

1. Introduction

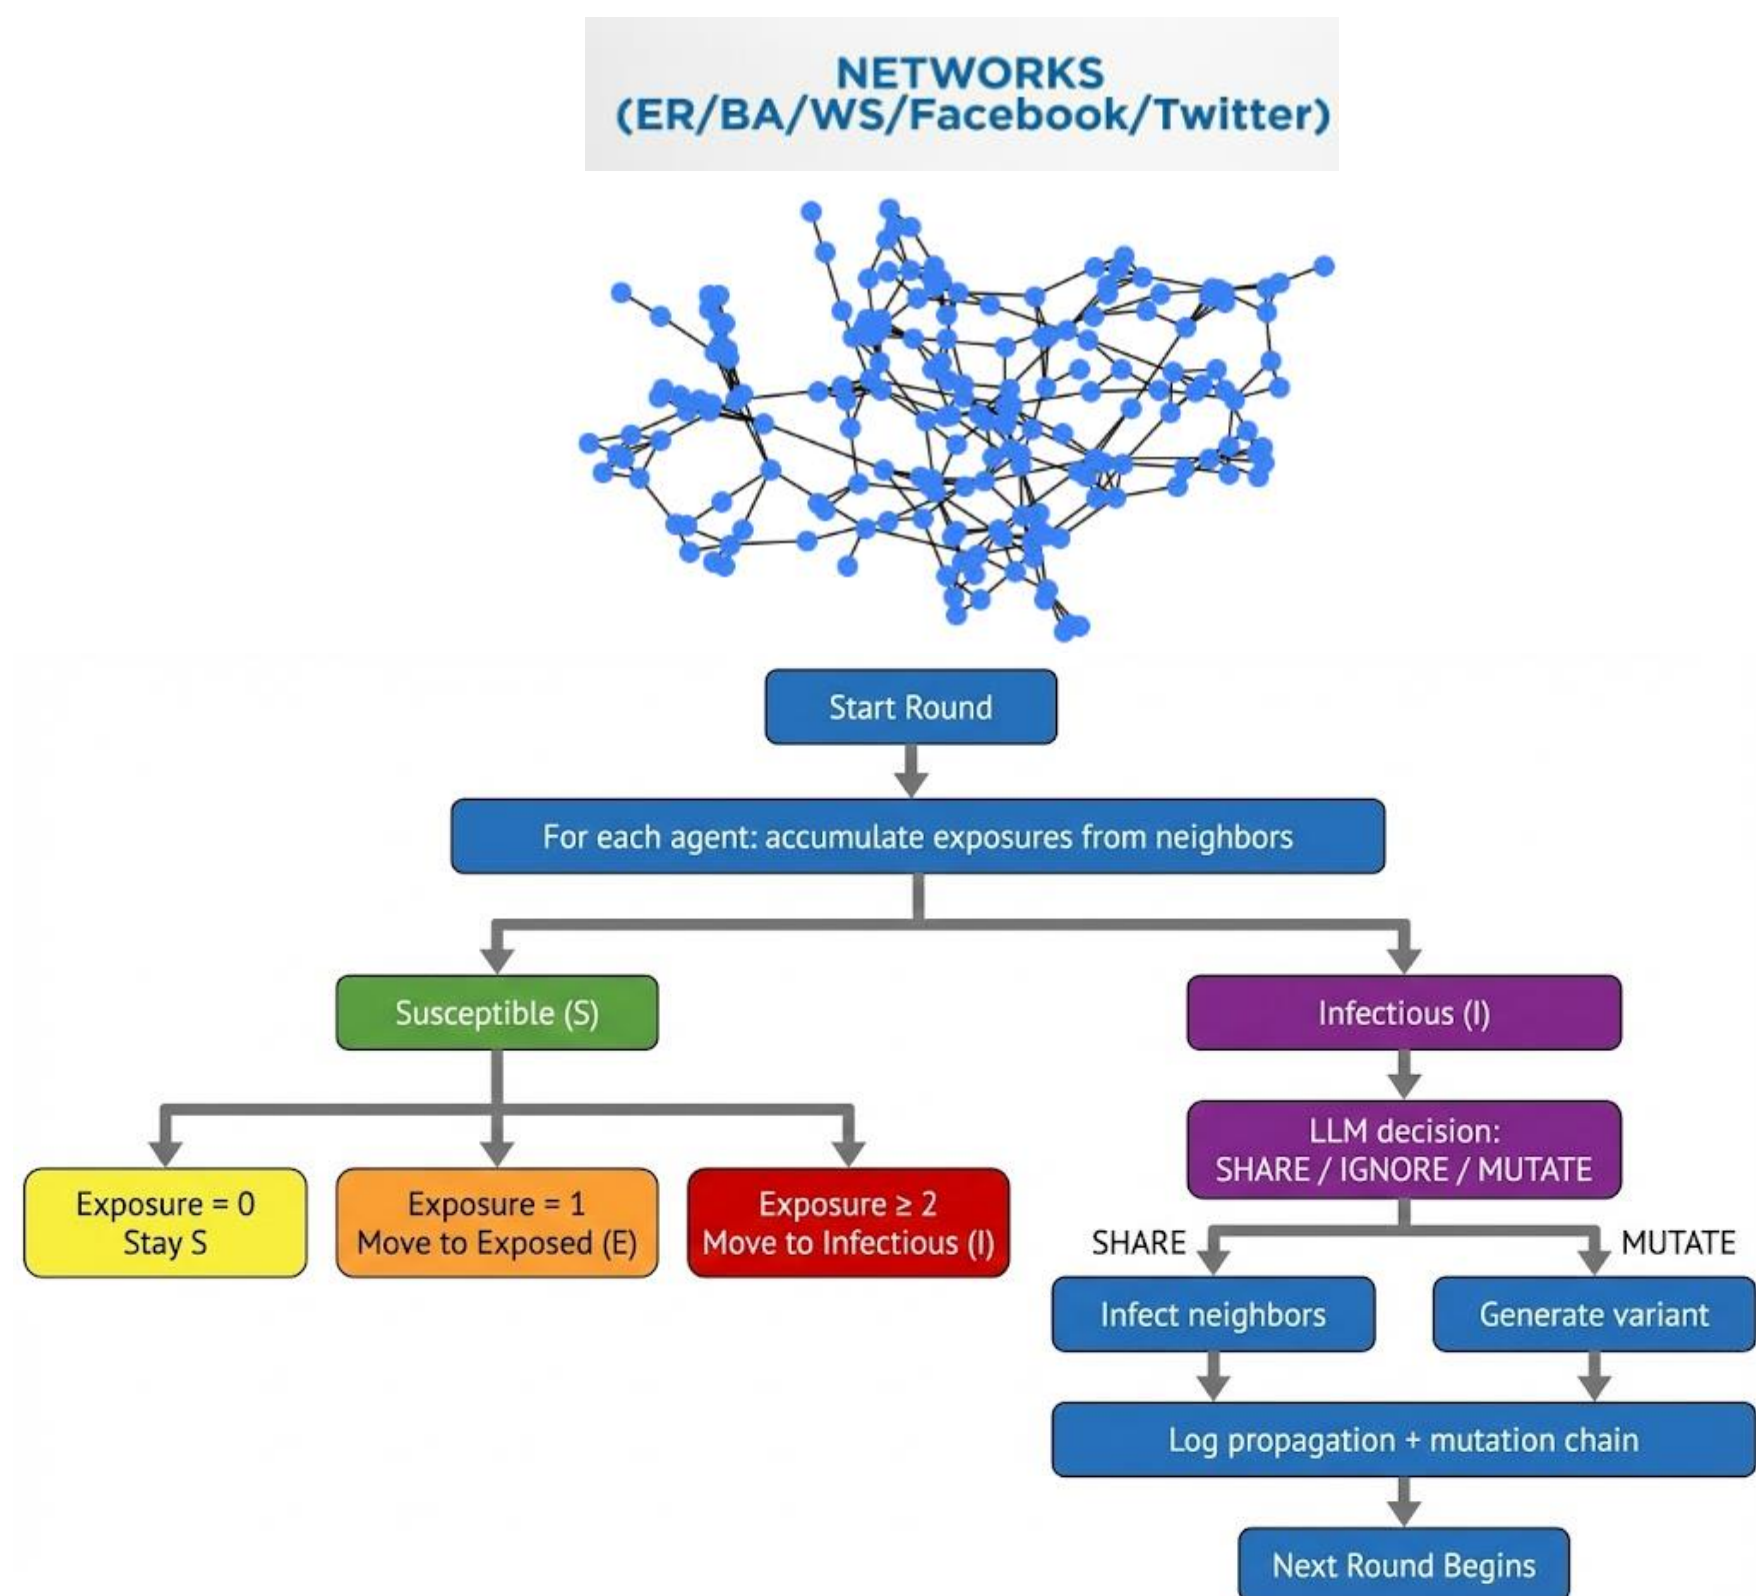
Abstract

Rumors and misinformation influence collective behavior, yet traditional SIR-style diffusion models reduce humans to simple probabilistic rules. Large Language Models enable agent-based simulations where individuals exhibit memory, emotion, persona biases, and spontaneous paraphrasing. We investigate: How do network structure, persona composition, and mutation mechanisms shape rumor spread and evolution?

Highlights

- ❖ Combined graph-theoretic modeling with LLM-based agents for realistic rumor propagation.
- ❖ Studied both spread and mutation of rumors across synthetic and real networks.
- ❖ Quantified semantic drift, emotional framing, and evasion of fake-news detectors across generations.

