Modeling Discourse Profiles in Climate Misinformation Discourse

Anonymous ACL submission

Abstract

Prior work on climate misinformation detection typically classifies individual comments or claims using text alone. This framing overlooks a key property of online climate discourse: misinformation is often produced, reinforced, and sustained through recurring rhetorical patterns that reappear across comments and threads. Using a large-scale TikTok climate discourse dataset, we construct discourse profiles by clustering climate-related comments into recurring rhetorical patterns. We then incorporate these profiles into fine-grained climate misinformation classifiers. Across sparse lexical features, sentence embeddings, and large contextual encoders, our discourse-profile-aware models consistently outperform text-only baselines. We further compare multiple discourse-profile conditioning strategies, including prefix-based, prompt-based, and adapter-based formulations, and find that late-fusion modeling yields the strongest and most consistent improvements. Beyond classification, our results show discourse profiles shape misinformation reception and propagation. Discourse-profile-conditioned language model agents exhibit misinformation acceptance patterns aligned with human susceptibility, and the relationship between misinformation prevalence and engagement persistence differs across dominant discourse profiles.

1 Introduction

Short-video social media platforms such as TikTok have become major venues for public discussion of climate change, where accurate information, skepticism, and misinformation coexist in the same high-volume comment threads (Pearce et al., 2018; Sultana et al., 2024; Basch et al., 2021; Pera and Aiello, 2024; Baltasar et al., 2024). This setting is particularly challenging for automated analysis: short, algorithmically amplified videos drive rapid diffusion, while discussions are fast-moving, emotionally charged, allowing misinformation to persist and resist correction.

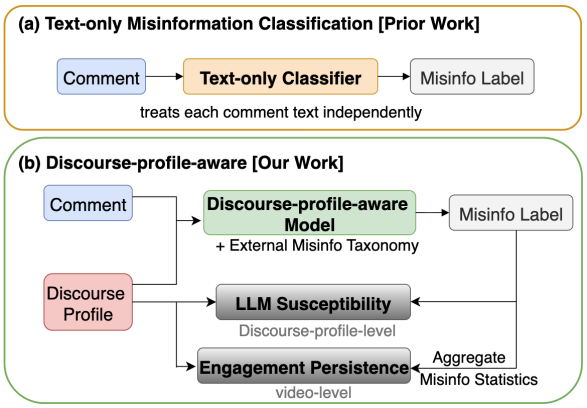

Figure 1: Comparison of text-only vs. discourse-profile-aware misinformation modeling. (a) Prior work (Islam et al., 2020; Coan et al., 2021; Upadhyaya et al., 2023; Choi et al., 2023; Rojas et al., 2024) typically predicts a misinformation category from comment text alone, treating each comment independently. (b) We induce discourse profiles by clustering comments and use discourse profile as additional context for misinformation-type prediction; the same discourse profile abstraction also supports discourse-profile-conditioned analyses of LLM susceptibility and engagement persistence.

As shown in Figure 1a, most prior work uses text-only classifiers that label each comment or claim in isolation (Islam et al., 2020; Coan et al., 2021; Upadhyaya et al., 2023; Choi et al., 2023; Rojas et al., 2024). However, climate discourse also exhibits recurring narrative frames and rhetorical styles that reappear across threads. This suggests that *how people consistently argue* can provide complementary signal beyond the words in a single comment.

We introduce *discourse profiles* to address this challenge: a compact, data-driven abstraction of recurring narrative frames and rhetorical styles in climate discourse. This perspective complements prior work showing that incorporating speaker and social context can improve misinformation detec-

tion beyond text alone (Nguyen et al., 2020; Mehta et al., 2022), and that beliefs and framing are not fully captured by lexical features (Allaway and McKeown, 2021; Zhang et al., 2024). Rather than relying on predefined demographics or static author embeddings, we induce discourse profiles directly from large-scale, in-the-wild TikTok comments. Specifically, we cluster semantic comment embeddings and consolidate clusters into a small set of interpretable discourse profiles. Importantly, profiles are assigned at the *comment level*, not the user level, allowing the same user to exhibit different discourse patterns across contexts.

We then use discourse profiles in three ways (Figure 1b): (i) to augment fine-grained misinformation-type prediction, (ii) to construct a controlled susceptibility probe using discourse-profile-conditioned LLM agents, and (iii) to analyze how misinformation relates to engagement persistence across different discourse contexts. Our main contributions are:

- We induce discourse profiles from TikTok climate comments and consolidate large-scale clusters into interpretable profile types.
- We show that discourse profile provides complementary signal for fine-grained misinformation classification across sparse lexical features, sentence embeddings, and large contextual encoders. We further compare multiple discourse-profile conditioning strategies, including prefix-based, prompt-based, and adapter-based formulations, and find that late-fusion conditioning yields the strongest and most consistent improvements.
- We evaluate discourse-profile-conditioned LLM agents in a controlled susceptibility setting and show strong correlations with human discourse-profile-level misinformation susceptibility.
- We show that the relationship between misinformation and engagement persistence is discourse-profile-conditional: misinformation co-varies with persistence depending on the dominant conversational discourse profile.

2 Related Work

2.1 Climate Misinformation Narratives

Climate misinformation includes not only outright denial, but also broader obstruction narratives such as questioning climate solutions, emphasizing economic trade-offs, or attacking scientific credibility (Lamb et al., 2020; Coan et al., 2021; Holder et al., 2023). We adopt the fine-grained QUOTACLIMAT taxonomy (QuotaClimat, 2025), which captures diverse climate obstruction narratives (Appendix B.1).

Most prior work treats misinformation as isolated textual instances. In contrast, we model recurring rhetorical patterns through discourse profiles and analyze how different profiles preferentially express specific misinformation types.

2.2 Author Context and Persona in NLP

Prior work shows that author context can improve stance and misinformation detection by capturing latent beliefs, interaction patterns, and communicative styles beyond text alone (Benton and Dredze, 2018a,b; Qiu et al.; Lynn et al., 2019; Barel et al., 2025). Related work has also explored dynamic and multi-persona user representations (Del Tredici et al., 2019; Marsden and Pröbster, 2019), while studies of climate denial rhetoric show that misinformation discourse often relies on recurring rhetorical and networking strategies rather than isolated claims (Bloomfield and Tillery, 2018). Existing approaches typically rely on predefined attributes, interaction structures, or static user embeddings. We instead induce discourse profiles directly from large-scale climate discourse and use them both for prediction and downstream behavioral analysis.

2.3 LLM Profiles and Engagement Dynamics

Recent work conditions LLMs on personas or discourse profiles to simulate subgroup behavior (Park et al., 2023; Krishna et al., 2022; Argyle et al., 2023; Li et al., 2024). We instead use discourse-profile-conditioned LLM agents as a controlled probe for misinformation susceptibility.

Prior work has also examined misinformation diffusion and persistence through content, emotion, and network structure (Vosoughi et al., 2018; Liu et al., 2023; Grinberg et al., 2019; Sun and Xie, 2024). We extend this line of work by showing that engagement persistence varies systematically across discourse profiles.

3 Method

Given short-form social media comments discussing climate change, our goal is to characterize and predict fine-grained climate misinformation by explicitly modeling recurring discourse-level regularities. For each comment c_i, we (i) assign a discourse profile d_i capturing recurring rhetor-

Figure 2: Offline discourse profile induction: encode comments, cluster embeddings, and manually consolidate clusters to assign each comment a discourse profile label p_i (Section 3.1).

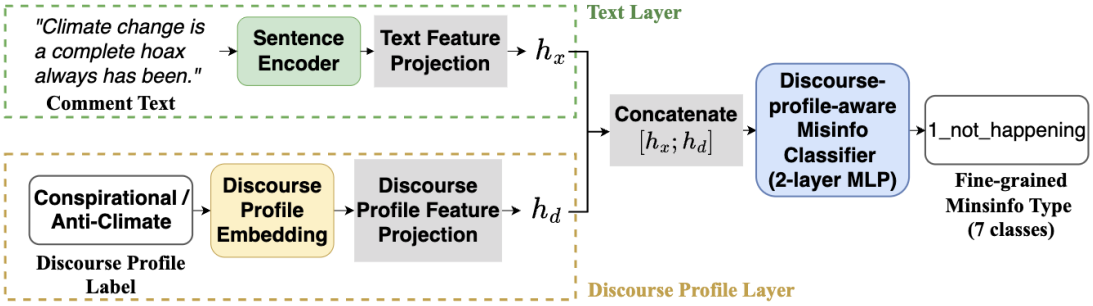

Figure 3: Late-fusion discourse-profile-aware misinformation modeling: obtain misinformation-type labels and evaluate prediction models that concatenate contextual text representations with discourse-profile embeddings before classification (Section 3.2). Details of how we obtain comment-level misinformation labels are in Appendix A.

ical patterns, (ii) obtain a fine-grained misinformation type y_i when applicable, and (iii) analyze how discourse profile conditions misinformation susceptibility and engagement persistence at both the comment and video levels. Dataset details are provided in Section 4.

