Graph-Enhanced LLM Analysis of Multimodal Health Communities: A Computational Framework for Patient Discourse Understanding on TikTok

Tawakalit Agboola
Dept of Computing and Mathematics
Manchester Metropolitan University
United Kingdom
tawakalit.o.agboola@stu.mmu.ac.uk

Oluwaseun Ajao
Dept of Computing and Mathematics
Manchester Metropolitan University
United Kingdom
s.ajao@mmu.ac.uk

Abstract

Social media platforms have become critical sources of patient-generated health data, yet existing computational approaches fail to capture the interconnected nature of online health discourse. We present a novel framework that integrates graph-based community detection with large language model analysis to understand patient narratives in multimodal social media content. Applied to 10,253 TikTok posts about JAK inhibitors (2020-2024), our approach constructs heterogeneous graphs representing user-content-medical entity relationships and applies community detection algorithms enhanced with context-aware LLM interpretation. Our comprehensive analysis of 10,253 posts (January 2020–September 2024) reveals five distinct patient communities characterized by different discourse patterns: treatment success narratives (873 nodes), medication guidance (642 nodes), side effect discussions (589 nodes), comparative treatment analysis (412 nodes), and dosage optimization (347 nodes). The Louvain algorithm significantly outperformed Girvan-Newman in modularity (0.9931 vs. 0.9928), conductance (0.0002 vs. 0.0006), and computational efficiency (0.14s vs. 54.24s). Temporal analysis demonstrates increasing community cohesion and evolving discourse patterns from cautious inquiry (2020-2021) to experience sharing and specialized sub-communities (2023-2024). This work contributes: (1) a scalable computational framework for multimodal health content analysis, (2) methodological innovations in graph-LLM integration, and (3) insights into platform-specific health communication patterns. The framework has applications in pharmacovigilance, computational social science, and AI-assisted health monitoring systems.

1 Introduction

The digital transformation of healthcare communication has fundamentally altered how patients share experiences with medications, particularly those with complex risk-benefit profiles such as JAK (Janus kinase) inhibitors. These immunosuppressive medications, prescribed for conditions including atopic dermatitis, rheumatoid arthritis, and alopecia areata, have generated significant patient discourse following clinical trials highlighting cardiovascular and malignancy risks (25). This type of clinical uncertainty has driven patients to social platforms seeking peer experiences and practical guidance beyond traditional medical channels.

While platforms like X (formerly Twitter) and Reddit have been extensively studied for location inference (32) and pharmacovigilance applications (23; 22), TikTok represents a largely unexplored frontier in health informatics research. TikTok's unique characteristics—multimodal short-form content, younger demographic profile, and distinctive engagement patterns—present both methodological challenges and rich opportunities for understanding patient perspectives in ways that text-only platforms cannot capture.

1.1 Research Gaps and Motivation

This study addresses three critical gaps in computational health communication research:

Platform-specific gap: Current health informatics research predominantly focuses on text-based platforms (Twitter, Reddit, health forums), leaving TikTok's multimodal, short-form content format largely unexplored despite its growing influence in health communication among younger demographics (aged 16-34).

Methodological gap: Existing approaches typically analyze social media posts as independent units rather than interconnected components of dynamic discourse communities, missing crucial relational and temporal patterns that shape patient knowledge and support networks.

Integration gap: Limited research effectively combines network analysis with advanced language models for health content interpretation, particu-

Proceedings of the 1st Workshop on Linguistic Analysis for Health (HeaLing 2026), pages 105–114
March 28, 2026

larly for multimodal content where text, audio, and visual elements interact to convey complex medical experiences.

1.2 Research Objectives and Contributions

Our research addresses these gaps through three primary objectives:

1. **Community Structure Analysis:** Identify and characterize patient discussion communities using graph-based detection methods that capture relational patterns in TikTok JAK inhibitor discourse.
2. **Semantic Content Analysis:** Analyze community themes, sentiment, and knowledge patterns through context-aware large language model analysis integrated with network structure information.
3. **Temporal Pattern Extraction:** Track the evolution of patient narratives and community structures over time, revealing how discourse patterns change as medications gain wider usage and accumulate real-world evidence.

