

Multi-agent Systems for Misinformation Lifecycle : Detection, Correction And Source Identification

Aditya Gautam

Independent Researcher
Seattle, Washington, USA
agautam7041@gmail.com

Abstract

The rapid proliferation of misinformation in digital media demands solutions that go beyond isolated Large Language Model(LLM) or AI Agent based detection methods. This paper introduces a novel multi-agent framework that covers the complete misinformation lifecycle: classification, detection, correction, and source verification to deliver more transparent and reliable outcomes. In contrast to single-agent or monolithic architectures, our approach employs five specialized agents: an Indexer agent for dynamically maintaining trusted repositories, a Classifier agent for labeling misinformation types, an Extractor agent for evidence based retrieval and ranking, a Corrector agent for generating fact-based correction and a Verification agent for validating outputs and tracking source credibility. Each agent can be individually evaluated and optimized, ensuring scalability and adaptability as new types of misinformation and data sources emerge. By decomposing the misinformation lifecycle into specialized agents - our framework enhances scalability, modularity, and explainability. This paper proposes a high-level system overview, agent design with emphasis on transparency, evidence-based outputs, and source provenance to support robust misinformation detection and correction at scale.

Introduction

Recent research underscores the growing sophistication of LLMs in identifying and countering misinformation, while also revealing critical gaps in their reliability (Wang et al. 2023b) , bias (Lin et al. 2025) and explainability (Cambria et al. 2024). Studies by chen2023combatingmisinformationagellms highlight the effectiveness of augmenting LLMs with external knowledge and tools for fact-checking, as well as fine-tuning and knowledge distillation. For instance, Wang et al. (2024) pioneered a knowledge distillation approach for multimodal misinformation detection, enhancing interpretability in complex image-text claims. Highlighting the benefits of multi-agent architectures, Li, Zhang, and Malthouse (2024) introduced FactAgent, breaking down fact-checking into specialized modules for evidence retrieval and source cross-referencing. To mitigate bias within these modules, Borah and Mihalcea (2024) employed ensemble methods, lever-

aging self-reflection to reduce discriminatory task assignments. Building on collaborative detection, Lakara et al. (2025b) proposed MAD-Sherlock, a debate-driven system that employs agent collaboration to reduce hallucinations and strengthen fact-verification, while Tian et al. (2024) integrated web-retrieval agents to boost detection over standalone models. Notably, Minici et al. (2025) refined disinformation campaign detection with IOHunter, achieving high precision in orchestrated network identification. Finally, Choudhary (2025) offered a comparative study on how LLMs handle political misinformation, revealing persistent challenges in grounding responses with credible sources. Collectively, these works confirm that LLM-driven techniques can significantly elevate detection capabilities for misleading content across multiple languages, modalities, and domains. In addition, Tang, Laban, and Durrett (2024) demonstrated how subtle linguistic manipulation can degrade LLM-based fact-checking, showcasing lingering vulnerabilities in single-agent systems.

Although existing agentic systems often excel in specialized tasks, they typically concentrate on detection alone and fall short of covering the entire misinformation lifecycle, which is needed to fully understand the proliferation and the extend of misinformation. Here the lifecycle refers to understanding different type of misinformation in the given claim, correcting it with other authentic sources, verification with robust reasoning model, getting all other sources where the misinformation is present on same context, and understanding the root cause i.e initial misinformation source (which may or may not be the input claim). This is a unique value addition that the proposed framework provides. To address the above mentioned gaps, this work introduce a five specialized agents—data source authentication and indexing, multi-class classification, evidence ranking with confidence scores, correction generation with cross validation, and source-focused verification. By separating source provenance tracking from core analytical tasks and implementing cross-agent audit trails, our framework aims to achieves higher modularity and interpretability than tightly coupled, domain-specific models. This end-to-end architecture directly tackles transparency and adversarial resilience challenges, establishing a comprehensive solution for managing misinformation from start to finish. Through dynamic knowledge indexing, focused agent specialization, and ex-

