

Quantum Perspectives on Persuasive Language in AI-Generated News: A QNLP-Based Analysis

Jung-Hua Liu

Department of Communication, National Chung Cheng University / No. 168, Sec. 1,
University Road, Minxiong Township, Chiayi County 621301, Taiwan
emileliu@ccu.edu.tw

Abstract

This study applies quantum natural language processing (QNLP) to 298 Chinese AI-generated YouTube news articles. Using IBM Qiskit, this study reveals multi-framing narratives with high frame competition but low conflict. Headlines employ emotion, content stays neutral or positive, showing strategic ambiguity. QNLP metrics highlight persuasive tactics and implications for communication theory and AI ethics.

Keywords: Quantum NLP, persuasive language, framing, agenda-setting, strategic ambiguity

1 Introduction

Generative AI now produces news articles, raising questions about authenticity, framing, and persuasion (Paviour, 2025). Building on Lippmann's (2017) idea that news constructs a "pseudo-environment," Yin & Liu (2025) and Reubold, and Campbell (2023) note that AI-driven journalism transfers gatekeeping from editors to algorithms. Classic theories of agenda-setting and framing remain relevant: while human editors once selected topics and angles, AI systems may inherit training-data biases or create new emphases (Mehrab et al., 2021; de-Lima-Santos & Jamil, 2024; Kuku et al., 2025; Singh & Ngu, 2025). Studies show AI news can differ in style and tone from human reporting but is not necessarily more biased (Nah et al., 2024; Sui, 2025). Recent studies on AI-generated discourse reveal that persuasion in digital communication extends far beyond surface-level sentiment or credibility measures. Goldstein et al. (2024) demonstrate that GPT-3 can generate propaganda nearly as persuasive as authentic state-backed content, particularly when human curators refine or select output. This finding underscores that persuasive efficacy in AI-generated text emerges not solely from factual accuracy but from rhetorical coherence, emotional framing, and contextual adaptability. Similarly, Pazzaglia et al. (2025) show that fine-tuned large language models reproduce polarized ideological rhetoric with high credibility and emotional

resonance. Their model's outputs were rated as both "provocative" and "human-like," suggesting that persuasive force arises from the capacity of language models to reproduce rhetorical alignment - a blending of ideological tone, emotional activation, and discursive context. Meanwhile, Al Giffari and Dermawan (2025) reveal through comparative rhetorical analysis that AI-generated religious messages, though formally coherent and citation-driven, lack the ethos, pathos, and kairotic timing that human preachers use to achieve moral and emotional persuasion.

These works indicate that persuasion in AI discourse depends not only on propositional content but on the quantum-like coexistence of multiple interpretive frames - logical, emotional, and ethical - that audiences navigate dynamically. Quantum Natural Language Processing (QNLP) provides a theoretical and computational framework for representing this multidimensional interplay. By encoding textual meaning as quantum states, QNLP models semantic superposition (simultaneous coexistence of conflicting frames), entanglement (interdependence among linguistic and contextual cues), and measurement collapse (resolution of ambiguity through interpretation). In persuasion analysis, these quantum phenomena map onto how readers oscillate between alternative framings or emotional cues before forming conviction—analogue to the probabilistic collapse of a quantum system upon observation. This approach captures what traditional NLP misses: that persuasive communication operates through contextual interference patterns among emotional tone, narrative perspective, and cultural resonance.

QNLP thus formalizes persuasion as an emergent property of narrative context rather than as a unidirectional rhetorical act. It illuminates how AI-generated news or propaganda can appear simultaneously neutral, credible, and manipulative - precisely because its semantic space allows multiple persuasive potentials to coexist until interpretively resolved by the audience. Through this lens, QNLP bridges computational linguistics and communication theory, offering a post-classical model for analyzing how machine-produced narratives shape belief, trust, and ideological alignment in the quantum field of discourse. Findings highlight strategic ambiguity,

emotional framing, and expanded agenda breadth as persuasive features. QNLP thus bridges communication theory and quantum semantics, offering new tools for detecting subtle persuasive strategies in AI-generated content.

2 Literature Review

2.1 Persuasion through Agenda-Setting, Framing, and Rhetorical Strategies

Persuasion in news discourse has long been theorized through the intertwined mechanisms of agenda-setting, framing, and rhetorical strategy. These perspectives, while often treated separately, all explain how media shape public attitudes not by direct argumentation but by structuring attention, interpretation, and affective response—the key ingredients of persuasion.

Agenda-setting theory (McCombs & Shaw, 1972) shows that the persuasive force of news lies in its power to prioritize certain topics over others, implicitly signaling their importance. At the first level, issue salience determines what the public thinks about; at the second level, attribute salience determines how they think about it. For example, emphasizing unemployment statistics rather than individual hardships frames the issue as technical rather than moral, guiding public concern and policy preferences. This process is persuasive because it conditions cognitive accessibility: repeated exposure elevates certain issues in collective awareness, creating perceived consensus and urgency.