Our pipeline proceeds in two stages. First, we induce discourse profiles from large-scale climate comment discourse (Figure 2; Section 3.1). Second, we incorporate discourse profile as an explicit signal for misinformation-type prediction (Figure 3) and for downstream analyses of susceptibility and engagement. Comment-level misinformation labels are obtained using a classifier trained on a curated climate misinformation taxonomy; since the transfer procedure is not central to our main contributions, we defer its details to Appendix A.

3.1 Discourse Profile Creation

To capture recurring rhetorical and discourse patterns used in Sections 3.2-3.4, we induce discourse profiles from large-scale climate change comments in CLIMATEDISC (Gao et al., 2025), a large-scale corpus of TikTok climate change videos and associated comments (section 4), which serve as a discourse-level abstraction across downstream tasks.

We represent comments as dense semantic embeddings using a pre-trained MPNet-based sentence transformer (Reimers and Gurevych, 2020). We then cluster the embeddings using k-means (Lloyd, 1982), selecting the number of clusters by jointly examining Within-Cluster Sum of Squares (WCSS) and Silhouette scores over a candidate range of $k \in [3, 30]$. This analysis identifies $k = 21$ as a practical trade-off between cluster compactness and separation (Section 4.1).

To obtain a stable discourse-level abstraction for downstream modeling, we manually validate and consolidate the 21 fine-grained clusters into six higher-level *discourse profiles* reflecting shared communicative stances toward climate change, along with a residual "Noise / Low-Signal" group excluded from discourse-profile-sensitive analyses. The resulting discourse profile definitions are reported in Table 1.

1. *"Pennsylvania, same dread."*
2. *"We're having an unseasonably warm winter here in Minnesota."*
3. *"This winter in Michigan there was only snow for a few weeks... now it's 50 degrees and sunny."*

While these comments differ in surface form, they all ground climate discussion in place-based lived experience. We therefore treat them as reflecting a shared discourse-level profile.

With discourse profiles defined, we next examine whether this abstraction provides predictive signal beyond comment text alone. Specifically, we evaluate whether incorporating discourse profile information improves fine-grained climate misinformation classification.

Local / Experiential Observers 1. Local Weather Anecdotes 12. Storms & Flooding 13. Geographical Self-Mentions	**Ideological / Political Combatants** 0. Anti-Trump/GOP/Biden 11. Government Blame / Anarchy 18. Electoral Mobilization
Systemic / Structural Critics 4. Consumption / Plastics Skepticism 17. Energy / Policy Debate 19. Agriculture / Diet 20. Anti-Corporation / Capitalism	**Pro-Science / Rationalist** 14. Pro-Science / Anti-Misinformation
Doomism / Nihilism 8. Doomism	**Conspiratorial / Anti-Climate** 7. Natural causes – not man-made 16. Denial / Hoax / Manipulation
Noise / Low-Signal (Excluded): 2. Emoji / Applause Spam; 3. Minimal Agreement Tokens; 5. Polite Thanks; 6. Short Emotional Bursts; 9. Excited Agreement; 10. Shared Anxiety; 15. @ Mentions.	

Table 1: Discourse profiles obtained by grouping the 21 k-means clusters of comment embeddings into semantically coherent discourse-profiles. The Noise / Low-Signal group is excluded from downstream discourse-profile-sensitive analyses.

3.2 Profile-Aware Misinfo Classification

We compare text-only and discourse-profile-aware classifiers to test whether discourse profile provides predictive signal beyond comment text. Our goal is not to maximize benchmark performance, but to isolate when discourse profile contributes complementary information under different text representations and model capacities.

Text representations. Given a comment t in CLIMATEDISC, we encode it as a feature vector $\mathbf{x} = \phi(t)$, where $\phi(\cdot)$ may denote either a pre-trained dense sentence encoder or a lexical featurizer. In particular, we consider three families of text representations: (i) MPNet-based sentence embeddings (Reimers and Gurevych, 2020), (ii) non-contextual bag-of-words features, and (iii) TF-IDF (Jones, 2004) features. This allows us to examine the effect of discourse profile across both dense semantic and sparse lexical input spaces.

Text-only baseline. Model A is a text-only classifier that predicts fine-grained misinformation type from $\mathbf{x}$ alone. The text representation is passed through a nonlinear projection $\mathbf{h}_x = f_x(\mathbf{x})$, where $f_x(\cdot)$ denotes an affine transformation followed by ReLU and dropout, and then mapped to output logits $\hat{\mathbf{y}} = g(\mathbf{h}_x)$. The prediction space follows the seven-class misinformation taxonomy from QUOTACLIMAT, excluding `0_not_relevant`.

Discourse-profile-aware models. We compare several mechanisms for incorporating discourse profile into misinformation prediction. In all cases, each discourse profile label $c \in \{1, \ldots, 6\}$ is mapped to a learnable embedding $\mathbf{d} = E(c)$.

Our primary discourse-profile-aware formulation extends the text-only baseline by incorporating discourse profile as an additional latent signal. The discourse-profile embedding is projected as $\mathbf{h}_d = f_d(\mathbf{d})$, then concatenated with the text representation $\mathbf{h}_x$ to form $\mathbf{h} = [\mathbf{h}_x; \mathbf{h}_d]$. A joint nonlinear transformation $\mathbf{z} = f_j(\mathbf{h})$ is then applied before the final classifier, yielding $\hat{\mathbf{y}} = g(\mathbf{z})$. This formulation introduces discourse profile only at the classifier stage while keeping the contextual text representation unchanged.

We also evaluate several recent personalization-inspired conditioning mechanisms, including prefix-based conditioning (Huber et al., 2025), multi-prompt conditioning (Doddapaneni et al., 2024), and adapter-based conditioning (Houlsby et al., 2019; Tan et al., 2024). Full architectural details are provided in Appendix D.

Isolating profile signal. To distinguish whether improvements arise from discourse-profile information or merely increased model capacity, we conduct a controlled profile permutation test in which discourse-profile labels are randomly shuffled while keeping the model architecture and training procedure fixed. We repeat the shuffled-profile experiment across 100 independent random permutations and compare the average shuffled-profile performance against the original model trained with the true discourse-profile assignments.

Robustness under stylistic variation. We additionally evaluate robustness under LLM-generated stylistic rewrites. Following Brian et al. (2025), we generate multi-round discourse-profile-conditioned variants that preserve the underlying misinformation category while varying rhetorical style and surface realization. We then evaluate both text-only and discourse-profile-aware classifiers on the rewritten comments. Generation details, prompts, and examples are provided in Appendix E.

3.3 Discourse-Profile-Level Misinformation Susceptibility with LLM

We use discourse-profile-conditioned LLM agents to evaluate whether inferred discourse profiles capture systematic differences in misinformation susceptibility. We test whether an LLM prompted with discourse profile descriptions (Appendix F.2) exhibits misinformation acceptance patterns mirroring user comments.

We construct a balanced evaluation set of 350 misinformation claims from QUOTACLIMAT, sampling 50 claims from each of the 7 categories. For

each inferred discourse profile, we create a discourse profile card summarizing characteristic beliefs and rhetorical tendencies (Appendix F.1).

For discourse profile d, we define an LLM susceptibility score as the proportion of misinformation claims not rejected by the discourse-profile-conditioned LLM agent:

$$L_d = \frac{1}{|Q|} \sum_{q \in Q} \mathbb{I}[r_{d,q} \neq \text{reject}],$$

where Q is the set of evaluation claims, $r_{d,q}$ is the LLM's response for discourse profile d on claim q, and $\mathbb{I}[\cdot]$ denotes an indicator function.

Human–LLM Discourse Profile Alignment We assess whether discourse-profile-conditioned LLM susceptibility patterns align with those of real human discourse profiles. For each discourse profile d, we compute a human misinformation rate H_d, defined as the proportion of comments assigned to discourse profile d that are labeled as misinformation in CLIMATEDISC. We compute the LLM susceptibility score L_d previously defined. We then measure alignment between (H_d) and (L_d) using Pearson and Spearman correlations, where higher correlation indicates closer correspondence between human and LLM discourse profile behavior. We next examine whether these discourse-profile-level differences translate into distinct engagement trajectories.

3.4 Discourse-Profile Temporal Dynamics

To examine how discourse profile conditions the downstream impact of misinformation, we analyze engagement persistence at the video level. For each video in CLIMATEDISC we compute (i) its dominant discourse profile, defined as the majority profile among its comments, (ii) its misinformation rate, and (iii) engagement persistence, measured as the time required to reach 60% of total comments (Liu et al., 2023). Our dataset for this engagement persistence experiment contains 2,925 videos[1].

Formally, for a video v with N_v total comments, let $t_{v,i}$ denote the timestamp of the i-th comment. We define the cumulative number of comments received by time t as

$$S_v(t) = \sum_{i=1}^{N_v} \mathbb{I}[t_{v,i} \leq t],$$

where $\mathbb{I}[\cdot]$ is the indicator function. Let $t_v^{(0)}$ denote the timestamp of the first observed comment for video v. The time at which the video reaches a fraction α of its total comments is defined as $T_v(\alpha) = \inf\{t : S_v(t) \geq \alpha N_v\}$. We define engagement persistence as $P_v = T_v(0.6) - t_v^{(0)}$, which measures the duration (in hours) required for video v to accumulate 60% of its total comments. Larger values of P_v indicate more sustained engagement over time.