Explicit source audits, it aims to outperform monolithic or single-agent systems in accuracy, adaptability, and clarity. By continuously tracing misinformation from its origin to final verification, the framework ensures timely and reliable oversight in an ever-changing information environment, culminating in a flexible yet complete multi-agent paradigm that holistically tracks misinformation, provides vetted evidence, and generates corrections with confidence scores. This system is envisioned for use by governmental bodies, fact-checking organizations, and media platforms to monitor and understand misinformation spread on specific topics (e.g., warfare, elections). By allowing users to define topics of interest (e.g., via keywords or BM25 filters), the system can identify and verify content against established ground-truth sources, tracing the lineage and origin of misinformation based on publication and modification timestamps in the metadata.

Related work

Recent advancements in multi-agent LLM frameworks demonstrate significant progress in addressing misinformation through collaborative architectures and overcoming the limitation with standalone LLMs. Hybrid systems, such as LLM-Consensus Lakara et al. (2025a), improve explainability and reduce hallucinations by combining multi-agent reasoning with external retrieval. MAD-Sherlock Lakara et al. (2025b) introduces a debate-driven system where multi-modal agents assess contextual consistency in out-of-context (OOC) misinformation, achieving state-of-the-art accuracy without domain-specific fine-tuning by leveraging external retrieval and collaborative reasoning. Complementing this, MACAW framework by WU et al. (2024) employs three specialized agents i.e. Retrieval, Detective and Analyst to cross-validate multi-granularity evidence, improving OOC detection accuracy through structured workflows. FactAgent by Li, Zhang, and Malthouse (2024) modularize fact-checking into evidence retrieval, temporal verification, and source cross-referencing, enhancing interpretability in veracity assessment. Expanding scope, Tian et al. (2024) integrate web retrieval agents with LLMs, boosting detection F1-scores compared to standalone models, while Wang et al. (2023a) expose vulnerabilities in graph-based detectors using adversarial multi-agent reinforcement learning. Minici et al. (2025) advanced coordination detection with IOHunter, combining graph neural networks and LLMs to identify orchestrated disinformation campaigns. In this proposal, comprehensive and modular approach is taken to cater to all aspects of misinformation life-cycle i.e. origin, proliferation, detection, correction. This flexibility would allow this framework to be applied in any particular domain with plug and play.

Proposed Multi-agents Framework

The proposed architecture comprises five distinct agents working in a pipeline. First, the Classifier Agent analyzes the input claim and classifies it into a specific types of misinformation. This classification guides the Extractor Agent in querying the Indexer Agent’s comprehensive database to

retrieve relevant sources and their lineage. The Extractor Agent then ranks these sources based on authenticity, and alignment with the claim. The Corrector Agent, leveraging advanced reasoning capabilities, cross-validates the information and generates an accurate correction if misinformation is detected, along with supporting sources. Finally, the Verification Agent validates the outputs of the other agents to ensure the overall accuracy and coherence of the misinformation management process. The proposed multi-agent framework is designed for an end-to-end latency on the order of minutes. This accounts for the indexing pipeline for new content and the retrieval of relevant information for misinformation classification. While each agent operates within seconds to minutes, the overall system’s timeliness is acknowledged as a critical factor, especially for sensitive topics like elections or conflicts. For inter-agent coordination, two design patterns are possible i.e. centralized and decentralized approaches. A centralized model would feature a master agent managing communication and policy enforcement, ensuring adherence to guidelines at the cost of potential latency. A decentralized model, where agents communicate directly, offers lower latency but requires robust mechanisms to prevent miscoordination. The choice of architecture will be guided by further research into the trade-offs between control, latency, and system resilience. A detailed discussion on the chosen orchestration design will be included in future work detailing the system’s implementation.