Framing theory deepens this account by showing that persuasion occurs through selection and emphasis. Entman (1993) defined framing as the act of selecting aspects of perceived reality to make them more salient, thereby promoting specific problem definitions, causal interpretations, moral evaluations, and policy recommendations. Frames thus operate as interpretive templates that steer reasoning. A protest described as a “law-and-order problem” activates threat and control schemas, whereas the same event framed as a “civil-rights struggle” evokes empathy and justice. In both cases, framing does not merely present facts—it organizes meaning in ways that predispose audiences toward particular attitudes or actions.

Rhetorical strategies complete the persuasive triad by illuminating how linguistic and stylistic choices translate cognitive framing into affective engagement. Classical rhetoric’s *ethos*, *pathos*, and *logos* correspond to credibility, emotion, and logic—the dimensions that sustain belief formation.

Even under norms of journalistic objectivity, subtle rhetorical cues such as evaluative adjectives, quotation patterns, or metaphoric phrasing convey stance and invite alignment. Ceccarelli’s (1998) concept of strategic ambiguity further explains how persuasion can arise from texts that support multiple plausible interpretations: ambiguity minimizes resistance by allowing diverse audiences to read agreement into the same message. Thus, persuasion in journalism is often implicit, operating through agenda prominence (what to think about), framing (how to think about it), and rhetorical form (how to feel about it). These mechanisms collectively blur the boundary between informing and influencing, creating an ecology of persuasion that relies on selection, emphasis, and affect rather than overt argumentation.

2.2 Persuasion in the Age of AI and NLP

The rise of artificial intelligence (AI) in journalism, often referred to as automated or robot journalism, has intensified scholarly attention to persuasion’s algorithmic dimensions. During the 2010s, outlets such as Reuters, the Associated Press, and The New York Times adopted rule-based generators for financial reports and sports summaries (Carlson, 2018; Diakopoulo, 2019; The Newsreel Project Consortium, 2021). By the 2020s, large language models (LLMs) like OpenAI’s GPT series enabled generative systems to produce multi-paragraph narratives that mimic human style and rhetorical nuance.

Recent studies highlight both opportunities and ethical challenges. A systematic review by Ioscote et al. (2024) notes that automation improves efficiency but introduces opacity and potential bias. Graefe (2016) found that algorithmic news was perceived as competent but emotionally flat, while Nah et al. (2024) observed that AI-generated stories differ in tone and coherence yet are not necessarily more biased. Nonetheless, AI’s capacity to synthesize persuasive patterns from vast corpora gives it unprecedented influence over public cognition. Goldstein et al. (2024) showed that GPT-3-generated propaganda can be nearly as persuasive as human-written material, especially when curated by humans. Pazzaglia et al. (2025) found that fine-tuned LLMs replicate polarized discourse with rhetorical realism, while Al Giffari and Dermawan (2025) demonstrated that AI reproduces logical appeals but lacks the adaptive *ethos* and emotional depth of human persuasion. These findings converge on one point: AI’s

persuasive power lies in its ability to simulate the *agenda-setting and framing patterns* that shape interpretive hierarchies in human journalism.

Even absent malicious intent, AI systems reproduce persuasive conventions—issue prioritization, emotional tone, narrative balance, and ambiguity—because these features are embedded in their training data. Conventional NLP tools such as sentiment analysis or topic modeling can capture tone and frequency but cannot fully represent how frames interact or compete within a narrative. Similarly, propaganda-detection systems focus on lexical signals but overlook the contextual superpositions that make messages persuasive across ideological lines.

QNLP provides a post-classical approach to this challenge. By encoding text as quantum states, QNLP models superposition (simultaneous activation of multiple frames), entanglement (interdependence among topics, emotions, and rhetorical cues), and measurement collapse (resolution of interpretive ambiguity during audience reception). These quantum concepts parallel how persuasion unfolds in narrative contexts: audiences oscillate between competing frames and affective interpretations before settling on belief or skepticism. QNLP thus allows researchers to formalize and visualize the non-linear, context-dependent nature of persuasion—how agenda-setting, framing, and rhetoric operate together to construct probabilistic meaning fields rather than fixed messages.

Classical theories reveal that persuasion in journalism emerges from the coordination of attention (agenda-setting), interpretation (framing), and affect (rhetoric). In the AI era, these mechanisms are not only replicated but amplified by generative systems capable of producing multi-frame, strategically ambiguous narratives at scale. QNLP offers a novel alternative. By encoding texts as quantum states, QNLP enables analysis of overlapping meanings, frame superpositions, and narrative entanglements. This study is among the first to employ QNLP to examine persuasive dynamics in AI-generated news, particularly focusing on frame competition, ambiguity, and agenda breadth (Wazni et al., 2024; Widdows et al., 2024). Integrating QNLP into this analytical framework offers a powerful means to decode the entangled semantics and contextual fluidity of persuasion in AI-generated news.

3 Methodology

3.1 Dataset and Corpus Preparation

This study generated and analyzed a dataset of 298 GPT-4o-generated Chinese-language news articles obtained from a YouTube channel that produces automated news videos. These videos drew content from Yahoo! News across domains such as politics, economy, technology, and society. Each news item contained three textual components:

- **News Title:** Averaging 16 Chinese characters, titles conveyed the story's core point or a teaser. They were designed to attract attention, often using emotion or framing (e.g., "*Tech CEO Promises Reform Amid Crisis*").
- **Video Dialogue (Transcript):** Averaging 334 characters, dialogues resembled talk-show or multi-speaker formats, simulating anchors and guests. This style incorporated multiple perspectives, quotes, and facts.
- **Video Description:** Averaging 256 characters, descriptions summarized key points and context, functioning as concise press-release style overviews.