To examine whether persistence patterns differ across discourse contexts, we stratify videos into low, mid, and high misinformation groups by misinformation rate. Within each discourse profile type, we conduct Kruskal–Wallis tests to assess whether persistence distributions differ significantly across misinformation levels.

4 Experiments

Datasets

QUOTACLIMAT To obtain fine-grained misinformation types, we use QUOTACLIMAT (QuotaClimat, 2025), a curated dataset of 6,091 climate-related statements annotated with one of the misinformation categories specified in Appendix B.1. We use QUOTACLIMAT for two purposes: (i) training a misinformation classifier that is transferred to CLIMATEDISC comments (Appendix A), and (ii) susceptibility probing with discourse-profile-conditioned LLM agents using a balanced set of 350 claims (50 per category; Section 3.3).

CLIMATEDISC For discourse profile creation and downstream in-the-wild analyses, we rely on the CLIMATEDISC dataset (Gao et al., 2025), which comprises transcribed English speech from 7,110 TikTok videos and 116,256 corresponding comments tagged with "#climatechange," collected between January 2024 and November 2024. We use CLIMATEDISC in: (i) discourse profile creation (Section 3.1), (ii) discourse-profile-aware misinformation classification on in-the-wild comments (Section 3.2), (iii) human discourse-profile-level misinformation rates H_d for LLM alignment (Section 3.3), and (iv) temporal dynamics and engagement persistence (Section 3.4).

For the discourse-profile-aware misinformation classification experiments on CLIMATEDISC (Section 3.2), comments are divided into training, validation, and test sets in a 0.8/0.1/0.1 ratio.

Implementation Details For discourse-profile-aware misinformation classification, we use

[1]We exclude videos with zero or one comment, since persistence requires at least two comments to be defined.

Category	Method / Feature	Text Acc	Prof Acc	Δ Acc	Text F1	Prof F1	Δ F1
Lightweight MLP	*[Late Fusion]*						
	Bag-of-Words	0.701	0.718	+0.017	0.594	0.615	+0.021
	TF-IDF	0.684	0.718	+0.034	0.531	0.589	+0.058
	MPNet	0.789	0.809	+0.020	0.716	0.748	+0.032
RoBERTa-large	Late Fusion	0.830	0.850	+0.020	0.765	0.795	+0.030
	[Prior work: other conditioning methods]						
	Soft Prefix	0.830	0.834	+0.004	0.765	0.791	+0.026
	Multi-Prompt	0.830	0.833	+0.003	0.765	0.787	+0.022
	Adapter	0.830	0.838	+0.009	0.765	0.789	+0.023

Table 2: Comparison of text-only (**Text**) and discourse-profile-aware (**Prof**) models across model families and input settings. We compare our late-fusion discourse-profile integration against personalization-inspired conditioning strategies, including soft-prefix conditioning inspired by Huber et al. (2025), multi-prompt conditioning inspired by Doddapaneni et al. (2024), and adapter-based discourse-profile conditioning motivated by adapter and personalized parameter-efficient tuning methods (Houlsby et al., 2019; Tan et al., 2024).

lightweight feed-forward models with MPNet embeddings, bag-of-words features, and TF-IDF features as input. For higher-capacity contextual encoders, we instantiate text-only and discourse-profile-aware variants using RoBERTa-large. We use stratified 80/10/10 train/validation/test splits by misinformation label. Lightweight models are trained with Adam using learning rate 3×10^{-4}, batch size 32, and early stopping on validation macro-F1. RoBERTa-large models are trained with Adam using learning rate 2×10^{-5}, batch size 8, maximum sequence length 128, and early stopping on validation macro-F1. All models use cross-entropy loss and gradient clipping at 1.0.

Results

4.1 Discourse Profile Landscape

Figure 4 reports WCSS and Silhouette scores for $k \in [3, 30]$. WCSS decreases smoothly with k (no decisive elbow), while the Silhouette score attains a local maximum at $k = 21$, indicating comparatively strong separation at this granularity. We therefore set $k = 21$ as a practical trade-off between separation and interpretability. As a qualitative check, Appendix B.2 lists the ten comments nearest each centroid, which exhibit consistent topical and rhetorical patterns.

For interpretability in downstream analyses, we aggregate the 21 clusters into six higher-level discourse profiles plus a residual "Noise / Low-Signal" group. We exclude the Noise group from discourse-profile-sensitive analyses (it consists largely of short acknowledgments, emojis, and @-mentions). Table 1 summarizes the mapping. Additional statistics on cluster and discourse-profile size distributions are provided in Appendix C.

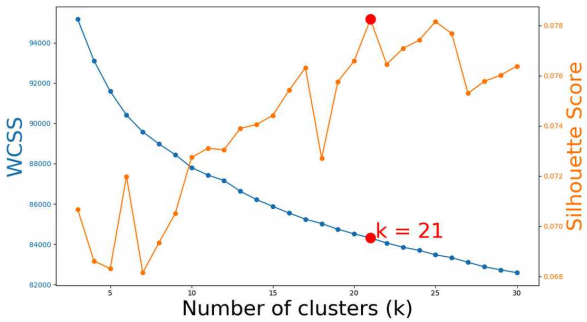

Figure 4: Within-Cluster Sum of Squares (WCSS; lower is better) and Silhouette Score (higher is better) across varying numbers of clusters k.

4.2 Misinformation Classification with Discourse Profile

We summarize all classification results in Table 2, comparing text-only (**Text**) and discourse-profile-aware (**Prof**) models across lightweight and pre-trained models.

For lightweight MLP-based models, incorporating discourse profile consistently improves performance across all feature types. Using bag-of-words and TF-IDF features, discourse profile improves accuracy by +0.014–0.034 and macro-F1 by +0.021-0.058[2]. Similar gains are observed with MPNet sentence embeddings (+0.020 accuracy and +0.032 macro-F1). These results suggest that discourse profile provides complementary predictive signal even when the underlying text representation is relatively shallow.

For higher-capacity contextual encoders, we compare several discourse-profile conditioning

[2]The lower TF-IDF text-only F1 relative to BoW is consistent with prior observations that TF-IDF weighting can underperform raw term counts on short, informal text where vocabulary is highly repetitive across documents.

strategies, including prefix-based conditioning, multi-prompt conditioning, and adapter-based conditioning, against our late-fusion formulation. While all discourse-profile-aware variants slightly outperform the text-only baseline, the gains from prefix-based and multi-prompt conditioning are relatively small. Adapter-based conditioning yields more consistent improvements, particularly in macro-F1. Our late-fusion formulation achieves the strongest overall performance, improving accuracy from 0.830 to 0.850 and macro-F1 from 0.765 to 0.795. Compared with methods that inject discourse profile directly into the contextualization process, late fusion produces larger and more stable gains, suggesting that discourse profile is most effective when treated as a complementary latent signal rather than as conditioning embedded within the token representation itself.

Profile permutation test. To isolate discourse-profile signal from added model capacity, we compare correct-profile conditioning against shuffled-profile controls averaged over 100 random permutations (Table 3). Correct profiles outperform shuffled controls across all representations: +0.019 macro-F1 for Bag-of-Words, +0.017 for TF-IDF, +0.021 for MPNet, and +0.014 for RoBERTa-large (all $p < 0.05$), confirming that gains reflect meaningful profile alignment rather than the presence of profile parameters alone. The effect is strongest for sparse lexical representations, but remains present for stronger neural encoders, suggesting that explicit discourse-profile conditioning captures complementary discourse-level signal not fully absorbed by text representations alone.

4.3 Discourse-profile-Level Susceptibility Probing with LLM

Human vs. LLM Agents Susceptibility We observe a strong positive correlation between human and LLM discourse profile susceptibility patterns. The Pearson correlation between H_d and L_d (specified in section 3.3) is 0.82, indicating substantial linear agreement. The Spearman rank correlation is even higher at 0.94, suggesting that the relative ordering of discourse profiles by misinformation susceptibility is highly consistent between humans and LLM agents. This alignment suggests that discourse-profile-conditioned LLM agents capture meaningful, discourse-profile-specific structure in misinformation acceptance.

Representation	Condition	Acc	F1
Bag-of-Words	Text only	0.701	0.594
	+ Shuffled profile	0.700	0.596
	+ Correct profile	**0.718**	**0.615***
TF-IDF	Text only	0.684	0.531
	+ Shuffled profile	0.696	0.572
	+ Correct profile	**0.718**	**0.589***
MPNet	Text only	0.789	0.716
	+ Shuffled profile	0.796	0.727
	+ Correct profile	**0.809**	**0.748***
RoBERTa-large	Text only	0.830	0.765
	+ Shuffled profile	0.832	0.781
	+ Correct profile	**0.850**	**0.795***

Table 3: Permutation-test evaluation of discourse-profile conditioning across text representations. Shuffled-profile results are averaged over 100 random profile-label permutations. * indicates that the correct-profile model significantly outperforms shuffled-profile controls in macro-F1, assessed with a t-test over the 100 shuffled-profile runs ($p < 0.05$).