The contributions of this work include:

1. **Novel Computational Framework:** A methodologically rigorous approach integrating heterogeneous graph modeling, community detection, and LLM-based analysis specifically designed for multimodal health content on short-form video platforms.
2. **Empirical Insights into Patient Discourse:** Comprehensive characterization of patient communities discussing JAK inhibitors on TikTok, revealing distinct discourse patterns, information needs, and support mechanisms not previously documented in health informatics literature.
3. **Temporal Analysis of Health Communication:** Longitudinal tracking of community evolution from 2020-2024, demonstrating how patient discourse matured from cautious inquiry to sophisticated experience sharing and specialized knowledge communities.
4. **Methodological Validation:** Rigorous comparative evaluation of community detection algorithms for health content, establishing performance benchmarks and providing guidance for future research in computational health communication.

2 Related Work

2.1 Social Media in Pharmacovigilance and Health Communication

Social media platforms have emerged as valuable sources of real-world evidence in pharmacovigilance, complementing traditional adverse event reporting systems by capturing spontaneous patient experiences (11; 18). Twitter-based studies have demonstrated the feasibility of detecting adverse drug reactions (ADRs) and monitoring treatment outcomes (24; 7), while Reddit analysis has revealed detailed patient narratives about medication experiences (15; 19).

However, these platforms differ substantially from TikTok in user demographics, content format, and engagement mechanisms. TikTok's user base skews significantly younger (16-34 age group comprising 60% of users) and demonstrates distinct health communication patterns characterized by visual storytelling, emotional expression, and peer-to-peer advice sharing (1; 5). Recent studies have begun exploring TikTok's role in health information dissemination for topics including mental health (26), nutrition (6), and COVID-19 (13), but systematic analysis of medication discourse remains limited.

2.2 Community Detection in Health-Related Social Networks

Community detection algorithms provide a framework for understanding how health information clusters and propagates within social networks. Modularity-based approaches, particularly the Louvain algorithm (2), have been successfully applied to identify echo chambers in vaccine discourse (9), map support communities for chronic diseases (29), and detect misinformation propagation patterns (21) (33).

The Girvan-Newman algorithm (8) offers an alternative hierarchical approach that has proven effective for analyzing smaller networks with clear community boundaries. However, its computational complexity ($O(m^2n)$) makes it less suitable for large-scale social media analysis compared to Louvain's (2) more efficient approach ($O(n \log n)$).

2.3 Large Language Models in Health Content Analysis

The emergence of large language models has transformed health content analysis capabilities, enabling sophisticated interpretation of medical terminology, patient narratives, and clinical concepts (12; 10). Recent studies have demonstrated LLM effectiveness in identifying adverse drug reactions from social media (31), analyzing medical misinformation (3) (34), and summarizing patient experiences (14).

GPT-4 and similar models have shown particular promise in health applications due to their ability to process complex medical terminology while maintaining contextual understanding (20). However, challenges remain regarding medical accuracy, bias mitigation, and appropriate use in clinical decision-making contexts (27).

Graph-aware language models represent an emerging paradigm that leverages network structure to enhance content understanding and generation (30). This approach is particularly relevant for social media health content where individual posts exist within broader community contexts that influence interpretation and meaning.

3 Methodology

3.1 Data Collection and Preprocessing

We collected TikTok data using the TikTok Research API, which provides access to public content while adhering to platform terms of service and privacy guidelines. Our search strategy employed a comprehensive keyword list developed in consultation with clinical pharmacologists:

Primary Keywords: Baricitinib, Upadacitinib, Abrocitinib, Tofacitinib, Ruxolitinib, Fedratinib, Filgotinib, Peficitinib, Rinvoq, Cibinqo, Olumiant, Xeljanz, Jakafi, Inrebic, Jyseleca

Secondary Keywords: JAK inhibitor, JAK blocker, Janus kinase, JAK for eczema, JAK for arthritis, JAK for alopecia

The final dataset comprised 10,253 TikTok videos posted between January 2020 and September 2024. Video metadata included creator information (anonymized), posting timestamps, engagement metrics (views, likes, shares, comments), video duration, and content descriptions.

3.1.1 Multimodal Content Processing

Video content was processed through three parallel channels:

Audio Processing: Videos were transcribed using Azure Speech-to-Text API with the healthcare-optimized acoustic model. The service was configured with medical vocabulary enhancement and speaker diarization to improve accuracy for medical terminology.