Together, these three layers provided a multi-tiered discourse structure: headlines framed events with emotional hooks, dialogues expanded perspectives through conversation, and descriptions offered neutral summaries. This layering enabled analysis of persuasive strategies at different textual levels.

The dataset covered diverse topics, ensuring generalizable findings beyond a single domain. While modest in size (n=298), the corpus allowed meaningful quantitative analysis while remaining computationally manageable.

To establish a baseline and strengthen empirical grounding, the QNLP pipeline was applied to a comparative dataset of professionally written news from Taiwan's Central News Agency (CNA), the nation's official wire service. The dataset comprised 20 paired samples of news titles and full articles, all published in 2020, thereby ensuring that the material predated the widespread adoption of AI-assisted or AI-generated writing. This corpus served as a human-authored benchmark for assessing whether the distinctive characteristics observed in AI-generated texts - such as high frame competition, low conflict, and mild positivity - are unique to algorithmic generation

or instead reflect broader conventions of traditional journalistic discourse.

Text Preprocessing: Since Chinese lacks spaces, word segmentation was essential. The Jieba tool was used to split text into lexemes (e.g., “人工智慧” as *artificial intelligence* rather than “人工” + “智慧”). After segmentation, standard cleaning included normalizing full-width to half-width characters, ensuring UTF-8 encoding, and removing non-textual artifacts. Stopword removal was not applied, as function words carry meaning important for QNLP. All analysis was conducted in Chinese. Each component (title, dialogue, description) was analyzed separately and comparatively to reveal differences in tone, framing, and entropy.

3.2 Quantum NLP Encoding with Qiskit

The Quantum Natural Language Processing (QNLP) pipeline was implemented using IBM’s Qiskit. Following the DisCoCat model (Coecke et al., 2010; Meichanetzidis et al., 2020), text was encoded as quantum states to represent semantic and narrative features.

The Jieba library performs segmentation and part-of-speech tagging. Each segmented and POS-tagged Chinese word is encoded as a quantum state $|\psi_{\text{word}}\rangle$, where its grammatical role determines the number of qubits used and how they interact within the circuit. Mapped tags follow the DisCoCat (Categorical Compositional Distributional) model types (see Table 1). Each part-of-speech category is assigned a **pregroup type** which is mapped to a **vector-space representation** $T_y(-)$ under the strong monoidal functor $F: \text{Pregroup} \rightarrow \text{FVect}$. The notation $T_y(n)$ denotes the vector space corresponding to the noun type under F . The tensor product symbol (\otimes) indicates the compositional combination of vector spaces (or linear maps) to represent joint meaning and grammatical interaction in the DisCoCat framework:

POS	Category	DisCoCat Type
n, nr, nt	Noun (N)	$T_y(n)$
v, vn	Verb (V)	$T_y(n)^r \otimes T_y(s) \otimes T_y(n)^l$
a	Adjective (A)	$T_y(n) \otimes T_y(n)^l$
d	Adverb (D)	$T_y(s) \otimes T_y(s)^l$
p	Preposition (P)	$T_y(n)^r \otimes T_y(n) \otimes T_y(n)^l$

Table 1: Pregroup Type \rightarrow Vector-Space Mapping.

Each qubit acts as a semantic container that can represent multiple potential meanings simultaneously just as a word such as “改革” (reform) may convey both positive and critical implications depending on context. Unlike classical bits that exist only as 0 or 1, a qubit can occupy a superposition

$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, representing a weighted combination of interpretive possibilities. In the QNLP model, each part of speech (POS) corresponds to a grammatical type that specifies how its meaning composes with others:

Noun (N): a single-qubit subsystem $T_y(n)$ representing an entity.

Verb (V): a composite subsystem $T_y(n)^r \otimes T_y(s) \otimes T_y(n)^l$ that links two noun qubits (subject and object) through entanglement.

Adjective (A): a two-qubit structure $T_y(n) \otimes T_y(n)^l$ that modifies a noun.

Adverb (D): a two-qubit structure $T_y(s) \otimes T_y(s)^l$ that modifies a verb or clause.

Preposition (P): a three-qubit subsystem $T_y(n)^r \otimes T_y(n) \otimes T_y(n)^l$ that introduces relational meaning.

There is an example as the below one,

Sentence: 「麥當勞性侵案後改革 董事長發聲承諾改善」

(*Màidāngláo xìngqān àn hòu gǎigé dǒngshìzhǎng fāshēng chéngnuò gǎishàn*; *After the McDonald’s sexual assault case, the chairman spoke out and promised reform.*)

Segmentation and POS tagging:

Output:

[(“麥當勞 (Màidāngláo, McDonald’s)”, ‘nt’), (“性侵 (xìngqān, sexual assault)”, ‘n’), (“案 (àn, case)”, ‘n’), (“後 (hòu, after)”, ‘f’), (“改革 (gǎigé, reform)”, ‘v’), (“董事長 (dǒngshìzhǎng, chairman)”, ‘n’), (“發聲 (fāshēng, to speak out)”, ‘v’), (“承諾 (chéngnuò, to promise)”, ‘v’), (“改善 (gǎishàn, to improve)”, ‘v’)]

Mapped to grammatical categories following DisCoCat: [N, N, N, F, V, N, V, V, V].