- **Classifier Agent:** The Classifier Agent is designed to process any incoming content or article, which may or may not be misinformation. Its initial role is to perform a multi-class classification to detect if given context is misinformation or not, and if yes, of what types. This classification will be handled by a fine-tuned encoder-based model (e.g., RoBERTa) trained on a proprietary, internally labeled dataset curated by fact-checking teams. The Classifier Agent utilizes the text understanding and multi-class classification capabilities of LLMs. Its primary function is to analyze an input claim and categorize it into a predefined set of misinformation classes. This set includes categories such as statistical error, cherry-picking, propaganda, misrepresentation, historical manipulation, logical fallacy, factual error etc. Establishing a well-defined and comprehensive taxonomy of misinformation types is crucial for accurate classification and for guiding the subsequent actions of other agents, particularly the Extractor Agent. Classifier agent can have one or multiple LLM models with varying capabilities to detect misinformation classes, and use the ensembles based voting mechanism for boosting multi-classification performance. Carefully crafted prompt instructions along with fine-tuning on specific datasets can significantly influence the accuracy of the classification. The choice of LLMs itself is also critical, balancing accuracy requirements with computational cost.
- **Indexer Agent:** The Indexer Agent is responsible for indexing a wide range of data sources, ensuring that the indexed content is suitable for addressing the various types of misinformation identified by the Classifier

Agent. This may involve utilizing LLM capabilities for text understanding to facilitate metadata extraction from the indexed content. The agent indexes web pages, news articles, scientific databases, statistical datasets, historical documents, fact-checking archives, etc based on the problem at hand. A diverse and comprehensive index is essential to provide the Extractor Agent with the necessary resources to verify different kinds of claims. New authentic data sources like Google Data commons etc. can be best leveraged by this agent along with official data from the government like data.gov and different organizations like WHO, UNESCO etc. The Indexer Agent ensures that the indexed content includes appropriate metadata, such as the source, publication date, topic, and a potential reliability score. This metadata enhances the searchability and relevance of the indexed data for the Extractor Agent. Possible indexing techniques include traditional keyword-based indexing, and more advanced semantic indexing using embeddings generated by language models like Retrieval Augmented Generation (RAG). Establishing ground truth for misinformation detection involves several approaches. This agent is also responsible for generating metadata like authenticity score, topics, category etc based on the data-source description, usage, columns structures and other information like data source site description, usage, research citations etc. Additionally, developing automated methods for ground truth generation and validation, potentially by cross-referencing information across multiple highly reliable sources, is crucial for scaling up misinformation detection efforts. The Indexer Agent employs a multi-step pipeline that includes data cleaning, standardization of various content formats (HTML, PDF etc.) into a uniform structure with rich metadata (title, author, domain, links etc.). Content is then segmented using appropriate chunking strategies (e.g., semantic, fixed-size) optimal for the data and task. These chunks are converted into vector embeddings using fine-tuned models and stored in a vector database (e.g., using FAISS or similar technologies) to enable efficient (O(1) retrieval time) top-K similarity searches. This scalable architecture is standard in industry for RAG systems.

- **Extractor Agent:** The Extractor Agent leverages LLM capabilities in text understanding, information extraction, semantic similarity assessment, and ranking 9. It receives the misinformation type classification from the Classifier Agent and uses this information to refine its search strategy within the indexed data. Based on the classification, the Extractor Agent prioritizes querying specific sections of the indexed data or dedicated databases that are most relevant to the identified misinformation type. For instance, if the claim is classified as a statistical error, the agent will focus its queries on statistical databases and reports. For example, If the claim is a recent news article about some politicians or celebrity, it would fetch all the news articles related to this, and the statement made by the actors in context. Furthermore, the Extractor Agent adjusts its re-ranking strategy for extracted documents to prioritize those most perti-

nent to the specific type of misinformation. For historical misrepresentation, primary historical sources and expert analyses will be given higher priority. The agent retains its original functionalities of assigning authenticity scores to sources, potentially based on traditional information retrieval and ranking features like domain reputation, page rank, word frequency analysis, author reputation etc. It also continues to determine the alignment of the extracted sources with the input claim using semantic similarity measures and to identify the lineage or original source of the input information based on publication timestamps and content similarity.

