Response to Reviewer 2

We thank Reviewer for the detailed and rigorous critique. The comments touch on genuinely substantive issues and have prompted us to reconsider several aspects of the manuscript's framing, validation, and interpretive claims. We respond to each major concern below.

**Comment 1 Abstract**

In the abstract, enhance clarity and conciseness, minimize metaphorical expressions (“flows,” “conduits,” “zones of empathic resonance”), and provide greater methodological detail.

**Response1:** We agree that expressions such as "flows," "conduits," and "zones of empathic resonance" privilege rhetorical vividness over methodological precision in the abstract. In the revision, we will replace these with more operationally grounded language—describing, for instance, the directional emotion transition inference procedure and the network construction logic—while condensing the abstract to foreground the analytical design and key quantitative outputs rather than conceptual imagery.

**Comment 2 Introduction**

The introduction is excessively long and contains repetitive critiques of traditional sentiment analysis. Multiple paragraphs reiterate points such as the overly aggregate nature of sentiment and the lack of semantic grounding, without providing additional nuance. The writing sometimes emphasizes theoretical elegance at the expense of analytical clarity, which may challenge reader engagement. Furthermore, the term “mechanism” is not clearly defined early in the text. While the introduction is generally persuasive, it would benefit from greater conciseness, a more precise problem definition, and improved conceptual focus.

**Response2:** The reviewer correctly identifies that the critique of aggregate sentiment approaches is reiterated across multiple paragraphs without substantive progression. In the revision, we will consolidate these into a single, more precisely formulated paragraph and redirect the recovered space toward a tighter problem definition. We will also introduce an explicit operational definition of "mechanism" at its first appearance—specifying that we use the term to denote identifiable, replicable structural pathways through which emotional transitions can be traced across entity co-occurrence configurations, rather than invoking causal determination in a strong philosophical sense. This distinction is important and we acknowledge it was insufficiently drawn in the current draft.

**Comment 3 Literature review**

The literature review demonstrates competence and currency; however, it requires a more focused scope, clearer conceptual boundaries, and more selective critical analysis.

**Response3:** We will restructure the literature review around three clearly demarcated conceptual threads—temporal emotion dynamics in disaster discourse, NER-based semantic analysis, and emotion propagation in social networks—and prune citations and commentary that do not directly bear on the gap our study addresses. The goal is to sharpen the intellectual lineage rather than demonstrate breadth of coverage.

**Comment 4 Methodology**

The overall workflow, encompassing data collection, annotation, model fine-tuning, and network construction, is logical and well structured. However, the absence of quantitative validation for both emotion classification and named entity recognition (NER), such as F1 score, accuracy, or inter-annotator agreement, is a significant limitation. Assigning a single “dominant emotion” to each post oversimplifies the emotional complexity present in disaster discourse. The inference of “emotion

transitions" based solely on entity co-occurrence lacks formal justification and may conflate association with propagation. Furthermore, key modeling decisions, including edge weighting, threshold selection, and network pruning, are not sufficiently justified. In summary, while the methodology is ambitious and well-designed at a conceptual level, it lacks adequate validation, transparency, and robustness checks to fully substantiate the paper's mechanistic claims.

**Response4:** Response: These are the most substantive methodological criticisms and we address them directly.

Regarding validation metrics: the omission of quantitative performance evaluation for both the emotion classifier and the NER model is a genuine gap. In the revision, we will report precision, recall, F1 scores, and inter-annotator agreement (Cohen's κ) for the hybrid annotation process and both fine-tuned BERT models on the held-out test set. These metrics exist from our experimental pipeline but were not included in the submitted draft, an oversight we will correct.

Regarding the single dominant emotion assignment: we acknowledge that reducing each post to a single emotion label involves a simplification. We will add a methodological note clarifying that this is a practical operationalization necessary for network-level analysis, and will discuss its implications as a limitation. We will also report the distribution of confidence scores across emotion categories for a representative sample, to indicate the degree of emotional ambiguity present in the data.

Regarding emotion transition inference: the reviewer correctly notes that inferring directional emotion transitions from entity co-occurrence conflates associative co-presence with directional propagation. We will revise the framing of this procedure explicitly: the transitions we infer are probabilistic associative patterns across the corpus rather than formally validated causal pathways. The language of "propagation" will be qualified throughout, and the discussion of this method's limitations will be strengthened.

Regarding modeling decisions: we will add a dedicated paragraph in Section 3.2 justifying the edge weighting scheme (co-occurrence frequency), the confidence threshold (0.85) chosen for entity extraction, and the rationale for network pruning decisions, with reference to comparable methodological choices in the existing literature.

**Comment 5 Results**

The results are presented in a clear sequence, moving from descriptive statistics to emotional dynamics and ultimately to semantic–affective networks. Nevertheless, several significant limitations are evident. The findings are predominantly descriptive and visual, with minimal quantitative testing or statistical validation. Assertions regarding "emotion flow," "amplification," and "buffering" are based on co-occurrence patterns rather than being formally substantiated. The extensive use of complex network visualizations increases the risk of interpretive overreach and reader subjectivity. Furthermore, the lack of baselines or null models hinders the assessment of whether the observed patterns are distinctive or could occur by chance. Although positive emotion clusters are emphasized, their relative magnitude and robustness are not systematically quantified. In summary, while the results are comprehensive, coherent, and visually engaging, they remain largely exploratory. The strength of interpretive claims is not consistently supported by analytical rigor.

