic Catalysts

Altruistic Catalysts are actors portrayed as challenging the status quo or enduring hardship to drive systemic change, liberation, or collective benefit.

- **Rebel:** Actors portrayed as challenging the status quo and fighting for major change, freedom, or liberation from oppression.

- **Martyr:** Actors portrayed as sacrificing their well-being, or even their lives, for a greater good or cause.

- **Underdog:** Actors portrayed as disadvantaged or unlikely to succeed, but striving against stronger forces or obstacles.

A.3.3 Overt Aggressors

Overt Aggressors are visible antagonists portrayed as using direct force, systemic power, or violent ideology to dominate, divide, or harm others.

- **Instigator:** Actors portrayed as initiating conflict, tension, unrest, or violence.
- **Tyrant:** Actors portrayed as abusing power, ruling unjustly, or oppressing those under their control.
- **Terrorist:** Actors portrayed as engaging in violence or terror to advance ideological goals, often by targeting civilians.
- **Bigot:** Actors portrayed as hostile or discriminatory toward specific groups, including racism, sexism, homophobia, antisemitism, Islamophobia, or other forms of hate speech.

A.3.4 Deceptive Subversives

Deceptive Subversives are hidden or manipulative actors portrayed as undermining trust, truth, or safety through secrecy, betrayal, deception, or strategic manipulation.

- **Conspirator:** Actors portrayed as participating in secret plots or covert plans to deceive or undermine others.
- **Spy:** Actors portrayed as gathering or transmitting sensitive information to a rival or enemy through secrecy or deception.
- **Traitor:** Actors portrayed as betraying a cause, community, or country.
- **Deceiver:** Actors portrayed as manipulating truth, spreading misinformation, or shaping public perception for strategic or self-interested purposes.
- **Saboteur:** Actors portrayed as deliberately damaging, obstructing, or weakening systems, organizations, or processes.
- **Foreign Adversary:** External actors portrayed as creating geopolitical tension or acting against another country's interests.

A.3.5 Institutional Toxins

Institutional Toxins are actors or entities portrayed as internal sources of organizational or systemic failure, caused by vice, incapacity, corruption, or incompetence.

- **Corrupt:** Actors portrayed as engaging in unethical or illegal behavior for personal, political, or financial gain.
- **Incompetent:** Actors portrayed as causing harm through ignorance, poor judgment, lack of skill, or inability to perform their duties.

A.3.6 Marginalized Sufferers

Marginalized Sufferers are actors portrayed as enduring harm, exploitation, neglect, exclusion, or unjust blame, often functioning as evidence of systemic failure or as targets for rescue.

- **Forgotten:** Marginalized or overlooked groups portrayed as ignored by society and lacking needed attention or support.
- **Exploited:** Individuals or groups portrayed as being used for others' benefit without adequate consent and with harm to their well-being.
- **Victim:** Actors portrayed as suffering physical or economic harm due to circumstances beyond their control.
- **Scapegoat:** Actors portrayed as being blamed unjustly for problems or failures in order to divert attention from the real causes or responsible parties.

B Prompts

The following section delineates the construction of prompts designed to guide LLMs in extracting taxonomic features across Persuasion, Emotion, and Narrative Roles. Specifically, the prompt configuration for eliciting Persuasion attributes is adopted from [Modzelewski et al. \(2025\)](#). Uppercase placeholders wrapped in hash-tags, such as #PERSUASION_DEFINITIONS#, #EMOTION_DEFINITIONS#, and #NARRATIVE_ROLE_DEFINITIONS#, are used as designators to denote the comprehensive taxonomic definitions detailed in [Appendix A](#).

B.1 Prompt 1: Extract Persuasion Taxonomic Features

System Prompt

You are an assistant who detects persuasion in text. Persuasive text is characterized by a specific use of language in order to influence readers. We distinguish the following high-level persuasion approaches:
#PERSUASION_DEFINITIONS#

You are the expert who detects high-level persuasion approaches: Attack on Reputation, Justification, Simplification, Distraction, Call, Manipulative Wording.

Context Prompt

Analyze the text and decide if the text contains any high-level persuasion approaches from the following: Attack on Reputation, Justification, Simplification, Distraction, Call, Manipulative Wording. For each high-level persuasion approach, provide an explanation of how it appears in the analyzed text. Be conservative in your final decisions; when you are not fully sure, answer No. Give your answer in the form of a dictionary:

```
{{
  "Attack_on_reputation": {"is_used": "Your answer. Use only Yes or No", "explanation": "If high-level persuasion Attack on Reputation appears, provide an explanation here."}},
  "Justification": {"is_used": "Your answer. Use only Yes or No", "explanation": "If high-level persuasion Justification appears, provide an explanation here."}},
  "Simplification": {"is_used": "Your answer. Use only Yes or No", "explanation": "If high-level persuasion Simplification appears, provide an explanation here."}},
  "Distraction": {"is_used": "Your answer. Use only Yes or No", "explanation": "If high-level persuasion Distraction appears, provide an explanation here."}},
  "Call": {"is_used": "Your answer. Use only Yes or No", "explanation": "If high-level persuasion Call appears, provide an explanation here."}},
  "Manipulative_wording": {"is_used": "Your answer. Use only Yes or No", "explanation": "If high-level persuasion Manipulative Wording appears, provide an explanation here."}}
}}
```

Text: {text}

Answer:

B.2 Prompt 2: Extract Emotion Taxonomic Features

System Prompt

You are a specialized linguistic analyst and expert in detecting emotional manipulation in online text. Your goal is to identify techniques designed to bypass rational thinking and influence readers by triggering specific emotional responses.