- **Corrector Agent:** The Corrector Agent utilizes a strong reasoning LLM to take inputs from the Extractor Agent (claim, sources with authenticity scores, lineage etc.) and misinformation categories as identified by the Classifier Agent along with confidence score. Based on these inputs, it performs in-depth research and cross-validation using the extracted sources, and conditionally conducting additional searches i.e. encompasses both querying the internal, pre-indexed database more extensively (e.g., by retrieving a larger number of context documents) and performing external online searches via integrated tools (e.g., Google Search APIs) or querying other accessible external databases to gather further validating evidence. The agent then generates accurate corrections tailored to the specific type of misinformation identified. For example, if the misinformation is a statistical error, the Corrector Agent will aim to provide the correct statistical information along with its source, re-think about potential other sources to cross validate it, and append different citations based on the authenticity score or other predefined criteria. The Corrector Agent also provides reliable information for all the sources with metadata like timestamps etc. to ensure a lineage of misinformation and source identification can be done.
- **Verification Agent:** The Verification Agent is responsible for validating the outputs of the other agents against predefined criteria 3. These criteria may include logical consistency, adherence to source reliability thresholds, format specification, tone and the alignment between the generated correction and the identified misinformation type. The primary goal of this agent is to ensure the overall coherence and accuracy of the misinformation detection and correction process. It acts as a final quality check, mitigating potential errors or biases that may have been introduced by the other agents. The Verification Agent may utilize LLM capabilities for reasoning and text understanding to perform these validation checks. This is where human-in-loop would be best utilized to understand the performance of the human based verification and labeling and through this multi-agent framework. Verification agents can be potentially merged with Corrector agents to do verification and correction, depending on the use case, domain and complexity of the problem. This agent cross validate any additional information or subjective instructions provided by the users like tone, number of doctors for cross checking,

presentation format etc. and prepare the final response and potentially use different tools at the disposable to generate reports, add lineage information to spreadsheet, display it through charts and diagrams etc.

Advantages of the Proposed Framework

- **Systematic Evaluation:** Each agent can be independently monitored and fine-tuned, allowing focused performance assessments across different tasks.
- **Freshness:** A dynamically updated index ensures real-time adaptation to newly emerging misinformation patterns and sources.
- **Adaptability:** Modular design lets practitioners add or revise agents (e.g., introducing new data sources) without overhauling the entire system.
- **Specialization:** Each agent targets a distinct task—indexing, classification, extraction, correction, or verification— instead of “jack-of-all-trades” model.
- **Cost Optimization:** Resource usage is allocated based on agent complexity, minimizing overall computational overhead.
- **Knowledge Sharing:** Agents share findings through a unified communication layer, allowing seamless collaboration and evidence cross-referencing.
- **Strength Maximization:** Researchers can focus on upgrading individual components, such as classification or correction, without destabilizing the entire pipeline.
- **Robustness and Reliability:** The decomposition of tasks into specialized agents reduces single points of failure and improves error detection and correction at each stage.

Discussion

The proposed LLM-based multi-agent framework offers a structured and scalable approach to managing the full misinformation lifecycle—from classification and detection to correction and source verification. Unlike monolithic models that attempt to handle all tasks within a single architecture, this system distributes responsibilities across specialized agents, improving transparency, explainability, and robustness. The modular design allows for independent upgrading of agents (e.g., improving the Extractor Agent with more advanced retrieval models) without overhauling the entire system. This enhances adaptability to emerging misinformation patterns and evolving data sources. Moreover, by integrating real-time evidence retrieval, citation generation, and reasoning capabilities, the framework not only identifies misinformation but also corrects it with traceable justification—critical for user trust and accountability. However, challenges remain, including managing inter-agent coordination, mitigating latency introduced by multi-stage processing, and ensuring reliability of indexed data. Future iterations could benefit from incorporating self-evaluation loops and reinforcement learning to dynamically improve agent collaboration and performance. There are several other things that needs to very well thought through in Multi-agent