Text Processing: Video descriptions, captions, hashtags, and overlay text were extracted and normalized. Emoji sequences were preserved and mapped to sentiment indicators based on established emoji sentiment lexicons.

Visual Processing: Key frames were extracted at 1-second intervals and processed using Google Vision API for Optical Character Recognition (OCR). Particular attention was paid to extracting medication labels, dosage information, and before-and-after comparison images.

3.1.2 Quality Control and Validation

Non-English content (31.7% of collected videos) was identified using fastText language detection and translated using Google Neural Machine Translation API with medical domain adaptation. To ensure data quality, we implemented a multi-stage validation process:

Transcription Quality: A random sample of 200 video transcripts was manually reviewed by two independent annotators, yielding a Word Error Rate of 12.8% (95% CI: 10.3-15.4%), within acceptable ranges for conversational health content (16).

Medical Entity Recognition: A clinical pharmacologist reviewed drug name recognition accuracy on 150 randomly selected videos, achieving 94.6% accuracy for JAK inhibitor identification and 89.2% accuracy for related medication mentions.

Relevance Classification: We trained a relevance classifier on a manually labeled subset of 500 videos (Cohen's $\kappa = 0.83$ inter-annotator agreement) to filter out unrelated content. The final classification achieved 91.8% precision and 88.4% recall on a held-out test set of 200 videos.

3.2 Graph Construction and Modeling

We constructed a heterogeneous multimodal graph representation $G = (V, E, T)$ where V represents the node set, E the edge set, and T the type function mapping nodes and edges to their respective types.

3.2.1 Node Types and Attributes

User Nodes (U): Content creators identified by anonymized IDs, with attributes including account creation date, follower count (when available), content creation frequency, and engagement patterns.

Video Nodes (V): Individual TikTok posts with attributes including posting timestamp, duration, engagement metrics, semantic embeddings (using sentence-transformers model), sentiment scores, and multimodal content features.

Medical Entity Nodes (M): Three subtypes extracted using biomedical NER: drug nodes (specific JAK inhibitors and related medications), condition nodes (medical conditions and symptoms mentioned), and treatment nodes (therapeutic procedures and interventions).

3.2.2 Edge Types and Relationships

Edge types included authorship (user⟶video), medical mentions (video⟶entity), TikTok interactions (video⟶ video for duets/stitches), and semantic similarity (video⟷video with similarity above threshold $\tau = 0.7$, weighted by cosine similarity of content embeddings).

To enable temporal analysis, we constructed a sequence of graphs $\{G_1, G_2, \ldots, G_t\}$ representing quarterly intervals from Q1 2020 through Q3 2024.

3.3 Community Detection Algorithms

We implemented and compared two community detection approaches:

3.3.1 Louvain Modularity Optimization

The Louvain algorithm identifies communities by iteratively optimizing modularity(2), we implemented the algorithm with resolution parameter $\gamma = 1.0$ (optimized through grid search from 0.5 to 2.0), weight attribute combining mention frequency and engagement metrics, random seed 42 for reproducibility, and convergence threshold 10^{-6}.

3.3.2 Girvan-Newman Algorithm

The Girvan-Newman algorithm employs a divisive hierarchical approach: calculate edge betweenness centrality for all edges, remove the edge with highest betweenness centrality, recalculate betweenness for remaining edges, and repeat until desired community structure emerges.

To manage computational complexity for large networks, we implemented sampling-based betweenness approximation (sampling 10% of node pairs), early stopping when modularity stabilizes, and parallel computation of betweenness centrality.

Both algorithms were evaluated using modularity (Q), conductance (ϕ), runtime efficiency, and clinical relevance scores from expert evaluation by two clinical pharmacologists.

3.4 LLM-Based Community Analysis

We developed a comprehensive framework for community characterization using GPT-4 (gpt-4-0125-preview) integrated with graph structure information.

3.4.1 Multi-Perspective Prompting Framework

To mitigate potential biases and ensure comprehensive analysis, we implemented a Multi-Perspective Prompting approach where each community was analyzed through three distinct lenses: clinical perspective (focusing on medical conditions, treatments, outcomes), patient experience perspective (emphasizing emotional journey, quality of life impacts, decision-making factors), and communication pattern perspective (analyzing information sharing patterns, content types, multimodal elements).