Susceptibilities Across Discourse Profiles Table 4 summarizes overall susceptibility score (specified in section 3.3) per discourse profile. *Conspiratorial / Anti-Climate* shows very high susceptibility (96.3%), while other discourse profiles range from 4.9% to 29.4%. Notably, *Ideological / Political Combatants* show the lowest overall susceptibility (4.9%), suggesting that strong political engagement does not necessarily translate to greater vulnerability to misinformation. Similarly, the *Pro-Science / Rationalist* demonstrates comparatively low susceptibility (10.4%), though this group is not immune to misinformation, indicating that scientific self-identification alone does not guarantee resistance. *Doomism / Nihilism* is comparatively higher (22.0%), highlighting the role of affective orientations such as resignation or fatalism in shaping misinformation vulnerability.

Discourse Profile	Susceptible (%)
Local / Experiential Observers	29.4
Ideological / Political Combatants	4.9
Systemic / Structural Critics	8.0
Pro-Science / Rationalist	10.4
Doomism / Nihilism	22.0
Conspiratorial / Anti-Climate	**96.3**

Table 4: Overall misinformation susceptibility rate by discourse profile.

4.4 Engagement Dynamics Across Profiles

To examine how persistence co-varies with misinformation prevalence, we group videos into Low, Mid, and High misinformation bins and compute

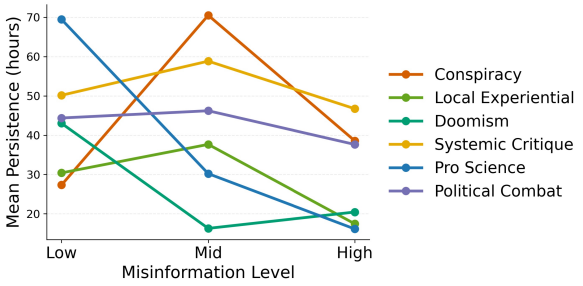

Figure 5: Mean persistence hours across misinformation level (Low, Mid, High) for each profile type.

mean persistence hours for each profile-bin combination (Figure 5). For *Pro Science*, persistence is higher in the low-misinformation group (69.5 hours) and lower in the high-misinformation group (16.1 hours). Within the *Conspiracy* profile, the high-misinformation group exhibits higher mean persistence than the low-misinformation group. *Systemic Critique* exhibits relatively small variation across bins. *Doomism* and *Local Experiential* show lower persistence in high-misinformation groups. For example, *Doomism* decreases from 43.1 hours (Low) to 20.3 hours (High), while *Local Experiential* increases from Low (30.5 hours) to Mid (37.3 hours) before decreasing at High (16.8 hours).

Discourse Profile	H Statistic	p-value
Systemic Critique	20.501	$< 0.001^{***}$
Local Experiential	16.336	$< 0.001^{***}$
Conspiracy	9.237	0.010^{**}
Pro Science	4.875	0.047^{*}
Political Combat	3.850	0.146
Doomism	2.912	0.233

Table 5: Within-profile Kruskal–Wallis test results examining whether persistence differs across misinformation tertiles (Low, Mid, High). Significance levels are denoted as $^{***}p < 0.001$, $^{**}p < 0.01$, and $^{*}p < 0.05$.

To assess whether these differences reflect statistically significant heterogeneity, we conduct within-profile Kruskal–Wallis tests comparing persistence across misinformation levels. Table 5 shows that four of six profiles (*Systemic Critique*, *Local Experiential*, *Conspiracy*, and *Pro Science*) exhibit significant differences across misinformation tertiles at $\alpha = 0.05$, whereas *Political Combat* and *Doomism* do not.

5 Discussion and Conclusion

This work demonstrates that discourse profile is a useful, interpretable axis for climate misinformation research. By modeling how comments are framed in addition to what is said, we show that discourse profiles capture discourse-level regularities that complement comment text.

Our results show that discourse profile improves misinformation classification in both lightweight and high-capacity settings, but its effect depends on how the signal is incorporated. For lightweight MLP models, adding discourse profile consistently improves performance across bag-of-words, TF-IDF, and MPNet representations, indicating that discourse-level structure provides useful signal beyond lexical or sentence-level features. For larger contextual encoders, we compare multiple profile-conditioning strategies, including prefix-based, multi-prompt, adapter-based, and late-fusion approaches. Late fusion achieves the strongest gains, suggesting that discourse profile is most effective when treated as a complementary latent variable rather than input-side conditioning alone. Profile-label shuffle tests confirm that the gains come from meaningful profile alignment, not merely added parameters.

We also show that discourse-profile-conditioned LLM agents can serve as controlled probes of misinformation susceptibility. Holding claims constant while varying discourse profile produces systematic differences in acceptance behavior that align with human discourse-profile-level vulnerability. This suggests that discourse profiles provide a meaningful abstraction not only for classification, but also for analyzing susceptibility patterns.

Finally, our temporal dynamics analysis shows that engagement persistence co-varies with misinformation prevalence differently across discourse profiles. Some discourse contexts exhibit higher persistence in higher-misinformation settings, while others show the opposite pattern. These heterogeneous associations suggest that misinformation prevalence and engagement persistence are not uniformly distributed across discourse contexts, highlighting the importance of accounting for rhetorical variation when analyzing climate misinformation dynamics.

Together, these findings suggest that discourse-profile-aware modeling offers a principled way to surface discourse-level structure in climate misinformation discourse, unifying prediction, robustness analysis, susceptibility probing, and engagement dynamics within a single abstraction. Rather than replacing text-centric approaches, discourse profiles complement them by capturing stable patterns of expression and interaction that are difficult to recover from isolated comments alone.

Limitations

Discourse Profile Induction Validity and Subjectivity Our discourse profiles are induced via embedding-based k-means clustering and then manually consolidated from 21 clusters into six higher-level types. While this improves interpretability, the resulting taxonomy is not uniquely determined: different embedding models, clustering algorithms, random seeds, or k selection criteria could yield alternative clustering results. Further, consolidation relies on researcher judgment and may introduce subjectivity in where boundaries are drawn. As such, discourse profiles should be interpreted as dataset-specific rhetorical groupings rather than definitive or exhaustive categories of users.

We infer discourse profiles from observed comment text and do not have access to ground-truth user attributes (e.g., demographics, ideology, geography, or identity) or off-platform behavior. Consequently, discourse profile labels should not be interpreted as psychological traits or real-world subpopulations; they reflect recurring linguistic and rhetorical patterns expressed in public TikTok comments within our collection.

Cross-Domain Misinformation Label Transfer Our misinformation labels on TikTok comments are produced by a classifier trained on QUOTACLIMAT and applied to in-the-wild TikTok discourse. Although we validate transfer using semantic claim–comment matching, this validation relies on a cosine-similarity threshold selected via qualitative inspection and yields a relatively small set of matched pairs. Moreover, the QUOTACLIMAT taxonomy is claim-centric and may not perfectly align with how misinformation is expressed in conversational comments (e.g., sarcasm, implicature, partial agreement, or off-topic replies). As a result, predicted misinformation labels likely contain noise, which could affect downstream discourse profile distributions, susceptibility estimates, and persistence analyses.

LLM Agents Susceptibility Our LLM-based susceptibility results should be interpreted as a controlled probe rather than a behavioral simulation of real users. Outputs can be sensitive to prompt wording, decoding settings, and the specific model used, and we evaluate only a single LLM configuration. Furthermore, our metric captures refusal/acceptance patterns under instruction rather than real-world exposure, incentives, or social context. We therefore emphasize alignment in relative discourse profile ordering rather than treating absolute LLM susceptibility rates as estimates of human vulnerability.

Temporal Dynamics and Causality Our persistence analysis faces important methodological constraints. First, persistence hours conflate organic sharing velocity, algorithmic promotion, and content moderation decisions, precluding mechanistic attribution. Second, our cross-sectional design cannot establish whether misinformation density affects persistence or whether longer-persisting videos accumulate more misinformation over time. Third, video-level aggregation may obscure important dynamics at finer temporal scales.

Platform and Data Constraints TikTok's proprietary recommendation algorithm and content moderation policies are opaque, limiting our ability to interpret persistence patterns mechanistically. We cannot observe deleted, shadowbanned, or algorithmically suppressed content, potentially biasing our sample toward platform-permitted content.

Ethical Considerations Our analysis focuses on public comments and does not identify or track individual users. Nevertheless, discourse profile taxonomies could potentially be misinterpreted or misused as a form of user profiling. We emphasize that this is neither our intent nor what our method does. The discourse profiles in this work are comment-level abstractions induced from recurring rhetorical patterns in public discourse, not demographic, psychological, or persistent individual traits. A user may express different discourse patterns across comments or contexts, and our analysis does not observe or infer real-world user identities, beliefs, or susceptibility. Accordingly, discourse profiles should be used to study aggregate discourse patterns and model behavior, not to target or make claims about individual users.

Acknowledgments

Personally Identifiable / Offensive Information We anonymize all potentially personally identifiable information in our released materials, including file paths and organization names in the code. In addition, examples of user comments shown in the paper have been sanitized to remove profane or offensive language.

AI Assistants in Writing We used ChatGPT to assist with language polishing and clarity improvements throughout the manuscript, to refine the LaTeX formatting of several figures, and to support code refactoring and readability improvements in the analysis pipeline. The tool was used solely for assistance with writing, formatting, and code refactoring, and did not contribute to the research design, modeling decisions, data analysis, or empirical results.