Here:

- **N (noun)** = organization / entity
- **V (verb)** = action or statement
- **F (function)** = adverbial or time marker (“後”, after)

The base model uses eight qubits representing major Chinese grammatical categories (noun,

verb, adjective, adverb, preposition, pronoun, conjunction, other). Additional qubits (up to four) are allocated proportionally to the number of unique part-of-speech tags and compositional transitions, ensuring that texts with richer syntactic variation yield more entangled circuits. The algorithm means the more **unique POS tags** (diversity of grammar) a sentence contains, And the more **category transitions** (e.g., $N \rightarrow V \rightarrow N \rightarrow F \rightarrow V$) occur, \rightarrow The higher the **compositional complexity**, and thus, more qubits are added.

These qubits don't represent specific parts of speech — they capture **semantic entanglement patterns** such as:

- **Noun–Verb entanglement:** subject–predicate dependencies.
- **Adjective–Noun entanglement:** modification dependencies.
- **Temporal–Action coupling:** time or cause-effect encoding.

For example, in the above case, “麥當勞性侵害案後改革 董事長發聲承諾改善”, there are 9 tokens, 4 major POS categories (N, V, F, A), and multiple inter-category transitions:

$N \rightarrow N \rightarrow N \rightarrow F \rightarrow V \rightarrow N \rightarrow V \rightarrow V \rightarrow V$ yielding a compositional complexity value high enough to allocate 4 extra qubits (Table 2).

Qubit	Linguistic Role	Example Representation
q ₀	Noun (Subject)	麥當勞 (McDonald's)
q ₁	Verb (Action)	改革 (reform)
q ₂	Function / Modifier	後 (after)
q ₃	Complement Noun	董事長 (chairman)
q ₄	Adjective / Evaluation	良好 (good)
q ₅	Adverb / Tone	積極 (actively)
q ₆	Contextual Frame	政治 / 經濟 (political/economic)
q ₇	Rhetorical Mode	Hopeful / Critical tone

Table 2: Qubit Type

This 8-qubit configuration allows this model to:

- Encode **frame superposition** (multiple meanings or framings coexisting).
- Maintain **semantic entanglement** (how grammatical roles affect each other).

- Simulate **interpretive collapse** (when a reader resolves ambiguity).

4 base qubits = structural grammar,
4 extra qubits = higher-order semantic entanglement. Together, they form a full 8-qubit quantum linguistic state:

$$|\psi\rangle = \sum_{i=0}^{2^8-1} \alpha_i |i\rangle$$

where each amplitude α_i corresponds to a possible interpretive configuration of the sentence. Each of the 256 possible configurations (from $|00000000\rangle$ to $|11111111\rangle$) represents a distinct combination of meanings. Each amplitude α_i captures the weight or probability of that interpretation. When measured (interpreted by a reader), the sentence collapses into one dominant interpretation. If many amplitudes are large, the sentence is ambiguous with multiple frames; if one dominates, the meaning is singular.

In the Qiskit implementation:

- **Hadamard (H)** gates initialize the emotional subsystem into superposition.
- **Rotation (R_Y)** gates encode each quantum weight as a rotation angle.
- **Controlled-NOT (CX)** and **Controlled-RZ (CRZ)** gates introduce entanglement when both positive and negative cues occur, simulating **emotional interference** between coexisting sentiments.

this design allows the circuit to capture complex emotional polarity interactions, e.g., optimism and anxiety expressed simultaneously within reform narratives.

Each sentence is converted to a quantum circuit through three main steps

1. **Initialization:** All qubits begin in superposition states via Hadamard gates: $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, representing interpretive openness.
2. **Category-Specific Rotations:** Each grammatical category applies rotation gates proportional to its frequency and semantic role. Rotation about the Y-axis, $R_Y(\theta)$, encodes meaning amplitude; phase rotations (R_Z) introduce semantic distinctions.

3. **Entanglement:** CNOT (CX) and controlled rotation (CRY) gates encode syntactic dependencies such as noun–verb or adjective–noun relationships.

Here, nouns form base qubits, verbs are modeled as multi-qubit subsystems, and grammatical dependencies such as noun–verb or adjective–noun pairs are represented through entanglement.

3.3 Quantum Metric Definitions

Defined as the normalized von Neumann entropy of the article’s density matrix (Widdows, et al., 2024; Agostino et al., 2025):

$$C = \frac{-\text{Tr}(\rho \log_2 \rho)}{\log_2 N}$$

where:

- ρ is the *density matrix*, which represents all possible meanings or frames encoded in the text’s quantum state. It is constructed as $\rho = |\psi\rangle\langle\psi|$, the outer product of the statevector with itself.
- $\text{Tr}(\rho \log_2 \rho)$ means taking the *trace* (sum of diagonal elements) of the matrix after applying a logarithm. This operation computes the weighted average uncertainty of the entire meaning distribution.
- $-\text{Tr}(\rho \log_2 \rho)$ gives the *von Neumann entropy*, a measure of how mixed or diverse the meanings are.
- **Dividing by $\log_2 N$** (where N is the number of possible interpretive frames) normalizes the result between 0 and 1.