system like agents coordination protocols, collusion, policy violation, bias amplification etc, Even though the system paves the way for more resilient and interpretable management of misinformation, there are some risk as mentioned in detail by Hammond et al. (2025) and complexity that comes with multi-agent systems as mentioned in Cemri et al. (2025). Some of the challenges in the proposed framework:

- **Data Quality and Freshness:** The effectiveness of the Indexer Agent is contingent upon the quality and recency of the underlying data sources.
- **Coordination and Collusion:** Ensuring seamless coordination among agents is complex, especially as the number of agents increases. Poor coordination can lead to conflicts or redundant actions. Additionally, there’s a risk of agents colluding, intentionally or unintentionally.
- **Cost and Scalability:** The cost considerations for continuous indexing, complex reasoning, and ensemble verification could be substantial, particularly for large-scale deployment.
- **Network Effects:** The interdependent nature of agents means that the behavior of one agent can influence others, sometimes leading to unintended consequences or emergent behaviors that are difficult to predict and control.
- **Security and Privacy:** MAS often involve extensive data sharing among agents, raising concerns about data privacy and security. Unauthorized access or data breaches can lead to significant vulnerabilities, especially when agents operate across different platforms or organizations

Future work

This paper only proposed the multi-agent framework for misinformation lifecycle in this paper as the foundational system architecture. Future research will focus on implementing and empirically evaluating the proposed multi-agent framework using benchmark misinformation datasets such as FakeNewsNet (Shu et al. 2018), and WELFake (Verma et al. 2021). We acknowledge the complexity of deploying such a system, which necessitates robust infrastructure including real-time indexing, scalable databases, and comprehensive data scraping capabilities to build and maintain a reliable ground-truth repository. The current paper presents a high-level framework, and detailed implementation specifics, including the orchestration of these components, are part of ongoing and future development. Experiments will assess each agent’s performance individually and in pipeline mode, measuring metrics like accuracy, precision, recall, latency, and explainability. Additionally, ablation studies can explore the impact of agent specialization and external retrieval on system robustness. Future extensions may also incorporate multilingual capabilities, cross-modal misinformation detection, and user feedback loops for adaptive learning and continuous improvement.

References

Borah, A.; and Mihalcea, R. 2024. Towards Implicit Bias Detection and Mitigation in Multi-Agent LLM Interactions. arXiv:2410.02584.

Agent	Functionality	Key Benefits	LLM Capabilities
Classifier	Initial claim categorization i.e. misinformation categories, topic, sentiment	Provides high-confidence triggers to extract sources and grounds for other agents	Text understanding, multi-class classification, few-shot/zero-shot learning
Indexer	Database indexing, finding new sources, embedding and chunking	Efficiently categorizes information with source authenticity score	Text understanding and metadata generation based on data sources
Extractor	Information retrieval, source detection (timestamp & metadata analysis, lineage tracing, confidence score)	Offers comprehensive understanding of misinformation origin and spread	Text understanding, information extraction, semantic similarity, ranking, targeted querying
Corrector	Misinformation correction generation with cross validation from additional sources	Produces accurate and contextually relevant corrections	Strong reasoning, knowledge generation, cross-validation, in-context learning
Verification	Final veracity determination, formatting, comprehensive reports generation etc.	Ensures coherence, accuracy, and consistency of the overall assessment	Reasoning, text understanding, fact verification, logical consistency checks

Table 1: Characteristics and Functional Capabilities of Different Agents

Cambria, E.; Malandri, L.; Mercorio, F.; Nobani, N.; and Seveso, A. 2024. XAI meets LLMs: A Survey of the Relation between Explainable AI and Large Language Models. arXiv:2407.15248.