**Response5**: We accept that the results section relies heavily on visual interpretation and that several key claims—particularly regarding amplification and buffering—are not formally tested. In the revision, we will take the following steps. We will report quantitative descriptive statistics for the emotion distribution across entity categories (proportions, relative frequencies), rather than

relying solely on visual network rendering. We will qualify the language of "amplification" and "buffering" to reflect that these are interpretive characterizations of observed structural patterns—high-degree emotionally homogeneous nodes and bridging nodes connecting emotion clusters of contrasting valence—rather than formally demonstrated dynamic processes. Regarding null models: we will add a brief discussion of whether the observed network properties (particularly the clustering of positive emotions around institutional entities) depart from what would be expected under a random emotion assignment baseline, acknowledging that a formal null model comparison represents a direction for future methodological development that falls outside the current paper's scope but should be acknowledged.

**Comment 6 Discussion**

The discussion appears to overinterpret descriptive findings, particularly when inferring mechanisms such as amplification and buffering without formal causal evidence. While claims regarding nationalism, trust construction, and geopolitical effect are plausible, they remain untested empirically. Alternative explanations, including media agenda-setting, platform effects, censorship, and posting norms, receive insufficient consideration. The discussion reiterates conceptual contributions at length, resulting in redundancy. Although the discussion is thoughtful and theoretically ambitious, it extends beyond what the results can robustly support. Furthermore, only two references are cited in the discussion, which limits its scientific support.

**Response6**: The reviewer raises legitimate concerns on multiple fronts. We will revise the discussion to observe greater epistemic discipline, distinguishing more carefully between what the data demonstrate, what they suggest, and what remains speculative. Claims regarding digital nationalism, trust construction, and geopolitical affect will be reframed as theoretically motivated interpretations requiring further empirical testing rather than findings. Importantly, we will add a substantive paragraph addressing alternative explanations that the current analysis cannot rule out—including media agenda-setting effects (the role of CCTV and Weibo's algorithmic curation in shaping which entities became dominant), platform-level censorship and content moderation norms, and the possibility that observed emotion-entity patterns reflect keyword-based sampling bias from our data collection procedure. The reference list in the discussion will be substantially expanded. Redundant reiteration of conceptual contributions will be consolidated.

**Comment 7 Conclusion**

The conclusion section reiterates claims from the Discussion but does not synthesize insights at a broader conceptual level. The use of mechanistic and causal language is overstated, given the analysis's primarily descriptive nature. Although limitations are acknowledged, their implications for interpretation are not thoroughly examined. Suggestions for future research are concise but lack specificity. Overall, while the conclusion is coherent and well written, it adds little beyond summarization.

**Response7**: We will revise the conclusion to more accurately characterize the study's contribution as providing a methodological framework for structurally situated emotion analysis—one that generates hypotheses about affective pathways rather than demonstrates causal mechanisms. Future research directions will be made more specific, including: the incorporation of user interaction dynamics (retweet networks, reply cascades) to test whether entity-level emotion transitions correspond to inter-user diffusion patterns; extension to multimodal content (images, video) for richer affective signal detection; and cross-disaster comparative analysis to assess the generalizability of the emotion-entity coupling framework.

**Comment 8 Reviewer Suggestions**

The authors' efforts in preparing this manuscript are appreciated.

The manuscript addresses a significant and timely topic in disaster management. However, several substantive concerns remain.

The study relies heavily on emotion classification but does not report essential validation metrics (e.g., accuracy, Scores, inter-annotator agreement). Without quantitative evidence of model performance, it is difficult to assess the reliability of the core analytical outputs on which the conclusions depend.

The manuscript frequently describes the analysis as "mechanistic" and interprets emotion as "flowing" or "propagating" through entity networks. However, the empirical foundation for these claims is primarily based on co-occurrence patterns rather than on formally specified mechanisms, causal inference, or diffusion modeling. This discrepancy results in a gap between the strength of the claims and the evidentiary support provided.

Although the results are rich and visually compelling, many interpretations, particularly those concerning emotional amplification, buffering, trust construction, and digital nationalism, extend beyond what can be robustly inferred from the analyses presented.

Key concepts such as emotion propagation, semantic routing, and affective pathways are used extensively, yet remain insufficiently operationalized. Consequently, the manuscript at times blurs the distinction between analytical metaphor and empirical demonstration.

Collectively, these issues necessitate substantial reconceptualization, additional validation, and methodological strengthening that exceed the scope of a standard revision process.

**Response8**: We take this overarching critique seriously. The reviewer is right that there is a tension between the mechanistic framing of the paper and the associative nature of the underlying evidence. Our position is that this tension can be substantially reduced—though not fully eliminated—through more precise language, explicit operationalization of key constructs, and the addition of validation metrics. Specifically, we will: (a) replace "mechanism" with more epistemically modest formulations where the evidence is associative; (b) define "emotion propagation," "semantic routing," and "affective pathways" operationally as observable patterns in entity co-occurrence and emotion transition frequency, rather than as inferred causal processes; (c) add a methodological limitations section that explicitly addresses the inferential gap the reviewer identifies. We believe these revisions, combined with the addition of validation metrics and alternative-explanation discussion, bring the manuscript's claims into appropriate alignment with its evidentiary base. We respectfully maintain that the core analytical contribution—the emotion-entity coupling framework and its application to transboundary disaster discourse—remains substantive and publishable with these revisions in place.

We hope this response demonstrates our commitment to addressing the reviewer's concerns rigorously.