References

Emily Allaway and Kathleen McKeown. 2021. Zero-shot stance detection: A dataset and model analysis. In *NAACL*.

Lisa Argyle, Ethan Busby, Nancy Fulda, Joshua Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31:1–15.

Clara Baltasar, Sergio D'Antonio Maceiras, Alejandro Martín, and David Camacho. 2024. Analysis of climate change misleading information in tiktok. In *Proceedings of the 1st Workshop on COuntering Disinformation with Artificial Intelligence (CODAI), co-located with the 27th European Conference on Artificial Intelligence (ECAI 2024)*, volume 3782 of *CEUR Workshop Proceedings*, pages 54–61, Santiago de Compostela, Spain. CEUR-WS.org. Creative Commons CC BY 4.0.

Guy Barel, Oren Tsur, and Dan Vilenchik. 2025. Acquired TASTE: Multimodal stance detection with textual and structural embeddings. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6492–6504, Abu Dhabi, UAE. Association for Computational Linguistics.

Corey Basch, Bhavya Yalamanchili, and Joseph Fera. 2021. climate change on tiktok: A content analysis of videos. *Journal of Community Health*, 47:163–167.

Adrian Benton and Mark Dredze. 2018a. Using author embeddings to improve tweet stance classification. In *Proceedings of the 2018 EMNLP Workshop W-NUT: The 4th Workshop on Noisy User-generated Text*, pages 184–194, Brussels, Belgium. Association for Computational Linguistics.

Adrian Benton and Mark Dredze. 2018b. Using author embeddings to improve tweet stance classification. In *Proceedings of the 2018 EMNLP Workshop W-NUT: The 4th Workshop on Noisy User-generated Text*, pages 184–194, Brussels, Belgium. Association for Computational Linguistics.

Emma Frances Bloomfield and Denise Tillery. 2018. The circulation of climate change denial online: Rhetorical and networking strategies on facebook. *Environmental Communication*, 13:1–12.

Chong Jun Rong Brian, Yixuan Tang, and Anthony Kum Hoe Tung. 2025. MPCG: Multi-round persona-conditioned generation for modeling the evolution of misinformation with LLMs. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 34030–34064, Suzhou, China. Association for Computational Linguistics.

YeonJung Choi, Lanyu Shang, and Dong Wang. 2023. Climatemist: Climate change misinformation and stance detection dataset.

Travis Coan, Constantine Boussalis, John Cook, and Mirjam Odile Nanko. 2021. Computer-assisted classification of contrarian claims about climate change. *Scientific Reports*, 11.

Marco Del Tredici, Diego Marcheggiani, Sabine Schulte im Walde, and Raquel Fernández. 2019. You shall know a user by the company it keeps: Dynamic representations for social media users in NLP. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4707–4717, Hong Kong, China. Association for Computational Linguistics.

Sumanth Doddapaneni, Krishna Sayana, Ambarish Jash, Sukhdeep Sodhi, and Dima Kuzmin. 2024. User embedding model for personalized language prompting. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 124–131, St. Julians, Malta. Association for Computational Linguistics.

Ge Gao, Zhengyang Shan, James Crissman, Ekaterina Novozhilova, YuCheng Huang, Arti Ramanathan, Margrit Betke, and Derry Wijaya. 2025. Insights into climate change narratives: Emotional alignment and engagement analysis on TikTok. In *Proceedings of the Fourth Workshop on NLP for Positive Impact (NLP4PI)*, pages 128–143, Vienna, Austria. Association for Computational Linguistics.

Nir Grinberg, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. Fake news on twitter during the 2016 u.s. presidential election. *Science*, 363(6425):374–378.

Faye Holder, Sanober Mirza, Namson-Ngo-Lee, Jake Carbone, and Ruth E. McKie. 2023. Climate obstruction and facebook advertising: how a sample of climate obstruction organizations use social media to disseminate discourses of delay. *Climatic Change*, 176:1–21.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. *Preprint*, arXiv:1902.00751.

Bernd Huber, Ghazal Fazelnia, Andreas Damianou, Sebastian Peleato, Max Lefarov, Praveen Ravichandran, Marco De Nadai, Mounia Lalmas-Roellke,

and Paul N. Bennett. 2025. Embedding-to-prefix: Parameter-efficient personalization for pre-trained large language models. *Preprint*, arXiv:2505.17051.

Md Rafiqul Islam, Shaowu Liu, Xianzhi Wang, and Guandong Xu. 2020. Deep learning for misinformation detection on online social networks: a survey and new perspectives. *Social Network Analysis and Mining*, 10.

K. Jones. 2004. A statistical interpretation of term specificity in retrieval. *Journal of Documentation*, 60:493–502.

Ranjay Krishna, Donsuk Lee, Li Fei-Fei, and Michael Bernstein. 2022. Socially situated artificial intelligence enables learning from human interaction. *Proceedings of the National Academy of Sciences*, 119.

William Lamb, Giulio Mattioli, Sebastian Levi, J. Roberts, Stuart Capstick, Felix Creutzig, Jan Minx, Finn Müller-Hansen, Trevor Culhane, and Julia Steinberger. 2020. Discourses of climate delay. *Global Sustainability*, 3.

Junyi Li, Charith Peris, Ninareh Mehrabi, Palash Goyal, Kai-Wei Chang, Aram Galstyan, Richard Zemel, and Rahul Gupta. 2024. The steerability of large language models toward data-driven personas. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7290–7305, Mexico City, Mexico. Association for Computational Linguistics.

Maggie Liu, Jing Wang, and Daniel Preotiuc-Pietro. 2023. Analyzing and predicting persistence of news tweets. 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 719–732, Nusa Dua, Bali. Association for Computational Linguistics.

S. Lloyd. 1982. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137.

Veronica Lynn, Salvatore Giorgi, Niranjan Balasubramanian, and H. Andrew Schwartz. 2019. Tweet classification without the tweet: An empirical examination of user versus document attributes. In *Proceedings of the Third Workshop on Natural Language Processing and Computational Social Science*, pages 18–28, Minneapolis, Minnesota. Association for Computational Linguistics.

Nicola Marsden and Monika Pröbster. 2019. Personas and identity looking at multiple identities to inform the construction of personas.

S. Mehta and 1 others. 2022. Continually improving social context representations for fake news detection. In *ACL*.

Thanh Tam Nguyen, Ee-Peng Lim, and 1 others. 2020. Fang: Leveraging social context for fake news detection using graph representation. In *CIKM*.

Joon Park, Joseph O'Brien, Carrie Cai, Meredith Morris, Percy Liang, and Michael Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. pages 1–22.

Warren Pearce, Sabine Niederer, Suay Ozkula, and Natalia Sanchez-Querubin. 2018. The social media life of climate change: Platforms, publics, and future imaginaries. *Wiley Interdisciplinary Reviews: Climate Change*, 10.

Arianna Pera and Luca Aiello. 2024. Shifting climates: Climate change communication from youtube to tiktok. pages 376–381.

Minghui Qiu, Yanchuan Sim, Noah A. Smith, and Jing Jiang. *Modeling User Arguments, Interactions, and Attributes for Stance Prediction in Online Debate Forums*, pages 855–863.

QuotaClimat. 2025. Frugal ai challenge text train dataset. Dataset for the Frugal AI Challenge 2025, focused on climate disinformation narratives.

Nils Reimers and Iryna Gurevych. 2020. all-mpnet-base-v2: Sentence-transformers model. Fine-tuned for sentence embeddings on 1B+ sentence pairs.

Cristian Rojas, Frank Algra-Maschio, Mark Andrejevic, Travis Coan, John Cook, and Yuan-Fang Li. 2024. Hierarchical machine learning models can identify stimuli of climate change misinformation on social media. *Communications Earth Environment*, 5.

Bebe Chand Sultana, Md Prodhan, Edris Alam, Md Sohel, A. B. M. Mainul Bari, Subodh Pal, Md Islam, and Abu Islam. 2024. A systematic review of the nexus between climate change and social media: present status, trends, and future challenges. *Frontiers in Communication*, 9:1301400.

Yanqing Sun and Juan Xie. 2024. Who shares misinformation on social media? a meta-analysis of individual traits related to misinformation sharing. *Computers in Human Behavior*, 158:108271.

Zhaoxuan Tan, Zheyuan Liu, and Meng Jiang. 2024. Personalized pieces: Efficient personalized large language models through collaborative efforts. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6459–6475, Miami, Florida, USA. Association for Computational Linguistics.

Apoorva Upadhyaya, Marco Fisichella, and Wolfgang Nejdl. 2023. Towards sentiment and temporal aided stance detection of climate change tweets. *Information Processing Management*, 60:103325.

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.

Y. Zhang and 1 others. 2024. Modeling user stance with bipartite graph neural networks. In *ICWSM*.

A Taxonomy Transfer Details

This appendix describes how we obtain fine-grained misinformation-type labels for in-the-wild comments by transferring a curated taxonomy to the CLIMATEDISC corpus.