We distinguish between the following five high-level emotional appeals:
#EMOTION_DEFINITIONS#

You are the expert who detects emotional manipulation types: Fear, Anger, Hope, Anxiety, Sadness.

Context Prompt

Analyze the text and decide if the text contains any emotional manipulation types from the following: Fear, Anger, Hope, Anxiety, Sadness. For each emotional manipulation type, provide an explanation of how it appears in the analyzed text. Be conservative in your final decisions; when you are not fully sure, answer No.

Give your answer in the form of a dictionary:

```
{{
  "Fear": {"is_used": "Your answer. Use only
    Yes or No", "explanation": "If the
    Fear type appears, provide a one-
    sentence explanation starting with '
    This post contains language that may be
    intended to make you feel fear by
    ...'"}},
  "Anger": {"is_used": "Your answer. Use
    only Yes or No", "explanation": "If the
    Anger type appears, provide a one-
    sentence explanation starting with '
    This post contains language that may be
    intended to make you feel anger by
    ...'"}},
  "Hope": {"is_used": "Your answer. Use only
    Yes or No", "explanation": "If the
    Hope type appears, provide a one-
    sentence explanation starting with '
    This post contains language that may be
    intended to make you feel hope by
    ...'"}},
  "Anxiety": {"is_used": "Your answer. Use
    only Yes or No", "explanation": "If the
    Anxiety type appears, provide a one-
    sentence explanation starting with '
    This post contains language that may be
```

```
    intended to make you feel anxiety by
    ...'"}},
  "Sadness": {"is_used": "Your answer. Use
    only Yes or No", "explanation": "If the
    Sadness type appears, provide a one-
    sentence explanation starting with '
    This post contains language that may be
    intended to make you feel sadness by
    ...'"}},
}}
Text: {text}
Answer:
```

B.3 Prompt 3: Extract Narrative Role Taxonomic Features

System Prompt

You are an expert at identifying narrative framing and role portrayal in text. Your goal is to detect whether a text utilizes specific high-level narrative categories -- Ethical Stabilizers, Altruistic Catalysts, Overt Aggressors, Deceptive Subversives, Institutional Toxins, or Marginalized Sufferers -- to frame the participants in a story.

This analysis focuses on functional roles -- how a person, group, or entity is positioned within the narrative -- independent of their actual moral standing.

You distinguish between six high-level main roles based on the following archetypal markers:
#NARRATIVE_ROLE_DEFINITIONS#

Your analysis must be based on the specific linguistic cues and narrative structure within the text.

Context Prompt

Analyze the text below to determine whether it contains any of the following high-level roles: Ethical Stabilizers, Altruistic Catalysts, Overt Aggressors, Deceptive Subversives, Institutional Toxins, or Marginalized Sufferers.

For each detected role, provide a concise explanation of how it is portrayed based on the narrative function and word choice. Be conservative in your final decisions; if you are not fully sure, answer No.

Return your answer as a JSON object in the following format:

```
{{
  "Ethical_Stabilizers": {"
    "is_used": "Yes or No",
```

```

    "explanation": "If Yes, provide
    evidence of how a participant is framed
    as a protector of values or a
    righteous figure (e.g., Guardian,
    Virtuous, Peacemaker)."
```

```

}},
"Altruistic_Catalysts": {{
  "is_used": "Yes or No",
  "explanation": "If Yes, provide
  evidence of how a participant is framed
  as driving positive change or
  sacrifice (e.g., Rebel, Martyr,
  Underdog)."
```

```

}},
"Overt_Aggressors": {{
  "is_used": "Yes or No",
  "explanation": "If Yes, provide
  evidence of how a participant is framed
  as an initiator of conflict or
  violence (e.g., Tyrant, Terrorist,
  Instigator, Bigot)."
```

```

}},
"Deceptive_Subversives": {{
  "is_used": "Yes or No",
  "explanation": "If Yes, provide
  evidence of how a participant is framed
  through secrecy, betrayal, or
  misinformation (e.g., Conspirator, Spy,
  Traitor, Deceiver, Saboteur, Foreign
  Adversary)."
```

```

}},
"Institutional_Toxins": {{
  "is_used": "Yes or No",
  "explanation": "If Yes, provide
  evidence of how a participant is framed
  as abusing power or as a source of
  systemic failure (e.g., Corrupt,
  Incompetent)."
```

```

}},
"Marginalized_Sufferers": {{
  "is_used": "Yes or No",
  "explanation": "If Yes, provide
  evidence of how a participant is framed
  as suffering harm, exploitation,
  neglect, exclusion, or unjust blame (e.
  g., Victim, Exploited, Forgotten,
  Scapegoat)."
```

```

}}
}}

Text: {text}
Answer:

```

C Method

Details of entropy-driven dynamic label smoothing. For each taxonomic facet, the binary LLM outputs are first aggregated by category. In the Emotion task, for example, $e_{i,k}^{(c)} \in \{0, 1\}$ indicates whether the k -th LLM assigns sample T_i to emotion category $c \in \mathcal{C}_E$. The positive vote proportion $p_{i,c}$ measures the fraction of LLMs predicting the presence of category c . We compute binary entropy over this Bernoulli distribution to quantify model

disagreement. The same computation is applied to the Persuasion and Narrative Role facets by replacing $e_{i,k}^{(c)}$ with the corresponding components of $\mathbf{p}_{i,k}$ and $\mathbf{n}_{i,k}$. The resulting smoothed labels across all samples and categories form:

$$\mathcal{Y}_{\text{smooth}} = \left\{ y_{\text{tax}}^{(i,c)} \mid i = 1, \dots, n, \right. \\ \left. c \in \mathcal{C}_P \cup \mathcal{C}_E \cup \mathcal{C}_N \right\}. \quad (15)$$

When all LLMs agree, $H_{i,c} = 0$ and the smoothed label equals the vote proportion $p_{i,c}$. When the models disagree maximally, $H_{i,c}$ is high and the target is shifted toward 0.5, reducing the effect of uncertain synthetic supervision.