If $C \approx 1$, the text contains multiple equally active frames (e.g., political, moral, and economic frames appearing together). If $C \approx 0$, one frame dominates and the article has a single clear angle. Thus, C quantifies how much interpretive “competition” exists in the text’s meaning structure.

Von Neumann Entropy

$$S(\rho) = -\text{Tr}(\rho \log_2 \rho)$$

where:

- $S(\rho)$ is the *von Neumann entropy*, quantifying semantic uncertainty or agenda diversity.

- ρ is the density matrix of the encoded quantum-linguistic state.
- Tr denotes the trace, and \log_2 computes information in bits.

Higher entropy values (approaching 0.9) imply broad topical or interpretive diversity; lower values indicate focused discourse.

3.4 Quantum-weighted Sentiment

Emotional tone analysis was implemented through a **heuristic quantum-weighted lexicon**, where each emotional token is assigned a polarity weight representing its affective intensity in the range 0.65–0.95. Because existing Chinese sentiment benchmarks do not support the QNLP encoding framework, this study constructed a self-calibrated emotional lexicon based on high-frequency evaluative terms observed in the dataset.

The complete emotion arrays used for analysis are listed below.

These weights serve as quantum amplitudes reflecting how strongly each emotional concept contributes to the sentence’s overall affective state before normalization.

After segmentation and POS tagging, each emotional word is matched with its corresponding weight w_i . Sentence-level emotional intensity is calculated as:

$$I_{\text{emotion}} = \sum_i (w_i \times \text{count}_i) / \text{total_words}$$

The resulting intensity is then mapped to a rotation parameter for quantum circuit encoding:

$$\theta_i = \pi \times I_{\text{emotion}}.$$

Thus, a word with weight $w_i = 0.78$ produces a rotation $R_Y(0.78\pi)$, generating the corresponding emotional amplitude in the quantum state.

Syntactic patterns further refine the emotional amplitude:

- Active-voice markers (“主動 zhǔdòng – active / initiative”, “積極 jījī – positive / energetic”, “推動 tuīdòng – to promote / to drive forward”) add up to +0.10 to strengthen positive orientation.
- Future markers (“將 jiāng – will / shall”, “會 huì – will / be likely to”, “計劃 jìhuà – plan / project”) add up to +0.05 to indicate optimism and anticipation.

Each sentence's affective profile is encoded as a normalized superposition:

$$|\psi_{emotion}\rangle = \alpha |positive\rangle + \beta |neutral\rangle + \gamma |negative\rangle,$$

where $|\alpha|^2 + |\beta|^2 + |\gamma|^2 = 1$.

4 Result

4.1 Overall Patterns: Multiple Framings and Frame Dynamics

The analysis confirmed that AI-generated news articles frequently exhibit “multiple framings” within their narratives. The multi-framing intensity averaged 0.7716 (on a 0–1 scale), suggesting that a single article typically encodes several possible interpretations simultaneously rather than offering a univocal story. This means that readers could plausibly reach different conclusions about events depending on which parts of the narrative they emphasize. Such a finding offers empirical support to the concept of quantum semantics in media: meanings remain in superposition until “collapsed” by reader interpretation. Rather than committing to one framing, AI-generated texts often include both optimism and skepticism, or conflict and harmony, side by side. This pattern challenges traditional expectations of objectivity in journalism and resonates with postmodern views of news as narrative construction, amplified by AI's probabilistic generation methods.

The comparative analysis between AI-generated content and CNA journalism reveals distinct differences in narrative structure and informational richness. In terms of Frame Competition, AI exhibits *perfect competition* (1.0000), meaning that all semantic frames coexist equally without dominance, reflecting a balanced multi-perspective discourse. In contrast, CNA demonstrates a *high but not perfect competition* (0.9173–0.9985), suggesting a slight frame hierarchy that produces more structured and coherent narratives. Examining von Neumann Entropy, AI maintains a consistent entropy of 4.0000, indicating uniform information density and even distribution of meanings. CNA, however, shows *variable entropy* ranging from 3.4378 to 7.3508, which is approximately 84% higher than AI, evidencing greater informational diversity and narrative complexity. Overall, CNA content is significantly more information-dense, while AI maintains ideal frame equality and supports multiple simultaneous interpretations. Both sources sustain a neutral tone,

but CNA achieves neutrality through editorial consistency, whereas AI achieves it through semantic averaging. These findings suggest that AI-generated content successfully models the “multiple framings” phenomenon characteristic of pluralistic discourse, yet this comes at the cost of reduced information density compared to the more hierarchically structured and detail-rich style of professional journalism.