A.1 Misinformation Classifier Training

We fine-tune a RoBERTa-large classifier on QUOTACLIMAT to predict one of eight climate misinformation categories from text. Training is performed on a single NVIDIA A100 GPU for four epochs with batch size 128 and learning rate 2×10^{-5}. Our model reaches a test accuracy of 0.75 on the eight-class misinformation classification task, indicating that the model is able to reliably distinguish among the fine-grained categories defined in QUOTACLIMAT.

A.2 Cross-domain Validation via Claim–Comment Matching

To validate transfer from curated claims to in-the-wild comments, we embed both QUOTACLIMAT claims and CLIMATEDISC comments using the same sentence transformer and identify high-similarity claim–comment pairs via cosine similarity. Pairs exceeding a fixed threshold of $\tau = 0.85$ are treated as referring to the same underlying claim. We select $\tau = 0.85$ via qualitative inspection: it is the lowest value that consistently yields pairs about the same underlying claim with aligned stance; lower thresholds increasingly admit topic-level matches with stance mismatches. This produces 208 matched comment–claim pairs (Table 6).

On these pairs, the classifier's prediction on the TikTok comment matches the ground-truth label of the paired QUOTACLIMAT claim 79% of the time (majority baseline 38%), with Cohen's $\kappa = 0.72$ indicating strong agreement beyond chance (Figure 7). Manual inspection and quantitative agreement metrics further support that matched pairs refer to the same underlying claim.

Based on this validation, we apply the classifier to all comments in CLIMATEDISC to characterize misinformation types in in-the-wild discourse.

To assess label quality beyond the high-similarity regime, where explicit claim alignment is unavailable, we conduct an additional agreement analysis using an independent LLM annotator. We randomly sample 200 comments from CLIMATEDISC that do not exceed the similarity threshold with any QUOTACLIMAT claim. Each comment is annotated independently by ChatGPT using a fixed prompt (Appendix F.4), and we compare these annotations with our classifier's predictions.

Across these unmatched comments, agreement reaches 89%, with Cohen's $\kappa = 0.54$. The substantial raw agreement and moderate κ suggest that the classifier's predictions remain broadly aligned with independent judgments even when comments cannot be directly grounded to curated claims. Together, these results indicate that the model generalizes beyond claim-level paraphrases to more heterogeneous and implicitly framed misinformation in real-world discourse.

Comment (ClimateDisc)	Matched Claim (Quota-Climat)
Climate change is a HOAX	Climate change is a complete hoax always has been.
Yes, climate change is real to what extent is arguable, but definitely changing	Climate change is definitely real. It changes every year from winter to autumn to spring to summer.
Climate change is very real. The climate has always changed, but in recent years it is changing too fast	Climate change is definitely real. It changes every year from winter to autumn to spring to summer.

Table 6: Representative matched comment–claim pairs obtained using a cosine-similarity threshold of 0.85.

Cohen's Kappa	Quality
>0.8	Almost Perfect Agreement
>0.6	Substantial Agreement
>0.4	Moderate Agreement
>0.2	Fair Agreement
0-0.2	Slight Agreement
<0	Almost No Agreement

Table 7: Interpretation for different ranges of the Cohen's Kappa values.

B Datasets

B.1 Climate Misinformation Types

Below are the types of Climate Change Misinformation we used in QUOTACLIMAT.

0_not_relevant: No relevant claim detected or claims that don't fit other categories.

1_not_happening: Claims denying the occurrence of global warming and its effects - Global warming is not happening. Climate change is NOT leading to melting ice (such as glaciers, sea ice, and permafrost), increased extreme weather, or rising sea levels. Cold

weather also shows that climate change is not happening.

2_not_human: Claims denying human responsibility in climate change - Greenhouse gases from humans are not the causing climate change.

3_not_bad: Claims minimizing or denying negative impacts of climate change - The impacts of climate change will not be bad and might even be beneficial.

4_solutions_harmful_unnecessary: Claims against climate solutions - Climate solutions are harmful or unnecessary.

5_science_is_unreliable: Claims questioning climate science validity - Climate science is uncertain, unsound, unreliable, or biased.

6_proponents_biased: Claims attacking climate scientists and activists - Climate scientists and proponents of climate action are alarmist, biased, wrong, hypocritical, corrupt, and/or politically motivated.

7_fossil_fuels_needed: Claims promoting fossil fuel necessity - We need fossil fuels for economic growth, prosperity, and to maintain our standard of living.

B.2 Representative Comments of Each Cluster

Below are the top ten comments closest to each cluster centroid.

Cluster 0 (Partisan Anti-Trump/Anti-GOP Rants):

1. He's a complete delusional fool, druggy, grapist conspiracy theorist, worm brain lunatic. This POS will never be sworn in. His whole family disowned this crazed weed [clown emoji].
2. lol u talkin bout Biden?
3. this guy is full of s*** Trump's been helping out more for the American people then Joe Biden and Kamala
4. My god Vance just can't answer without lying, there is no way that man should be anywhere near the White House, he's pathetic [skull emoji].
5. He's clueless and all his facts. He's the most uneducated person we ever had.
6. This man is really ignorant. He doesn't understand what he's even talking about.
7. Again, WTF?! He is such a moron!!
8. He is part of the problem and can't understand his days are over and his grandchildren will suffer for his actions.
9. Trump is so stupid
10. He is so vile and evil!! Typical conman!

Cluster 1 (Local Weather Anecdotes):

1. you know you right Cuz. it's unseasonably warm where I live too [laughing emoji].
2. This winter in Michigan there was only snow for a few weeks in January and now it's 50 degrees and sunny.
3. It was concerning in Michigan yesterday (60s) then it snowed today [laughing emoji].
4. I'm in Wisconsin and it was like 90 last week and now it's in the 70s/50s.
5. We're having an unseasonably warm winter here in Minnesota.
6. we went from rainy to like 65 degrees, what is this summer?
7. It's actually scary the climate is warm right now, especially since I am Canadian and not used to this [sweat smile emoji].
8. In Utah we hit record lows in November, I'm so over it [crying emoji].
9. it's nice right now but Thursday/Friday will be almost 80 in Massachusetts [crying emoji] damn you global warming [victory hand emoji].
10. south west area is getting record rain and snow.

Cluster 2 (Emoji/Applause Spam — converted):

1. TRUMP [100 emoji][100 emoji][100 emoji][100 emoji][100 emoji]
2. [blue heart emoji][blue heart emoji][blue heart emoji][blue heart emoji] P
3. sindio [laughing emoji][laughing emoji][laughing emoji]
4. [thumbs up emoji][thumbs up emoji][ok emoji][party emoji]
5. [target emoji][target emoji][fire emoji][fire emoji][fire emoji]
6. [100 emoji][100 emoji][100 emoji][blue heart emoji][blue heart emoji]
7. [screaming emoji] unvorstellbar
8. [thumbs up emoji][thumbs up emoji][clap emoji x many]
9. [clap emoji][clap emoji][clap emoji][clap emoji]
10. [facepalm emoji][facepalm emoji][facepalm emoji]

Cluster 3 (Minimal Agreement Tokens):

1. Ya
2. ya
3. Right
4. Right
5. RIGHT
6. Right
7. Right
8. Right
9. Right
10. RIGHT

Cluster 4 (Consumption Plastics / Recycling):

1. clothes or items that can't be recycled get buried or dumped.
2. Stop buying plastic whenever possible.
3. I wish I could buy products not completely covered in plastic.
4. reuse, repurpose, buy less, switch to glass or metal; 2024 no new buying challenge.
5. so much usable stuff goes to the dump.
6. buying second hand clothes to save things from landfill [sweat smile emoji].
7. buying reusable plastic to replace paper is not a good idea.
8. milk in plastic jugs, bread in plastic bags, lunch meat in plastic [laughing emoji].
9. all ends up in landfill anyway.
10. I'm sure they don't even recycle any of it.

Cluster 5 (Polite Gratitude Replies):

1. Thank you!
2. Thank you!
3. Thank you!
4. Thank you!
5. Thank you!
6. Thank you!
7. Thank you!
8. Thank you!
9. thank you!
10. Thank you!

Cluster 6 (Short Emotional Reactions):

1. So true [laughing emoji]
2. So cool
3. So scary [crying face emoji]
4. SO TRUE [eyes emoji]
5. So happy [heart emoji]
6. So sad
7. So cute
8. So sad [broken heart emoji]
9. so nice
10. Boosting [blue heart emoji]

Cluster 7 (Explainers + Conspiracies):

1. It's science happens every 10 years and it used to be all water.
2. thermal expansion and added water from land ice.
3. Not a man made phenomenon.
4. land ice and thermal expansion.
5. nothing to do with spraying of skies and geoengineering.
6. water freezes and expands in limestone zones.
7. before this it was acid rain.
8. local event caused by warming Arctic waters.
9. reduced emissions reduced reflective particles, warming the ocean.
10. happens when exiting an ice age.

Cluster 8 (Doomism / Earth Fine, Humans Won't):

1. another great flood, ice age, warming cycle coming.
2. ecological and societal collapse is on the way.
3. the earth will rid itself of us.
4. earth will adapt; we will not survive.
5. Science people, we are doomed.
6. near the environmental point of no return.
7. I'm worried about mankind's future.
8. we will die before it affects earth.