Details of adaptive gated pooling. Given the LLM hidden states $\mathbf{S}_i \in \mathbb{R}^{L_i \times D}$, Adaptive Gated Pooling uses a learnable parameter vector $\mathbf{w} \in \mathbb{R}^{L_{\text{max}}}$, where L_{max} is the maximum truncation length. For an input with length L_i , we slice the first L_i elements of $\mathbf{w}$ as $\mathbf{w}_{1:L_i}$ and compute the normalized gating weights:

$$\mathbf{g}_i = \text{softmax}(\mathbf{w}_{1:L_i}) \in \mathbb{R}^{L_i}, \quad (16)$$

$$g_{i,h} = \frac{\exp(w_h)}{\sum_{j=1}^{L_i} \exp(w_j)}. \quad (17)$$

The pooled representation is then obtained by a weighted sum:

$$\mathbf{v}_i = \sum_{h=1}^{L_i} g_{i,h} \mathbf{S}_{i,h}. \quad (18)$$

All elements of $\mathbf{w}$ are initialized to 1, making the initial pooling behavior equivalent to average pooling. During training, the gating vector is optimized to assign higher weights to more informative token positions.

Details of taxonomy-aligned representation learning. Each taxonomic facet is modeled using an independent MLP head:

$$\mathbf{h}_i^{(c)} = \text{LN}(\text{GELU}(\text{Linear}_1^{(c)}(\mathbf{v}_i))), \quad (19) \\ f_c(\mathbf{v}_i) = \text{Linear}_2^{(c)}(\text{Dropout}(\mathbf{h}_i^{(c)})),$$

where $\text{Linear}_1^{(c)} : \mathbb{R}^D \rightarrow \mathbb{R}^{D_h}$ and $\text{Linear}_2^{(c)} : \mathbb{R}^{D_h} \rightarrow \mathbb{R}^{|\mathcal{C}_c|}$. The outputs are passed through sigmoid activations to obtain multi-label probabilities.

For each facet, we optimize Binary Cross-Entropy (BCE) loss against the smoothed taxonomic labels. Taking Persuasion as an example:

$$\mathcal{L}_P = -\frac{1}{n} \sum_{i=1}^n \sum_{c \in \mathcal{C}_P} \left[y_{\text{tax}}^{(i,c)} \log(\hat{p}_{i,c}) + (1 - y_{\text{tax}}^{(i,c)}) \log(1 - \hat{p}_{i,c}) \right]. \quad (20)$$

The Emotion and Narrative Role losses are defined analogously.

Details of text feature transformation. The transformation module maps each token vector $\mathbf{S}_{i,h} \in \mathbb{R}^D$ into the 17-dimensional taxonomic space using a two-layer MLP:

$$\begin{aligned} \mathbf{h}_{\text{mid}} &= \text{Linear}_1(\text{LN}(\mathbf{S}_{i,h})), \\ \tilde{\mathbf{S}}_{i,h} &= \text{Linear}_2(\text{GELU}(\mathbf{h}_{\text{mid}})), \end{aligned} \quad (21)$$

where $\text{Linear}_1 : \mathbb{R}^D \rightarrow \mathbb{R}^{d_{\text{ff}}}$ and $\text{Linear}_2 : \mathbb{R}^{d_{\text{ff}}} \rightarrow \mathbb{R}^{17}$.

D Experiments

D.1 Dataset

- **CoAID** (Cui and Lee, 2020): A COVID-19 health disinformation dataset containing both news articles and social media posts, including over 4,000 news articles and more than 1,000 posts mainly from Twitter. The content covers pandemic-related topics such as public health policies, vaccines, treatment measures, and other healthcare claims. In our setting, it contains 5,011 samples across train, validation, and test splits, with a relatively stable genre distribution in each split: posts account for approximately 44%, while articles account for approximately 56%.
- **ISOT Fake News** (Ahmed et al., 2018, 2017): A large-scale fake news article dataset containing over 44,000 news articles collected from both reliable mainstream media, such as BBC and CNN, and unreliable or fake news websites. The articles span multiple domains, including politics, entertainment, technology, and health. Following the setting in PCoT, we use its test split with 499 samples, including 274 fake news instances and 225 real news instances.

- **ECTF** (Paka et al., 2021; Bansal et al., 2021): An extended version of the CTF (COVID-19 Tweet Fake) dataset for COVID-19 fake tweet detection. It consists of tweets annotated as fake or real and covers pandemic-related topics such as virus transmission, protective measures, public responses, and vaccine controversies. This dataset provides train, validation, and test splits. In our setting, it contains 4,152 samples, with 3,001, 751, and 400 samples in the train, validation, and test sets, respectively.
- **MultiDis** (Modzelewski et al., 2025): A researcher-constructed multi-topic disinformation dataset containing nearly 2,000 English news articles. The articles cover a broad range of political and social topics, including Anti-Europeanism, Anti-Atlanticism, Anti-migration, Xenophobia, Climate Change, Energy Crisis, Health, Institutional and Media Distrust, Gender Issues, Ukraine War, Refugees, and LGBT+. In our setting, this dataset is used as a test-only dataset and contains 499 test samples.
- **EUDisinfo** (Modzelewski et al., 2025): A disinformation dataset built upon the EUvsDisinfo database, focusing on pro-Kremlin propaganda narratives and political disinformation. The articles mainly address issues such as EU policies, the situation in Ukraine, NATO expansion, and related geopolitical narratives. In our setting, this dataset is used as a test-only dataset and contains 359 articles, including 241 credible articles and 118 disinformation articles.