3. Approach



- ❖ Unlike traditional SIR models where nodes are passive relays, our agents are powered by Llama-3.1-8B.
- ❖ Agents are assigned behavioral profiles that dictate their rewriting logic: Neutral (45%): Objective paraphrasing. Skeptical (30%): Adds doubt or requests evidence. Conformist (15%): Aligns with the perceived majority. Sensationalist (10%): Amplifies emotional and dramatic framing. Fact-Checkers: 5% of nodes are designated as "Correctors"
- ❖ Unified system measuring both how far a rumor travels and how much its meaning changes as it spreads.
- ❖ Introduce metrics to quantify information distortion
- ❖ Semantic Drift: Cosine similarity between the original and evolved embeddings (SBERT all-MiniLM-L6-v2).
- ❖ Cognitive Load: Text perplexity and readability scores (Flesch-Kincaid)


2. Problem description

Real misinformation evolves as people reinterpret or sensationalize claims. We seek to model both diffusion (who gets infected) and evolution (how the rumor changes).


Research Questions:

- ❖ How do personas (skeptical, neutral, conformist, sensationalist) affect spreading?
- ❖ How do different network topologies (ER, BA, WS, Facebook, Twitter, SBM) accelerate or inhibit cascades?
- ❖ How do mutations alter semantic meaning, readability, sentiment, and detectability?


4. Results



Network structure governs timing, not outcome. While hubs accelerate saturation (<3 rounds), high Modularity only delays the global outbreak. Once a rumor escapes the 'echo chamber,' the final reach is unavoidable.



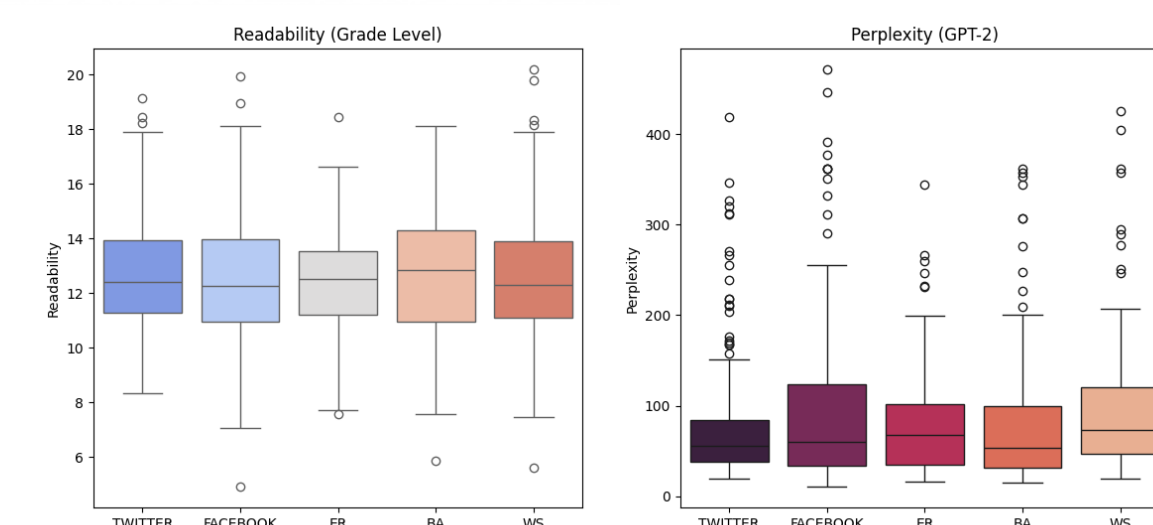
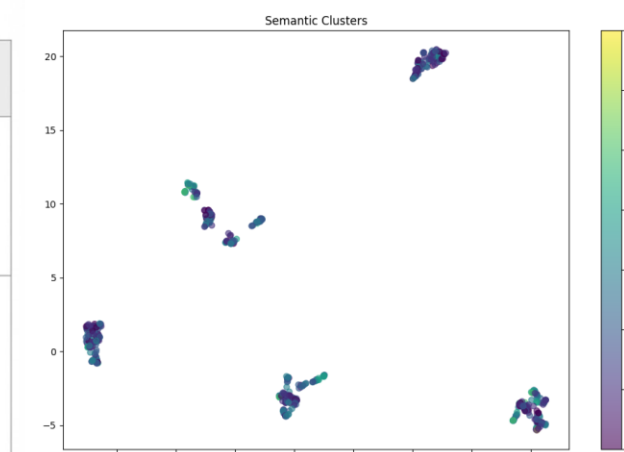
Misinformation is bottlenecked by agents checking for credibility, true content passes these filters, achieving ~10% higher saturation across all topologies.



In the absence of community clusters to reinforce the original context, semantic similarity drops ~2x faster than in clustered social networks, where local redundancy acts as a 'Semantic Anchor' for the truth.

Qualitative Evolution: How the Rumor Mutated

Original Rumor	Evolved Variant (Generation 3)	Persona Driver
Staff at Gold Coast Hospital danced for a man in quarantine.	"The ENTIRE staff... ERUPTED into a choreographed dance party... the drama was simply EPIC!"	Sensationalist (Amplifies Drama)
Chinese doctors confirmed African people are resistant to COVID.	"Reports suggest that some African populations may have a genetic advantage... but this requires further evidence."	Skeptic (Sanitizes & Adds Doubt)



Rumors do not degrade into random noise; they follow a directed evolutionary path, converging into distinct 'generational dialects' at each stage of the spread. The rumor evolves to become linguistically 'natural' (low perplexity) while maintaining a high grade for maximum credibility.