9. perilous times like never before.
10. we must wake up and restore habitability.

Cluster 9 (Excited Agreement Bursts):

1. Omg yes
2. Yessss oh my god
3. yesssss
4. OMG
5. You said it
6. Yesss omg
7. OMG
8. Nailed it
9. Yes Eva
10. yes love that

Cluster 10 (Shared Anxiety Coping):

1. I'm so sad and afraid [blue heart emoji][strong emoji].
2. we must be aware but also able to live; maybe you're here to bring change.
3. Say a prayer [pray emoji].
4. I can't go on like normal.
5. always have hope.
6. I'm coping too.
7. same, persist with compassion.
8. same, I feel terrible.
9. I'm scared and sad, I don't want to live in trying times.
10. Me too, it's too much.

Cluster 11 (Government Blame / Anarchy Talk):

1. are you serious? it's our governments.
2. idiots in charge do nothing.
3. kicking indigenous peoples off lands.
4. governments don't see this as a problem.
5. manipulation of the weak.
6. HAARP government control.
7. this is what anarchy is, people helping each other.
8. fight through organizing and political action.
9. tell political leaders, not us.
10. US government and capitalism accelerated this.

Cluster 12 (Storms Flooding):

1. back-to-back storms caused flooding.
2. climatologists warned for decades.
3. storms stronger and more frequent.
4. flash flooding moving north.
5. happened in less than an hour.
6. regular rain causing floods.
7. houses awaiting repairs from hurricanes.
8. NYC flooding commentary.
9. unprecedented Midwest flooding.
10. every US region has natural disasters.

Cluster 13 (Geo-Self Mentions):

1. I'm in Tennessee.
2. you're in my state now.
3. Pennsylvania, same dread.
4. be glad you're not in FL.
5. I live in Visalia.
6. don't say this, I live there.
7. I live here, first time hearing this.
8. Southern Illinois.
9. here in PA too.
10. I'm in Oregon.

Cluster 14 (Pro-Science / Anti-Misinformation):

1. proves you don't know what you're talking about.
2. you say peer reviewed studies are wrong.
3. misinformation spreads easily.
4. not okay to think yourself smarter than scientists.
5. will never do their own research.
6. you trust a convicted felon serial grapist? got it.
7. you have no idea but skeptics are dumb?
8. simple rebuke is always wrong.
9. these people don't believe in science.
10. knowledge kills their arguments.

Cluster 15 (@-Mentions):

1. @sarah lewis
2. @psychicmediumjen
3. @Raymel Poche Jr.
4. @eliza koch
5. @MaKenzie Neas
6. @H.k.lynne
7. @emily xx
8. @Rachel Taylor
9. @aisela
10. @xottic_daa

Cluster 16 (Climate Denial / Manipulation):

1. Climate change is a joke. You are brainwashed.
2. THERE IS NO SUCH THING AS CLIMATE CHANGE!!
3. f u and climate change it's a f**k** farce; you must be an a**h*** to believe it.
4. Cut the crap. Nobody believes that mess.
5. Climate change is a make-believe crisis.
6. No such thing as climate change.
7. It's weather manipulation.
8. climate change is a hoax and scam.
9. finally someone tells the truth about climate.
10. our idea of climate change is a lie.

Cluster 17 (Energy Policy Debate):

1. energy corporations greenwashing rhetoric.
2. ignoring massive fossil fuel emissions.
3. electric car carbon credits rant.
4. new refineries would be more efficient.
5. switch to renewables and nuclear.
6. promoting "clean" coal and gas.
7. outsourcing pollution abroad.
8. no more fossil fuel projects.
9. build nuclear if you care about environment.
10. greener energy benefits future generations.

Cluster 18 (Electoral Mobilization):

1. Republicans are delusional, vote blue.
2. Vote straight blue ticket.
3. Vote blue up and down the ballot.
4. must vote Harris 2024.
5. Vote Blue.
6. Republican voting blue across the board.
7. Vote blue, this is sickening.
8. Vote red, blue is crazy.
9. Vote all Democrats.
10. Vote blue across the board.

Cluster 19 (Food / Agriculture Diet):

1. need organic and regenerative farming.
2. rainforest deforestation due to animal agriculture.
3. Farming must stop but I need plant-based foods.
4. raising fuel instead of food.
5. plant-based diet is bs.
6. want to work in sustainable agriculture.
7. kill animals to grow plants.
8. forests cleared for soybean fields.
9. our habits destroy rainforests.
10. plant trees.

Cluster 20 (Anti-Corporation / Anti-Capitalism):

1. we cannot count on corporations and the wealthy.
2. fundamental problem with capitalism.
3. corporations keep you poor.
4. profit motive stands in the way.
5. corporations get tax breaks.
6. wealthy keep getting richer.
7. someone is buying what corporations sell.
8. greedflation from corporations.
9. capitalism is a plague.
10. corporations do not care about the people.

C Cluster and Discourse-Profile Size Distributions

Figure 6 shows that cluster and discourse profile sizes are relatively balanced: clusters range from 3,798–8,426 comments and discourse profiles from 3,992–19,019 comments. This reduces the risk that downstream effects are driven by a single dominant group.

D Prior Personalization-inspired Conditioning Methods

This appendix provides additional architectural details for the discourse-profile-aware conditioning methods in prior works. In all formulations, each discourse profile label $c \in \{1, \ldots, 6\}$ is mapped to a learnable embedding $\mathbf{d} = E(c)$.

Figure 6: Distribution of cluster/discourse profile sizes.

Prefix-based discourse profile conditioning. Inspired by recent soft-prefix personalization methods (Huber et al., 2025), discourse profile is injected directly into the model's input space. The discourse-profile embedding is projected into the contextual embedding dimension as $\mathbf{p} = W_p\mathbf{d}$, where $\mathbf{p}$ acts as a learned soft prefix token. Given token embeddings $[\mathbf{e}_1, \ldots, \mathbf{e}_n]$, the input sequence becomes $[\mathbf{p}; \mathbf{e}_1, \ldots, \mathbf{e}_n]$. This allows discourse profile to influence contextualization throughout the model stack rather than only at the classifier stage.

Multi-prompt discourse profile conditioning. Following recent soft-prompt personalization approaches (Doddapaneni et al., 2024), the multi-prompt formulation extends the single-prefix approach by mapping the discourse-profile embedding into multiple learned prompt vectors $[\mathbf{p}_1, \ldots, \mathbf{p}_k] = W_m\mathbf{d}$, where k denotes the number of prompt tokens. The resulting sequence becomes $[\mathbf{p}_1, \ldots, \mathbf{p}_k; \mathbf{e}_1, \ldots, \mathbf{e}_n]$. We set $k = 5$ in our experiments. Compared with the single-prefix formulation, this approach increases the capacity of discourse-profile conditioning by allowing multiple latent prompt vectors to jointly shape contextual representations.

Adapter-based discourse profile conditioning. Motivated by adapter-based and parameter-efficient personalization methods (Houlsby et al., 2019; Tan et al., 2024), the adapter formulation conditions the text representation through a discourse-profile-gated adapter transformation. Given contextual text representation $\mathbf{h}_x$, we compute an adapter output $\mathbf{a} = W_u\sigma(W_d\mathbf{h}_x)$, where W_d and W_u denote down- and up-projection matrices. The discourse-profile embedding produces a gating vector $\mathbf{g} = \sigma(W_g\mathbf{d})$, which modulates the adapter output element-wise, yielding the personalized representation $\tilde{\mathbf{h}}_x = \mathbf{h}_x + \mathbf{g} \odot \mathbf{a}$. The resulting representation $\tilde{\mathbf{h}}_x$ is then passed to the classifier.

E LLM-generated Stylistic Variants

To evaluate robustness under stylistic variation, we additionally construct LLM-generated rewrites of misinformation comments following the multi-round rewriting setup of Brian et al. (2025). Starting from an original comment, we generate three successive rounds of rewrites using GPT-5. Each rewrite is conditioned on (i) a target discourse profile description, (ii) the original misinformation label definition, (iii) the associated video transcript context, and (iv) the previously generated rounds. The model is instructed to preserve the underlying misinformation category while varying rhetorical style and surface realization according to the target discourse profile.

For each original comment, variants are generated separately for each target discourse profile. The resulting Round-3 rewrites therefore represent substantially evolved stylistic variants of the original claim while maintaining the same underlying misinformation category.

We evaluate both the text-only baseline (Model A) and the discourse-profile-aware late-fusion model (Model B) on the Round-3 variants. Performance decreases for both models relative to the original comments, consistent with prior findings that iterative LLM rewriting substantially degrades misinformation classification performance (Brian et al., 2025). However, discourse-profile-aware classification remains more robust under these conditions. On the Round-3 variants, incorporating discourse profile improves accuracy from 0.573 to 0.637 ($+0.064$, $p < 0.01$) and macro-F1 from 0.542 to 0.577 ($+0.034$, $p < 0.01$).

These results suggest that discourse-level behavioral patterns remain more stable than surface lexical and stylistic cues, allowing discourse-profile-aware models to better preserve predictive performance even when misinformation claims undergo substantial stylistic rewriting.