D.2 Baseline Introduction

- **RoBERTa** (Liu et al., 2019): A BERT variant optimized via dynamic masking, larger batches, and expanded corpora without Next Sentence Prediction, serving as a competitive contextualized baseline.
- **DRES** (Farhangian et al., 2025): A framework routing instances to optimal feature spaces based on hardness measures to dynamically ensemble classifiers.
- **MSynFD** (Xiao et al., 2024): A syntax-aware framework capturing long-range dependency graphs and debiasing topical priors to detect fine-grained contextual distortions.

- **RAEmoLLM** (Liu et al., 2025): An emotion-based retrieval-augmented framework using a specialized LLM to dynamically select affective exemplars for cross-domain detection.
- **FACTUAL** (Li et al., 2025): A counterfactual augmented calibration framework disentangling robust features from domain shortcuts to dynamically calibrate model confidence scores.
- **DYNAMO** (Jin et al., 2025): A framework utilizing Monte Carlo Tree Search to verify factual paths and dynamically update knowledge graphs for real-time news detection.
- **PCoT** (Modzelewska et al., 2025): A zero-shot prompting method that first analyzes persuasive strategies in the input text and then uses this persuasion-aware reasoning to detect disinformation.

D.3 Metrics

To quantitatively evaluate our framework, we employ three standard metrics: **Accuracy**, **Macro Recall** and **Macro F_1 -score**.

Accuracy measures the global proportion of correctly classified instances:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (22)$$

where TP, TN, FP, and FN represent True Positives, True Negatives, False Positives, and False Negatives, respectively.

To account for potential class imbalances, we also use the **Macro F_1 -score**. It calculates the F_1 -score for each class independently and takes their unweighted arithmetic mean. For a specific class k , its precision (P_k), recall (R_k), and F_1 -score ($F_{1,k}$) are defined as:

$$P_k = \frac{\text{TP}_k}{\text{TP}_k + \text{FP}_k} \quad (23)$$

$$R_k = \frac{\text{TP}_k}{\text{TP}_k + \text{FN}_k} \quad (24)$$

$$F_{1,k} = \frac{2 \cdot P_k \cdot R_k}{P_k + R_k} \quad (25)$$

Given the binary class set $\mathcal{C} = \{0, 1\}$ (denoting reliable news and disinformation), the final Macro F_1 -score is computed as:

$$\text{Macro } F_1 = \frac{1}{|\mathcal{C}|} \sum_{k \in \mathcal{C}} F_{1,k} \quad (26)$$

Similarly, we report the **Macro Recall** to evaluate the model’s ability to identify instances from each class regardless of class frequency. Macro Recall computes the recall for each class independently and then takes their unweighted arithmetic mean. Given the binary class set $\mathcal{C} = \{0, 1\}$, Macro Recall is defined as:

$$\text{Macro Recall} = \frac{1}{|\mathcal{C}|} \sum_{k \in \mathcal{C}} R_k \quad (27)$$

D.4 ExTax Implementation

During the training process, we perform an exhaustive hyperparameter search to ensure optimal convergence. To assess robustness to random initialization, we conduct multiple independent training and evaluation runs across three random seeds: 43, 434, and 445. We subsequently report the mean and standard deviation of the core performance metrics. For the divergent attribute synthesis operation, we leverage a diverse suite of frontier models, including GPT-5-mini, DeepSeek-V4-flash, Qwen3.6-flash, and Grok-4.1-fast-reasoning. The label smoothing hyperparameter α is strictly grounded at 0.1896 to mitigate potential decision noise. All experiments and model training phases are computationally lightweight and can be fully executed on a single NVIDIA RTX 4090 (24 GB VRAM) GPU. For further details regarding the comprehensive training configurations and architectural model profiles, please refer to Table 10 and Table 9, respectively.

D.5 Baseline Implementation

Below, we report the detailed implementation of the reproduced baseline models. All baselines are evaluated under the same binary label space and dataset split protocol as ExTax. Specifically, CoAID and ECTF are used as source-domain datasets with training and validation splits, while CoAID, ECTF, EUDisinfo, ISOTFakeNews, and MultiDis are used for test evaluation. All labels are normalized into real and fake. For trainable baselines, hyperparameters are selected using the validation split. For LLM-based baselines, the source-domain training split is used as the support pool when retrieval, counterfactual generation, calibration, or knowledge construction is required. For each reproduced method, we retain the original high-level algorithmic design whenever possible and explicitly document the implementation-aligned adaptations required for our unified evaluation setting, including dataset loading, input/output formatting, label

Component	Parameter	Value
Taxonomies-Aligned	Adopted model	RoBERTa-large (Liu et al., 2019)
	Maximum token length	512
	Pooling strategy	Adaptive gated pooling
	Adaptive initialization	initialized to 1
	Aligned MLP hidden dims	[786, 32, 32]
	Aligned MLP dropout rate	0.3290
Disinformation Detection	Aligned MLP output dimension	Persuasion&Narrative Role: $d = 6$, Emotion: $d = 5$
	Number of HeteroAttention Block	$N_{att} = 1$ (hidden dim = $6+6+5 = 17$)
	Classification Head	hidden dim = 128
Total parameters	Activation	GELU
		$\approx 357M(\approx 267K \text{ trainable})$

Table 9: Full architecture details of ExTax

Category	Parameter	Value
Taxonomies-Aligned	lr	0.00065
	Weight decay	0.00069
	Epochs	10
	Patience	7
	Seed	[43, 434, 445]
	Batch size	128
Disinformation Detection	lr	0.00096
	Weight decay	0.00018
	Epochs	50
	Patience	3
	Seed	[43, 434, 445]
	Batch size	128
Infra	Hardware	an NVIDIA RTX 4090 (24 GB VRAM) GPU (approximately 1–2 h for training and evaluation)

Table 10: Training hyper-parameters and configuration for ExTax

normalization, computational approximations, and unavailable-component substitutions.