Frame analysis revealed an additional pattern: very high frame competition (average 0.8891) but low frame conflict (average 0.1640). AI news tends to present numerous frames simultaneously but arranges them to minimize contradiction. For instance, a controversial policy article might include both “public safety” and “personal freedom” frames without resolving which is correct. Each frame is presented discretely, often by different speakers, allowing peaceful coexistence. Unlike traditional journalism, where competing frames often clash, AI-generated narratives appear to place frames side by side. This reflects a distinctive “high competition, low conflict” framing style that broadens interpretive possibilities without forcing resolution.

From a persuasion standpoint, this polyvalence can be read as strategic ambiguity. By leaving interpretation open, AI news accommodates varied audiences, each of whom may find their own perspective validated.

4.2 Emotional Tone and Sentiment Use

Emotional tone analysis showed that AI-generated news maintains a largely neutral to slightly positive register, with negative sentiment being rare. The mean positive sentiment intensity was ~0.2065, while only about 23.4% of articles carried significant negative language. Neutrality dominated across the corpus, suggesting a style that favors factual exposition peppered with subtle positivity.

Breaking down by section revealed important differences. Titles carried the strongest emotional charge (average 0.2760), often employing dramatic or evaluative words such as “重大突破” (major breakthrough) or “嚴重警告” (stern warning). About 37% of titles included stronger sentiment than the body, aligning with journalistic practices of crafting attention-grabbing headlines.

Dialogues, which made up the body text, were the most neutral (average sentiment 0.1566). Emotive expressions were frequently balanced by counterpoints in simulated multi-speaker exchanges. This dynamic reduced variance and created an impression of neutrality, reinforcing credibility through balanced voices.

Descriptions were mostly factual, resembling wire-service summaries. When sentiment appeared, it leaned positive, often framing problems alongside hopeful solutions. For example, disaster coverage frequently pivoted to recovery measures, mitigating negativity.

Persuasively, this pattern suggests AI news seeks credibility through neutrality while using selective positivity to foster reassurance. Rather than overtly directing audience emotions, it subtly steers interpretation toward optimism.

4.3 Agenda Breadth and Information Density

Another key finding was the broad agenda breadth of AI-generated news. Articles often included wide-ranging contextual information but lacked strong emphasis on priority issues. Von Neumann entropy was highest in descriptions (0.8937), indicating dense, information-rich content. Titles, by contrast, had low entropy, while dialogues fell in between.

Descriptions also scored highest in frame competition (~0.9050). They frequently included multiple angles—political, economic, social, and historical—in a single paragraph. For example, a corporate scandal description referenced ethical implications, financial effects, prior incidents, and investor reactions, leaving the reader to decide which angle mattered most.

This encyclopedic style contrasts with traditional journalism, where editors foreground particular aspects to guide audience focus. AI-generated news instead outsources agenda-setting to readers by presenting numerous perspectives without hierarchy. From a persuasive standpoint, agenda breadth can increase credibility by conveying thoroughness but risks diluting focus. It may also create an “illusion of depth,” where sheer quantity of details fosters trust even if no clear conclusion is provided.

5 Analysis

5.1 Rethinking Media Theory in a Quantum Framework

Classical theories of agenda-setting and framing assume linear effects: media highlight issues to shape public focus and frame them in ways that guide interpretation. AI-generated news disrupts this model. Instead of a singular agenda, AI texts exhibit agenda multiplicity—a wide array of issues included without a clear hierarchy. This suggests a need for an “algorithmic agenda-setting” concept, where priorities emerge from data frequency or algorithmic design rather than editorial judgment. Readers may be told “many things to think about” without guidance on which matter most.

Similarly, framing becomes pluralistic. Rather than privileging one interpretive angle, AI news embeds multiple frames within a single article. This polysemy resonates with postmodern media theory, particularly John Fiske’s work on polysemic texts and Leah Ceccarelli’s notion of strategic ambiguity. The AI is not a rhetor with intent, but the effect mirrors deliberate ambiguity: conflicting audiences can each find validation. A conservative and a liberal might interpret the same AI-generated political story differently, confirming their predispositions. This parallels theories of selective perception and hostile media effect, where ambiguity fosters divergent interpretations.

From a quantum perspective, meaning exists in superposition until “collapsed” by the reader. Different audiences measure the text differently, producing varied interpretations. Unlike traditional journalism, which assumes a preferred reading, AI-generated journalism may lack any singular intended meaning.

5.2 Strategic Ambiguity and Persuasion

Strategic ambiguity emerges as a core persuasive element. By presenting multiple perspectives without adjudication, AI-generated news broadens acceptability. Ceccarelli (1998) noted that ambiguity unites conflicting audiences, and AI articles function similarly. This ambiguity can diffuse polarization by avoiding outright conflict, but it also dilutes clarity. Readers may leave with less certainty about what truly matters.

The persuasive outcome is paradoxical: ambiguity may reduce backlash but also reduce impact. Articles that hedge on every angle may keep audiences engaged without deeply swaying them. In polarized environments, such ambiguity

could stabilize discourse by avoiding provocation, but it may equally risk fostering complacency.

5.3 Emotional Tone and Comfort Bias

Emotional analysis revealed that AI-generated news leans heavily neutral to slightly positive, with negative sentiment rare. This positivity bias, though subtle, may enhance persuasion by creating psychological comfort. Readers often prefer constructive or optimistic narratives, making them more receptive. By emphasizing reforms or solutions, AI-generated articles may foster goodwill toward institutions and authorities.