Cemri, M.; Pan, M. Z.; Yang, S.; Agrawal, L. A.; Chopra, B.; Tiwari, R.; Keutzer, K.; Parameswaran, A.; Klein, D.; Ramchandran, K.; Zaharia, M.; Gonzalez, J. E.; and Stolica, I. 2025. Why Do Multi-Agent LLM Systems Fail? arXiv:2503.13657.

Choudhary, T. 2025. Political Bias in Large Language Models: A Comparative Analysis of ChatGPT-4, Perplexity, Google Gemini, and Claude. *IEEE Access*, 13: 11341–11379.

Hammond, L.; Chan, A.; Clifton, J.; Hoelscher-Obermaier, J.; Khan, A.; McLean, E.; Smith, C.; Barfuss, W.; Foerster, J.; Gavenčiak, T.; Han, T. A.; Hughes, E.; Kovařík, V.; Kulveit, J.; Leibo, J. Z.; Oesterheld, C.; de Witt, C. S.; Shah, N.; Wellman, M.; Bova, P.; Cimpeanu, T.; Ezell, C.; Feuillade-Montixi, Q.; Franklin, M.; Kran, E.; Krawczuk, I.; Lamparth, M.; Lauffer, N.; Meinke, A.; Motwani, S.; Reuel, A.; Conitzer, V.; Dennis, M.; Gabriel, I.; Gleave, A.; Hadfield, G.; Haghtalab, N.; Kasirzadeh, A.; Krier, S.; Larson, K.; Lehman, J.; Parkes, D. C.; Piliouras, G.; and Rahwan, I. 2025. Multi-Agent Risks from Advanced AI. arXiv:2502.14143.

Lakara, K.; Channing, G.; Sock, J.; Rupprecht, C.; Torr, P.; Collomosse, J.; and de Witt, C. S. 2025a. LLM-Consensus: Multi-Agent Debate for Visual Misinformation Detection. arXiv:2410.20140.

Lakara, K.; Sock, J.; Rupprecht, C.; Torr, P.; Collomosse, J.; and de Witt, C. S. 2025b. MAD-Sherlock: Multi-Agent Debates for Out-of-Context Misinformation Detection.

Li, X.; Zhang, Y.; and Malthouse, E. C. 2024. Large Language Model Agent for Fake News Detection. arXiv:2405.01593.

Lin, L.; Wang, L.; Guo, J.; and Wong, K.-F. 2025. Investigating Bias in LLM-Based Bias Detection: Disparities between LLMs and Human Perception. In Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B. D.; and Schockaert, S., eds., *Proceedings of the 31st International Conference on Computational Linguistics*, 10634–10649. Abu Dhabi, UAE: Association for Computational Linguistics.

Minici, M.; Luceri, L.; Fabbri, F.; and Ferrara, E. 2025. IO-Hunter: Graph Foundation Model to Uncover Online Information Operations. arXiv:2412.14663.

Shu, K.; Mahudeswaran, D.; Wang, S.; Lee, D.; and Liu, H. 2018. FakeNewsNet: A Data Repository with News Content, Social Context and Dynamic Information for Studying Fake News on Social Media. *arXiv preprint arXiv:1809.01286*.

Tang, L.; Laban, P.; and Durrett, G. 2024. MiniCheck: Efficient Fact-Checking of LLMs on Grounding Documents. arXiv:2404.10774.

Tian, J.-J.; Yu, H.; Orlovskiy, Y.; Vergho, T.; Rivera, M.; Goel, M.; Yang, Z.; Godbout, J.-F.; Rabbany, R.; and Pelrine, K. 2024. Web Retrieval Agents for Evidence-Based Misinformation Detection. arXiv:2409.00009.