- **RoBERTa** (Liu et al., 2019): For the baseline implementation, we utilize the roberta-large variant as the backbone encoder network. Given an input sequence, the token-level contextualized embeddings are extracted from the encoder, and a mean-pooling operation is subsequently applied across the sequence length to derive a unified sequence representation. This pooled representation is then fed into a sequence-level classification head structured as a multi-layer perceptron (MLP). Specifically, the classification head comprises a linear projection layer mapping the hidden size into a 256-dimensional intermediate space, followed by a Gaussian Error Linear Unit (GELU) activation function and Layer Normalization. To prevent overfitting, a Dropout layer with a regularization rate of 0.25 is applied before the final linear layer, which maps the features directly into the 2-

dimensional classification target space.

- **DRES** (Farhangian et al., 2025): We reproduce DRES using the provided Dynamic Representation and Ensemble Selection implementation and adapt its input/output interface to our unified data splits. The training set is constructed by merging the CoAID and ECTF training splits, and the validation set is constructed by merging the corresponding validation splits. The test set contains CoAID, ECTF, EUDisinfo, ISOTFakeNews, and MultiDis. Following the dynamic representation selection principle of DRES, we construct multiple TF-IDF feature spaces with maximum dimensions of 100, 200, 300, 500, 700, and 1000, respectively, using English stop-word removal. For each representation, we train a pool of five classifiers, including Logistic Regression, Bernoulli Naive Bayes, k -Nearest Neighbors, Random Forest, and Support Vector Machine. Hyperparameters are selected according to validation-set

performance: Logistic Regression searches $C \in \{0.1, 1, 10\}$ with maximum iteration 2000; Bernoulli Naive Bayes searches $\alpha \in \{0.5, 1.0\}$; k -Nearest Neighbors searches $k \in \{3, 5, 7\}$; Random Forest uses 100 estimators and searches maximum depth in $\{5, 10, 20\}$; and SVM searches $C \in \{0.1, 1, 10\}$ with an RBF kernel. Instance hardness is estimated by the k -Disagreeing Neighbors criterion with $k = 5$. For each test instance, the representation with the lowest estimated hardness is selected dynamically. Dynamic ensemble selection is then performed using KNORA-Eliminate (KNORA-E; implemented as KNORAE in deslib2), with the region of competence size set to 5. Following the local-oracle criterion of KNORA-E, only base classifiers that correctly classify all samples in the local region of competence are selected; if no classifier satisfies this criterion, the region is iteratively reduced and re-evaluated. The final prediction is obtained by majority voting among the classifiers selected by KNORA-E. The experiment is repeated with three random seeds, $\{43, 434, 445\}$.

- **MSynFD** (Xiao et al., 2024): We reproduce MSynFD as a syntax-aware neural fake news detection baseline based on the released implementation structure. The model uses bert-base-uncased as the textual encoder, followed by a Semantic Transformer module, a Graphformer-based multi-hop syntactic aggregation module, and a keyword debiasing branch. The hidden dimension is set to 768, the number of attention heads is 12, the dropout rate is 0.1, and the hop number is set to 3. The maximum input sequence length is 200, and the maximum keyword sequence length is 50. In our reproduced implementation, the syntactic adjacency matrix is approximated by a local dependency-distance window with window size 3, rather than by running an external dependency parser. The final logits are computed by combining the main prediction branch and the keyword debiasing branch with weights 0.9 and 0.1, respectively. The training objective is the sum of the cross-entropy loss on the final prediction and 0.1 times the cross-entropy loss on the keyword debiasing output. The model is trained for 3 epochs with batch size 8 and learning

rate 1×10^{-5} using AdamW. Following the actual optimizer parameter grouping in the implementation, non-bias and non-LayerNorm parameters use weight decay 0.01, while bias and LayerNorm parameters use no weight decay. For each run, up to 2000 training samples and 500 validation samples are sampled using the corresponding random seed. Evaluation is performed with batch size 16. We repeat the experiment with seeds $\{43, 434, 445\}$.

- **RAEmoLLM** (Liu et al., 2025): We implement RAEmoLLM as an emotion-aware retrieval-augmented LLM baseline. The support pool is built from the CoAID and ECTF training splits. For each test instance, the model extracts affective information and retrieves source-domain examples as in-context demonstrations. The reproduced implementation uses keyword-based affective features and, when available, transformer-based emotion or sentiment models. For retrieval, semantic similarity is computed with all-MiniLM-L6-v2 when the encoder is available, and affective similarity is computed from the extracted emotion-oriented representation. The retrieval score combines semantic and affective similarity with weights 0.45 and 0.55, respectively. If neural sentence encoders are unavailable, the implementation falls back to a TF-IDF retriever with unigram and bigram features. The retrieval size is set to $\text{top-}K = 1$. The retrieved example, affective analysis, and target news text are inserted into a JSON-constrained prompt, and the LLM is required to output the predicted label, confidence score, and reasoning. The generated label is then normalized into the unified binary label space before evaluation.
- **FACTUAL** (Li et al., 2025): We reproduce FACTUAL as a counterfactual augmented calibration baseline and adapt it to binary fake news detection. The implementation first obtains an original LLM prediction and two debiased predictions, where the debiased prompts explicitly reduce sentiment-driven and target-prior-driven shortcuts. It then generates two counterfactual variants of the input: a causal counterfactual that minimally changes the text to flip the authenticity judgment, and a non-causal counterfactual that paraphrases