At the same time, the absence of strong negative framing reduces the risk of outrage-based virality. This could make AI-generated news less prone to fueling polarization but also less effective at holding power accountable. In terms of ethics, neutrality and optimism may seem impartial, yet they introduce a subtle pro-status-quo bias.

5.4 Practical Implications: Media Literacy and Regulation

For media literacy, these findings imply that readers must learn to navigate ambiguity. Rather than identifying a single bias, they should detect multiple frames and question what is absent. Educators may teach critical reading strategies for AI-generated texts: What perspectives are included? Which are omitted? Who benefits from this framing?

For regulation, quantum semantic metrics could aid content moderation. High multi-framing scores might flag overly contradictory or confusing texts, while high competition paired with high conflict could indicate propagandistic extremes. Automated monitoring could complement fact-checking in identifying problematic AI-generated content at scale.

5.5 Comparing Quantum and Classical NLP Approaches

Classical NLP techniques like BERT and LDA are effective at identifying dominant topics and frames. They assign fixed labels or topic proportions to text, capturing surface-level patterns of content. However, they struggle to represent ambiguity (Liu et al., 2023), competing interpretations (Waldon et al., 2025), or the contextual dynamics of persuasion (Saha et al., 2021; Bozdog et al., 2025). When faced with contradictory signals, such as praise and criticism in the same sentence, classical models tend to

average or disambiguate, forcing a singular reading.

QNLP, in contrast, encodes language as quantum states capable of representing multiple meanings simultaneously. Using superposition, QNLP captures coexisting frames; entanglement models dependencies between semantic elements; and measurement simulates reader-driven interpretation, collapsing the state to a context-specific meaning. These features enable QNLP to reflect the uncertainty and multiplicity inherent in persuasive language. Rather than replacing classical methods, QNLP complements them - adding depth in cases of ambiguity, strategic framing, or interpretive variability where classical NLP falls short.

6 Conclusion

Future research should investigate how audiences actually interpret multi-frame AI-generated news. Do readers experience it as balanced and informative, or as vague and non-committal? Controlled experiments could measure which frames readers recall, which interpretations they adopt, and whether strategic ambiguity unites audiences or simply enables selective perception. Comparative analyses with human-written news on identical events would also clarify systematic differences, such as AI's tendency toward broader context or more neutral tone. Expanding to larger, cross-lingual corpora across domains like finance, sports, and health would further test the generality of these patterns and identify whether cultural or stylistic contexts alter persuasive dynamics.

On the technical side, QNLP methods can be refined to enable automatic frame detection, with advances in quantum machine learning and hardware allowing the encoding and analysis of larger, more complex semantic states. Practical applications may include monitoring tools that use metrics such as multi-framing intensity and frame entanglement to flag overly ambiguous or potentially polarizing articles, assisting editors in enhancing clarity. Finally, the study underscores the need for ethical guidelines in AI journalism, ensuring that neutrality does not come at the expense of omitting critical moral or evaluative frames. In sensitive domains like public health, balancing neutrality with clarity is essential for trustworthy communication.

References

Agostino, C. J., Thien, Q. L., Apsel, M., Pak, D., Lesyk, E., & Majumdar, A. (2025). A quantum