Verma, P. K.; Agrawal, P.; Amorim, I.; and Prodan, R. 2021. WELFake: Word Embedding Over Linguistic Features for Fake News Detection. *IEEE Transactions on Computational Social Systems*, 8(4): 881–893.

Wang, H.; Dou, Y.; Chen, C.; Sun, L.; Yu, P. S.; and Shu, K. 2023a. Attacking Fake News Detectors via Manipulating News Social Engagement. arXiv:2302.07363.

Wang, L.; Xu, X.; Zhang, L.; Lu, J.; Xu, Y.; Xu, H.; Tang, M.; and Zhang, C. 2024. MMIDR: Teaching Large Language Model to Interpret Multimodal Misinformation via Knowledge Distillation. arXiv:2403.14171.

Wang, W.; Haddow, B.; Birch, A.; and Peng, W. 2023b. Assessing the Reliability of Large Language Model Knowledge. arXiv:2310.09820.

WU, Y.; Zhang, Z.; WANG, F.; Luo, Y.; Xiong, H.; and Tang, N. 2024. Detecting Out-of-Context Misinformation via Multi-Agent and Multi-Grained Retrieval.

Ethical Consideration

While this paper primarily proposes a conceptual multi-agent framework for addressing the misinformation lifecycle and outlines high-level implementation aspects without presenting experimental data, it is crucial to acknowledge the inherent ethical considerations. The future development and deployment of such a system demand careful scrutiny of each specialized agent. For instance, the Indexer Agent must address potential biases in the selection and maintenance of "trusted" repositories. The Classifier Agent's categorization of misinformation types requires safeguards against mislabeling and perpetuating harmful stereotypes. The Extractor Agent's evidence retrieval and ranking mechanisms must ensure fairness and avoid amplifying certain viewpoints disproportionately. Furthermore, the Corrector Agent's generation of corrections carries the responsibility of accuracy and neutrality, while the Verification Agent's processes for validating outputs and source credibility must be transparent and robust against manipulation. Future work will need to rigorously evaluate and mitigate these and other ethical challenges to ensure responsible application of this framework. Below are the ethical checklist responses for this paper.

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **NA, No data is used in this paper**
- (e) Did you describe the limitations of your work? **Yes, there is a section on potential limitation**
- (f) Did you discuss any potential negative societal impacts of your work? **Yes, though this is little beyond the scope, it is mentioned that the agents and data-sources needs to be well vetted and the labeling would be done by experts and authorized body to avoid any societal bias and adverse impact**
- (g) Did you discuss any potential misuse of your work? **No, this is the beyond the scope of the framework proposed**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **NA, No data is used**

- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes, I have read them**
- ### 2. Additionally, if your study involves hypotheses testing...
- (a) Did you clearly state the assumptions underlying all theoretical results? **NA, there is no experimentation in this paper**
 - (b) Have you provided justifications for all theoretical results? **NA**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**
 - (e) Did you address potential biases or limitations in your theoretical framework? **Yes, it has been discussed in one section**
 - (f) Have you related your theoretical results to the existing literature in social science? **NA, there is no relevance of social science**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA**
- ### 3. Additionally, if you are including theoretical proofs...
- (a) Did you state the full set of assumptions of all theoretical results? **NA**
 - (b) Did you include complete proofs of all theoretical results? **NA**
- ### 4. Additionally, if you ran machine learning experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **NA, no experimentation in this proposal just the framework**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **NA**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **NA**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **NA**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **NA**
 - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? **NA**
- ### 5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
- (a) If your work uses existing assets, did you cite the creators? **NA**
 - (b) Did you mention the license of the assets? **NA**
 - (c) Did you include any new assets in the supplemental material or as a URL? **NA**

- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating?
NA
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? NA
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see ?)? NA
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see ?)? NA
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**
- (a) Did you include the full text of instructions given to participants and screenshots? NA
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
 - (d) Did you discuss how data is stored, shared, and de-identified? NA