the input while preserving the original stance. These counterfactual texts are re-evaluated by the LLM to estimate counterfactual fake probabilities. The final prediction probability is calibrated using the original prediction, the two debiased predictions, the counterfactual predictions, and an explicit bias signal. The calibration strength is set to $\alpha = 0.35$. To avoid using target-test labels for calibration, the lightweight Logistic Regression calibration module is treated as an optional component and is used only when a calibrator trained on source-domain or validation data is explicitly available. Otherwise, we use the heuristic calibration rule implemented in the reproduced code. The final calibrated probability is thresholded at 0.5 to obtain the binary prediction.

- **DYNAMO** (Jin et al., 2025): We implement DYNAMO as a dynamic knowledge-update-driven LLM baseline with knowledge graph verification and Monte Carlo Tree Search. The source-domain support pool is built from the CoAID and ECTF training splits. The initial knowledge graph is constructed from reliable source-domain records labeled as real; if the number of reliable records is insufficient, the implementation falls back to the full support pool. In the reproduced setting, up to 10 support records are used for initial knowledge graph construction, and at most 2 entity-centered triples and 2 event-centered triples are extracted from each selected text. For each test instance, DYNAMO decomposes the input into verifiable subclaims through Monte Carlo Tree Search. The reproduced MCTS configuration uses maximum depth 2, maximum branching factor 2, one simulation, and UCB1 exploration coefficient 1.414. Each subclaim is verified against the constructed knowledge graph and assigned one of three verdicts: supported, contradicted, or unverifiable. The verification statistics and reasoning traces are then inserted into a JSON-constrained prompt, from which the LLM outputs the final binary label, confidence score, and reasoning. For efficiency, online knowledge updates are deferred during batch evaluation.

- **PCoT** (Modzelewski et al., 2025): We imple-

mented the baseline using the official code released by PCoT at <https://github.com/ArkadiusDS/PCoT>. During evaluation, we ran all three prompting strategies, namely VaN, Z-CoT, and DeF-SpeC, and selected the best-performing strategy, DeF-SpeC, as the final reported result.

For the three LLM-based baselines, RAE-moLLM, FACTUAL, and DYNAMO, we use the same OpenAI-compatible chat-completion backend with the configured model name gpt-5-mini. The decoding temperature is set to 0.1, the maximum generation length is 256 tokens, and the request timeout is set to 30 seconds. Each LLM-based baseline is evaluated over 5 repeated inference runs with fixed base seed 42 for local sampling and preprocessing when applicable; variation across runs mainly comes from non-zero-temperature LLM generation. All LLM outputs are parsed through JSON-constrained generation and then normalized into the same binary label space as ExTax.

D.6 External Validation on Human-Annotated Taxonomies

Evaluation setting. We further evaluate whether the taxonomy-aligned representations learned in Stage 1 transfer to independently human-annotated taxonomic signals. Importantly, the taxonomy-aligned heads are trained only using the LLM-generated weak supervision constructed from our original disinformation datasets and are not re-trained or fine-tuned on either external benchmark. The evaluation therefore constitutes an out-of-distribution external validity test.

For Persuasion, we use the English development split of SemEval-2023 Task 3, Subtask 3, containing 90 articles. The original annotations are paragraph-level multi-label annotations over 23 fine-grained persuasion techniques. We map these fine-grained labels into the same six high-level persuasion categories used by ExTax: Attack on Reputation, Justification, Simplification, Distraction, Call, and Manipulative Wording. Paragraph-level labels are then aggregated to the article level: a category is considered positive for an article if it occurs in at least one paragraph.

For Narrative Roles, we use the English development split of SemEval-2025 Task 10 (Piskorski et al., 2025b), Subtask 1, containing 27 documents. The original annotations are defined at the entity-

Text: #Hope #NYC #NY #USA #World #Health vigilance a #vaccine is a year away #Pandemic harmful is misinformation side effects discrimination before #Ebola #SARS #MERS there was #AIDS a community marginalize...

Label: Disinformation **Genre:** Post

Detected Taxonomies Type:

- Persuasion: Justification, Call, Manipulative wording.
- Emotion: Hope, Sadness.
- Narrative Roles: Deceptive Subversives, Marginalized Sufferers.

Table 11: Case study 1.

Text: Title: Trump Attempts New Year’s Eve Do-Over Tweet And Gets CRUSHED By Americans He Called ‘Enemies’ Content: If Donald Trump had sent proper New Year’s greetings to the American people to begin with, he wouldn’t be getting humiliated. But he royally f*cked up on New Year’s Eve morning by referring to millions...

Label: Disinformation **Genre:** News Article

Detected Taxonomies Type:

- Persuasion: Attack on reputation, Justification, Simplification, Manipulative wording.
- Emotion: Anger.
- Narrative Roles: Overt Aggressors, Deceptive Subversives, Institutional Toxins.

Table 12: Case study 2.