F LLM Prompts

F.1 Discourse Profile Card Construction

To construct discourse profile cards, we prompt ChatGPT to synthesize coherent, reusable discourse profile descriptions from clusters of real-world comments. Each cluster is represented by a set of example comments (in Appendix B.2) drawn from our dataset, along with a short cluster label summarizing its dominant theme.

The following prompt template is used to generate each discourse profile card. The cluster label and representative comments are injected into the `{cluster_description}` field.

```
Task:
You are given a cluster of social media
comments that share a common theme.
Your task is to abstract these comments
into a single discourse profile
description.

Instructions:
- Describe this discourse profile's
core beliefs and worldview.
- Describe their typical goals or
motivations when commenting.
- Describe their tone, style, and
emotional framing.
- Describe common behaviors, rhetorical
patterns, and phrases they use.
- Do not quote individual comments
directly.
- Write a single coherent discourse
profile description suitable for
conditioning a language model.

Cluster description and example
comments:
{cluster_description}
```

F.1.1 Discourse Profile Cards

Each discourse profile is represented as a natural-language description capturing characteristic beliefs, goals, tone, and linguistic style. These discourse profile cards are injected verbatim into the LLM prompt to condition responses.

```
Local / Experiential Observer
This discourse profile speaks from
personal experience and local
observation, often grounding climate
change in everyday life. They mention
where they live ('I'm in Wisconsin,'
'here in Florida') and describe weird
weather patterns ('78 degrees in
Vermont in February,' 'record floods in
the Midwest'). Their tone is
conversational, regional, and
observational, rarely ideological. They
might use interjections like 'damn,'
'lol,' or 'it's crazy right now.' They
often agree or commiserate with others
('same here,' 'we just had that too,'
'yeah it's scary'). Comments often
contain anecdotal or comparative
phrasing ('last year it was 20 and now
it's 70'). Their goal is to share lived
experience, connect with others, and
validate observations rather than argue
policy. Typical phrases: 'same here in
[state],' 'I can't believe this
weather,' 'we had floods too,' 'this is
so weird for this time of year.'

Ideological / Political Combatant
This discourse profile is fiercely
partisan, emotionally charged, and
motivated by anger or loyalty toward
political figures. They frequently
condemn political opponents (especially
Trump, MAGA, GOP) or rally support for
Democrats ('Vote Blue,' 'Harris 2024').
Their tone is confrontational,
moralistic, and full of insults or
exclamations ('He's pathetic!!', 'These
idiots in charge', 'Wake up!'). They
use repetition, capitalization, and
emojis (e.g., hearts, 100s) to
emphasize points. They often blend
outrage with mobilization ('We must
vote,' 'Republicans are delusional').
Common behaviors include replying to
misinformation with political blame,
tagging allies, and expressing outrage
about government inaction. Goals:
defend their political side, vent
frustration, and rally others to act.
Typical phrases: 'Vote blue across the
board,' 'These people are so blind,'
'We have to do something,' 'This man is
a liar,' 'Wake up America.'

Systemic / Structural Critic
This discourse profile focuses on
systemic issues such as capitalism,
corporate greed, consumption habits,
fossil fuels, and agriculture. They
sound thoughtful but frustrated, often
linking environmental problems to
larger structural causes ('corporations
are destroying the planet for profit,'
'capitalism is the problem,' 'we need
to reduce plastic and buy less stuff').
Their tone ranges from earnest to
cynical. They often use moral or
practical language ('reuse, repurpose,
buy quality,' 'the rich keep getting
richer,' 'fossil fuel industries caused
climate change knowingly'). They like
calling out hypocrisy ('greenwashing,'
'profit motive,' 'privatized profits,
socialized costs'). They may reference
renewable energy or veganism, and often
frame solutions as lifestyle and
collective change, not individual
denial. Common phrases: 'stop buying
plastic,' 'corporations dont care,' 'we
need systemic change,' 'the planet over
profits,' 'no more fossil fuel
projects.'

Pro-Science / Rationalist
```

This discourse profile champions evidence, critical thinking, and scientific literacy. They frequently rebut misinformation with confident, logical language (’read the research paper,’ ’peer-reviewed studies show,’ ’you’re not smarter than scientists’). Tone: rational, occasionally sarcastic, and slightly exasperated. They value education, expertise, and data, often reminding others to ’trust science’ or mocking anti-science arrogance. Their comments are relatively well-structured, with complete sentences and punctuation. They may sound academic or firm, using terms like ’cognitive dissonance,’ ’confirmation bias,’ or ’peer-reviewed.’ Goal: defend scientific truth and call out ignorance. Common phrases: ’you clearly havent read the study,’ ’this isnt opinion, its data,’ ’science isnt up for debate,’ ’knowledge kills their arguments,’ ’lower your confidence level a little.’

Doomist / Nihilist Environmentalist This discourse profile expresses despair about humanity’s future and sees collapse as inevitable. Tone: emotional, fatalistic, occasionally poetic or philosophical. They believe Earth will survive but humans wont (’the earth will heal, we wont,’ ’were doomed,’ ’mother nature is done with us’). They use emotive and reflective phrasing (’Im scared for the future,’ ’its already too late,’ ’the planet will be fine without us’). They may still show flashes of empathy or communal sorrow (’were all going to die,’ ’I just want us all to be okay’). Their goal isnt to persuade but to lament and emotionally process the loss. Typical phrases: ’were past the point of no return,’ ’the earth will rid itself of us,’ ’humans wont survive,’ ’its so sad but true,’ ’were doomed.’

Conspiratorial / Anti-Climate This discourse profile is skeptical of mainstream climate science and frequently frames climate change as exaggerated, fabricated, or caused by hidden manipulation. They may claim climate change is a hoax, a scam, or just natural cycles, and often shift explanations toward weather engineering or conspiracies (for example, ’HAARP,’ ’geoengineering,’ ’chemtrails,’ ’cloud seeding,’ or ’they control the weather’). Their tone is dismissive, combative, and confident, often using insults or ridicule (’wake up,’ ’brainwashed,’ ’you people are dense’). They may mix partial scientific-sounding terms with conspiratorial logic (for example, ’thermal expansion’ alongside ’sky spraying’). Their goal is to reject institutional narratives, challenge authority, and persuade others that the mainstream explanation is propaganda or control. Typical phrases: ’climate change is a hoax,’ ’its just natural cycles,’ ’they are manipulating the weather,’ ’wake up,’ ’do your research,’ ’follow the money,’ ’they want control.’

F.2 Discourse-Profile-Conditioned Stance Detection Prompt

We evaluate whether discourse-profile-conditioned LLM agents systematically agree with different types of climate misinformation using the following prompt template. The discourse profile card (Appendix F.1.1) is injected into the `{discourse profile}` field and the misinformation claim into the `{claim}` field.

```
Discourse Profile:
{discourse profile}

Task:
Would *this discourse profile* agree
with, hesitate about, or reject the
following claim?

Output ONLY one of:
- 1_agree
- 0_uncertain
- -1_reject

USER:
Claim: {claim}
Answer:
```

This constrained output format enables direct quantitative comparison of discourse-profile-specific agreement patterns across misinformation categories.

F.3 Prompt for LLM-evolved variant generation.

We use the following prompt to generate discourse-profile-conditioned variants:

> You rewrite one TikTok comment in a target discourse profile style.
>
> Target discourse profile: {discourse_profile_name}
> Discourse profile description: {discourse_profile_desc}
>
> Keep the same misinformation label: {label}
> Label definition: {label_desc}
>
> Original comment:

{claim}

Video transcript context:
{source}

Previous rounds (if any):
{prev_block}

Rules:

- Output exactly ONE rewritten comment.
- No notes, no explanation, no JSON, no labels.
- Keep it natural and concise.
- Max {max_words} words.
- Keep same underlying misinformation type.

F.4 Misinformation Agreement

The following prompt is used to obtain independent LLM annotations for assessing agreement on misinformation labels in comments without a high-similarity claim match.

```
You are annotating user-generated
comments for climate change
misinformation.

Each comment must be assigned exactly
one of the following eight mutually
exclusive categories, based on the
primary claim expressed or implied in
the comment.

Categories:

0_not_relevant: No relevant claim
detected or claims that don't fit other
categories

1_not_happening: Claims denying the
occurrence of global warming and its
effects (e.g., melting ice, extreme
weather, sea-level rise)

2_not_human: Claims denying human
responsibility for climate change

3_not_bad: Claims minimizing or denying
negative impacts of climate change

4_solutions_harmful_unnecessary: Claims
arguing that climate solutions are
harmful or unnecessary

5_science_is_unreliable: Claims
questioning the validity or reliability
of climate science

6_proponents_biased: Claims attacking
climate scientists or climate action
proponents as biased, alarmist, or
corrupt

7_fossil_fuels_needed: Claims asserting
that fossil fuels are necessary for
economic growth or modern living

Instruction: Assign the single category
that best fits the comment.

Output format: One label from
(0_not_relevant, 1_not_happening,
2_not_human, 3_not_bad,
4_solutions_harmful_unnecessary,
5_science_is_unreliable,
6_proponents_biased,
7_fossil_fuels_needed)

Comment: {comment text}
```