- semantic framework for natural language processing. arXiv preprint arXiv:2506.10077.
- Al Giffari, H. A., & Dermawan, A. (2025). AI vs. Human-Led Da'wah: A Comparative Rhetorical Analysis of Islamic Preaching in the Digital Age. *DINIKA: Academic Journal of Islamic Studies*, 10(1), 107-130.
- Bozdag, N. B., Mehri, S., Yang, X., Ha, H., Cheng, Z., Durmus, E., You, J., Ji, H., Tür, G., & Hakkani-Tür, D. (2025, April 23). Must Read: A systematic survey of computational persuasion (arXiv:2505.07775v1). arXiv. <https://arxiv.org/abs/2505.07775v1>
- Campbell, C. H. (2023). *Automated Journalism at the Intersection of Politics and Black Culture: The Battle Against Digital Hegemony*. Lexington Books.
- Carlson, M. (2018). The robotic reporter: Automated journalism and the redefinition of labor, compositional forms, and journalistic authority. In *Journalism in an era of big data* (pp. 108-123). Routledge.
- Ceccarelli, L. (1998). Polysemy: Multiple meanings in rhetorical criticism. *Quarterly Journal of Speech*, 84(4), 395-415.
- Coecke, B., Sadrzadeh, M., & Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. arXiv preprint arXiv:1003.4394.
- de-Lima-Santos, M. F., & Jamil, S. (2024). Bridging the AI divide: human and responsible AI in news and media industries. *Emerging Media*, 2(3), 335-346.
- Diakopoulos, N. (2019). *Automating the news: How algorithms are rewriting the media*. Harvard University Press.
- Du Bois, J. W. (2003). Discourse and grammar. In *The new psychology of language* (pp. 47-88). Psychology Press.
- Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4), 51-58.
- Esteban, A. J. (2024, December). From Rules to Meaning Making: Teaching Grammar through Discourse Analysis as an Approach. In *Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation* (pp. 1047-1054).
- Goldstein, J. A., Chao, J., Grossman, S., Stamos, A., & Tomz, M. (2024). How persuasive is AI-generated propaganda?. *PNAS nexus*, 3(2).
- Graefe, A. (2016). *Guide to Automated Journalism*. Tow Center for Digital Journalism.
- Ioscote, F., Gonçalves, A., & Quadros, C. (2024). Artificial intelligence in journalism: A ten-year retrospective of scientific articles (2014-2023). *Journalism and Media*, 5(3), 873-891.
- Kuku, D., Charlotte Ojukwu, N. N., Onyoko Omali, T., & Shestakova, A. (2025). Algorithmic Bias in AI-Driven News Production and Dissemination: The Dynamics of X Misinformation During Elections. Lippmann, W. (2017). *Public opinion*. Routledge.
- Liu, A., Wu, Z., Michael, J., Suhr, A., West, P., Koller, A., ... & Choi, Y. (2023). We're afraid language models aren't modeling ambiguity. arXiv preprint arXiv:2304.14399.
- Matz, S. C., Teeny, J. D., Vaid, S. S., Peters, H., Harari, G. M., & Cerf, M. (2024). The potential of generative AI for personalized persuasion at scale. *Scientific Reports*, 14(1), 4692.
- McCombs, M. E., & Shaw, D. L. (1972). The agenda-setting function of mass media. *Public opinion quarterly*, 36(2), 176-187.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6), 1-35.
- Meichanetzidis, K., Gogioso, S., De Felice, G., Chiappori, N., Toumi, A., & Coecke, B. (2020). Quantum natural language processing on near-term quantum computers. arXiv preprint arXiv:2005.04147.
- Nah, S., Luo, J., Kim, S., Chen, M., Mitson, R., & Joo, J. (2024). Algorithmic Bias or Algorithmic Reconstruction? A Comparative Analysis Between AI News and Human News. *International Journal of Communication*, 18.
- Paviour, B. (2025, October 16). AI-generated news sites spout viral slop from forgotten URLs. Nieman Journalism Lab. Retrieved October 18, 2025, from <https://www.niemanlab.org/2025/10/ai-generated-news-sites-spout-viral-slop-from-forgotten-urls/>
- Pazzaglia, S., Vendetti, V., Comencini, L. D., Deriu, F., & Modugno, V. (2025). Passing the Turing Test in Political Discourse: Fine-Tuning LLMs to Mimic Polarized Social Media Comments. arXiv preprint ArXiv:2506.14645 v1 [cs. CL], 17.
- Reubold, J. (2023). *The Democratization of News-Analysis and Behavior Modeling of Users in the Context of Online News Consumption* (Doctoral dissertation, Dissertation, Karlsruhe, Karlsruher Institut für Technologie (KIT), 2022).
- Saha, S., Kalra, K., Patwardhan, M., & Karande, S. (2021, December). Performance of BERT on Persuasion for Good. In *Proceedings of the 18th International Conference on Natural Language Processing (ICON)* (pp. 313-323).
- Salvi, F., Horta Ribeiro, M., Gallotti, R., & West, R. (2025). On the conversational persuasiveness of GPT-4. *Nature Human Behaviour*, 1-9.
- Singh, K., & Ngu, W. (2025, May). Bias-Aware Agent: Enhancing Fairness in AI-Driven Knowledge Retrieval. In *Companion Proceedings of the ACM on Web Conference 2025* (pp. 1705-1712).
- Sui, M. (2025). Writing with emotion? Assessing emotional valence and appeals in AI-generated vs. human-written articles. *AI & SOCIETY*, 1-15.
- The Newsreel Project Consortium. (2021). *Newsreel2: New teaching fields for the next generation of journalists* [Research report]. Erich Brost Institute

- for International Journalism.
<https://doi.org/10.17877/DE290R-22455>
- van Dijk, T. A. (2015). Critical discourse analysis. *The handbook of discourse analysis*, 466-485.
- Waldon, B., Schneider, N., Wilcox, E., Zeldes, A., & Tobia, K. (2025, February 3). Large language models for legal interpretation? Don't take their word for it. *Georgetown Law Journal*, 114 (forthcoming).
<https://doi.org/10.2139/ssrn.5123124>
- Wazni, H., Lo, K. I., McPheat, L., & Sadrzadeh, M. (2024). Large scale structure-aware pronoun resolution using quantum natural language processing. *Quantum Machine Intelligence*, 6(2), 60.
- Widdows, D., Aboumrad, W., Kim, D., Ray, S., & Mei, J. (2024). Natural Language, AI, and Quantum Computing in 2024: Research Ingredients and Directions in QNLP. *arXiv Prepr. arXiv2403.19758*.
- Yin, R., & Liu, X. (2025, May). From Technological Alienation to Value Regression: Ethical Regulation of Algorithmic Bias in the Post-Truth Era. In *2025 IEEE 10th International Conference on Smart Cloud (SmartCloud)* (pp. 32-37). IEEE.
- Zhang, J., He, R., Guo, F., & Liu, C. (2024, March). Quantum interference model for semantic biases of glosses in word sense disambiguation. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 38, No. 17, pp. 19551-19559).