Benchmark	Category	AUC
SemEva’23 Task 3 Persuasion	Attack on Reputation	0.831
	Justification	0.635
	Simplification	0.571
	Distraction	0.735
	Call	0.666
	Manipulative Wording	0.886
	Macro	0.720
SemEval’25 Task 10 Narrative Role	Ethical Stabilizers	0.533
	Altruistic Catalysts	0.600
	Overt Aggressors	0.856
	Deceptive Subversives	0.606
	Institutional Toxins	0.602
	Marginalized Sufferers	0.482
	Macro	0.613

Table 13: Category-level external validation of the taxonomy-aligned representation heads on independently human-annotated datasets. AUC is computed for each of the six mapped categories, and Macro denotes the unweighted average across categories.

mention level using three coarse roles and 22 fine-grained roles. We map the fine-grained roles into the six Narrative Role groups adopted by ExTax: Ethical Stabilizers, Altruistic Catalysts, Overt Aggressors, Deceptive Subversives, Institutional Toxins, and Marginalized Sufferers. Entity-level labels are aggregated to the document level, where a category is marked positive if any entity mention within the document belongs to a corresponding fine-grained role.

After aggregation, each document is associated with a six-dimensional binary human-annotated target vector. We directly apply the corresponding

Stage 1 taxonomy-aligned head to each document and use its six predicted probabilities as category-level scores. Since label frequencies are highly imbalanced and no decision threshold is calibrated on the external datasets, we report the area under the ROC curve (AUC) for each category and their macro average.

Results. Table 13 shows that the Persuasion head exhibits clear transferability to independent human annotations, achieving a Macro AUC of 0.720. The strongest alignment is observed for Manipulative Wording (AUC = 0.886) and Attack on Reputation (AUC = 0.831), while Distraction also shows substantial discriminative ability (AUC = 0.735). These results suggest that the representations learned from LLM-generated supervision capture several persuasion dimensions that remain discriminative under a different annotation source and data distribution.

The Narrative Role head achieves a Macro AUC of 0.613. In particular, Overt Aggressors transfers strongly (AUC = 0.856), while Altruistic Catalysts, Deceptive Subversives, and Institutional Toxins remain above random-level discrimination. Marginalized Sufferers is more challenging (AUC = 0.482). We attribute the weaker overall transfer partly to the substantially smaller evaluation set ($n = 27$), as well as differences between the original entity-level framing annotations and ExTax’s document-level representation.

Overall, these results do not establish per-

fect correspondence between LLM-generated supervision and human judgments. Rather, they provide independent evidence that the learned taxonomy-aligned representations retain meaningful alignment with human-annotated persuasion and narrative-role signals across datasets and annotation schemes.

Label-space alignment. Because the external benchmarks provide annotations at a finer granularity than the six-dimensional heads used by ExTax, we map their official fine-grained labels to the corresponding high-level categories before evaluation. For example, fine-grained persuasion techniques such as Name Calling/Labeling and Casting Doubt are mapped to Attack on Reputation, while narrative roles such as Conspirator, Spy, and Deceiver are mapped to Deceptive Subversives. The complete mapping is provided with our released evaluation code.

D.7 Interpretability

D.7.1 Global interpretability analysis

This appendix provides additional analysis of the interpretability results reported in the main text. The global distribution of the three taxonomic dimensions—persuasion strategies, emotional manipulation, and narrative roles—is shown in Figure 2. As illustrated, disinformation and real news exhibit markedly distinct structures within the constructed 17-dimensional taxonomic attribute space, providing intuitive evidence of latent manipulation patterns in deceptive texts.

Specifically, concerning persuasion strategies, fake news relies heavily on “Attack on Reputation”, “Simplification”, and “Manipulative Wording”, reflecting a typical discourse pattern that incites audiences by denigrating opposing targets through extreme language or misleads readers by omitting crucial information. This directly corroborates the findings in PCoT (Modzelewski et al., 2025). Regarding the emotional dimension, disinformation displays skewed distributions toward “Anger” and “Fear”—two high-arousal negative emotions—which closely aligns with psychological theories (Jamieson et al., 2025; Zhao et al., 2024) suggesting that intense emotional stimuli accelerate the dissemination of misleading content. Furthermore, in shaping narrative roles, disinformation strongly tends to frame narrative subjects as negative archetypes, such as “Deceptive Subversives” and “Institutional Toxins”, whereas real news leans

more toward positive or constructive roles such as “Ethical Stabilizers”.

To further elucidate how taxonomic attributes across different dimensions jointly manifest within texts, the Sankey diagram in Figure 3 reveals their strong co-occurrence relationships. By tracing the convergence and divergence of the data flows, we observe several typical attribute binding patterns:

- *High-frequency co-occurrence of aggressive persuasion and negative emotions.* From the left to the middle of the diagram, the persuasion strategy of “Attack on Reputation” exhibits a strong binding relationship with “Anger” and “Fear”. This indicates that when disinformation employs aggressive discourse, it is frequently accompanied by intense, high-arousal negative emotional stimuli; the two act as closely coupled co-occurrence components within inciting texts.
- *Mutual corroboration of rational defense and constructive roles.* In contrast to the extreme co-occurrence patterns of disinformation, “Justification”—a more neutral and objective persuasion strategy—has a high-probability co-occurrence with the “None” category in the emotional dimension. Texts without extreme emotional manipulation are also consistently associated with the positive narrative role of “Ethical Stabilizers”, forming a core feature cluster of authentic news.
- *Chain transmission of non-manipulative characteristics.* When explicit manipulative signals are absent at one level, i.e., labeled as “None”, this absence is likely to co-occur in subsequent levels. For instance, a large flow originating from “None” in persuasion strategies transitions into “None” in emotions and subsequently into “None” for narrative roles. This suggests that authentic news maintains a consistent non-manipulative pattern across multiple cognitive dimensions.

Overall, these distributional disparities show that ExTax not only improves the transparency of neural disinformation detection, but also captures taxonomy-aligned manipulation evidence, providing structured views of the manipulation patterns associated with model predictions.

Text: We cannot prevent infectious diseases from emerging. But we can and must double down on developing #vaccines. Vaccines are our most powerful tool in fighting epidemic diseases. #coronavirus #publichealth ...
Label: Real **Prediction:** Disinformation ($P(\text{real}) = 0.007$) **Genre:** Post

Detected Taxonomies Type:

- Persuasion: Justification, Call.
- Emotion: None.
- Narrative Roles: Ethical Stabilizers.

Table 14: Case study 3.

Text: Title: UN Celebrates Its 70th Anniversary With Communist Statue In NYC Park. Content: In case you hadn't heard, Obama signed a new UN development plan to make America a socialist nation by 2030. Yes, it's true and the UN is celebrating by unveiling a statue celebrating our "global citizenship"... For example, on the subject of universal healthcare, on page 6, section 26 of "Transforming our World: The 2030 Agenda for Sustainable Development"... it states: 26. To promote physical and mental health...

Label: Disinformation **Prediction:** Real ($P(\text{real}) = 0.9998$) **Genre:** News Article

Detected Taxonomies Type:

- Persuasion: None.
- Emotion: None.
- Narrative Roles: None.

Table 15: Case study 4.

D.7.2 Case Study Analysis

To further examine the instance-level interpretability and failure modes of ExTax, we present four representative case studies. Cases 1 and 2 illustrate taxonomy-aligned profiles for correctly identified disinformation, while Cases 3 and 4 examine representative false-positive and false-negative errors.

Case Study 1: Health-related social media post.

Table 11 presents a short pandemic- and vaccine-related social media post. Although the text is fragmented and hashtag-driven, ExTax identifies a coherent manipulation profile. In the persuasion dimension, the detected categories include *Justification*, *Call*, and *Manipulative Wording*, indicating that the post attempts to legitimize its stance through emotionally charged public-health references while encouraging readers to adopt a particular viewpoint. Emotionally, ExTax detects both *Hope* and *Sadness*: vaccine-related references evoke hopeful expectations, whereas mentions of marginalized communities and severe disease histories trigger vulnerability and sympathy. In terms of narrative roles, *Deceptive Subversives* and *Marginalized Sufferers* suggest that the post frames some actors as potentially misleading while positioning affected communities as vulnerable victims. This case shows that ExTax can extract interpretable manipulation signals even from short and noisy social media posts.

Case Study 2: Political news article.

Table 12 shows a long-form political article about Donald Trump's New Year's Eve tweet. Compared with the first case, this example contains more explicit evaluative and adversarial language. ExTax detects *Attack on Reputation*, *Justification*, *Simplification*, and *Manipulative Wording*, suggesting that the article relies on reputational attacks and reduces a complex political communication event into a morally one-sided conflict. The detected emotion, *Anger*, is consistent with the article's confrontational and mocking tone. The narrative-role predictions, including *Overt Aggressors*, *Deceptive Subversives*, and *Institutional Toxins*, further indicate that political actors are framed as visible antagonists and sources of institutional or social harm. Thus, the disinformation signal arises not from a single lexical cue, but from the alignment of aggressive persuasion, anger-based emotional manipulation, and negative role assignment.

Comparison between the two cases.

Together, the two cases highlight the genre-adaptive interpretability of ExTax. The first case represents an implicit, compressed, and hashtag-driven manipulation pattern in short-form social media, while the second exhibits a more explicit pattern based on reputational attacks, anger, and antagonistic role construction in long-form political news. Despite these genre differences, both cases are mapped into the same 17-dimensional taxonomic space, enabling comparable explanations across posts and

articles. These examples show that ExTax goes beyond binary classification by exposing how deceptive texts manipulate readers through rhetoric, emotion, and narrative framing.

Case Study 3: False positive on public-health advocacy. Table 14 presents a real public-health advocacy post that is incorrectly classified as disinformation with $P(\text{real}) = 0.007$. The post employs rhetorical patterns that overlap with those frequently observed in deceptive content, including an explicit call to action (“we can and must double down”) and a strongly positive framing of vaccines as “our most powerful tool.” These cues activate the Justification and Call dimensions, while the public-health framing is associated with Ethical Stabilizers. This case suggests that ExTax may over-rely on rhetorical and framing signals when they occur in legitimate advocacy. More generally, the taxonomy captures how a message is communicated rather than whether its claims are factually correct, and persuasive but truthful content may therefore trigger false positives.

Case Study 4: False negative on stylistically neutral disinformation. Table 15 shows a disinformation article that is incorrectly classified as real with $P(\text{real}) = 0.9998$. The article advances a conspiratorial interpretation of the UN 2030 Agenda, but adopts a superficially factual news style. In particular, it cites concrete document locations such as “page 6, section 26” and quotes genuine policy language, giving the text an appearance of documentary support while avoiding overtly emotional or aggressively manipulative wording. None of the 17 taxonomic categories is activated in this example. This exposes a complementary failure mode of ExTax: when disinformation is rhetorically restrained and mimics conventional reporting, taxonomy-based manipulation signals may provide little evidence for detection. The model therefore struggles with content whose misleading effect depends primarily on selective quotation, contextual distortion, or false inference rather than explicit persuasive, emotional, or narrative-role cues.

Together, Cases 3 and 4 reveal two complementary failure modes of ExTax: legitimate but persuasive content may trigger false positives, while stylistically neutral disinformation may evade the taxonomy when its misleading mechanism lies primarily in factual distortion or contextual misrepresentation rather than overt manipulation cues.