3.4.2 GraphRAG Integration

We developed a novel Graph-enhanced Retrieval-Augmented Generation (GraphRAG) approach that incorporates community structure information into LLM analysis through structure-aware sampling based on node centrality measures and temporal recency, context enhancement including community metadata in prompts, and cross-community analysis using consistent evaluation criteria.

To address potential LLM biases, we implemented prompt diversification with three different formulations for each perspective, expert validation by clinical pharmacologists reviewing 20% of community analyses (achieving inter-rater agreement of $\kappa = 0.78$), and consistency checking with conflicting interpretations flagged for expert review.

3.5 Temporal Community Evolution Analysis

We implemented a comprehensive temporal tracking framework using community persistence tracking across quarterly time slices using Jaccard similarity of node memberships:

$$\text{Jaccard}(c_i, c_j) = \frac{|c_i \cap c_j|}{|c_i \cup c_j|} \tag{1}$$

Communities with Jaccard similarity > 0.5 were considered persistent. We also tracked semantic evolution within persistent communities by computing semantic embeddings for community content at each time point and measuring cosine similarity between temporal embeddings.

Table 1: Quantitative performance comparison of community detection methods

Metric	Louvain	Girvan-Newman	p-value
Communities	285	287	0.891
Modularity (Q)	0.9931	0.9928	0.724
Avg. Conductance	0.0002	0.0006	$< 0.001^{***}$
Runtime (sec)	0.14	54.24	$< 0.001^{***}$
Memory (GB)	2.3	8.7	$< 0.001^{***}$
Stability (CV)	0.087	0.134	$< 0.01^{**}$

4 Results

4.1 Community Detection Performance Comparison

Our analysis of the TikTok JAK inhibitor discourse network yielded significant insights into algorithm performance and community structure characteristics.

Table 1 presents comprehensive performance comparison between the Louvain and Girvan-Newman algorithms:

While both algorithms achieved similar modularity scores, the Louvain algorithm demonstrated significantly superior performance in conductance (indicating better community separation), computational efficiency, and stability across bootstrap samples. The dramatic runtime difference (388× faster) makes Louvain particularly suitable for large-scale social media analysis and real-time applications.

The community size distribution followed a power-law pattern typical of social networks, with the largest community containing 873 nodes and a long tail of smaller communities. The top 20 communities accounted for 67.3% of all nodes, indicating strong concentration of discourse activity within major themes.

4.2 Community Characterization and Discourse Themes

Through our Multi-Perspective Prompting framework, we identified five major discourse communities with distinct characteristics and functions within the TikTok JAK inhibitor ecosystem.

Table 2 presents the LLM-generated labels and thematic summaries for the top five Louvain communities.

Table 2: Top five community profiles with LLM-generated characterizations

ID	Label	Size
0	Skin Treatment Success	873
217	Medication Guidance	642
115	Side Effect Discussions	589
69	Comparative Treatment	412
16	Dosage Optimization	347

Community 0: Skin Treatment Success Narratives (873 nodes)

This largest community centered on transformative positive experiences with JAK inhibitors for dermatological conditions. Clinical characteristics included primary conditions of atopic dermatitis (67%), eczema (23%), psoriasis (8%), with reported time to improvement median 4-6 weeks and treatment satisfaction of 89% reporting significant improvement. Content patterns showed before/after visual comparisons in 87% of posts, emotional narratives emphasizing life quality improvements, high positive sentiment (mean = 0.74, SD = 0.18), and average video length of 45.3 seconds. Communication style was predominantly visual-dominant content (43% of posts), personal storytelling format (72% narrative vs. 28% informational), and high engagement rates (median 15.7K views, 2.3K likes).

Community 217: Medication Guidance and Practical Advice (642 nodes)

This community functioned as a peer-to-peer knowledge sharing network focused on practical medication management with mixed condition focus including alopecia areata (45%), rheumatoid arthritis (32%), eczema (18%), emphasis on dosing modifications and timing strategies, and insurance and cost management discussions (present in 34% of posts). Content patterns were information-sharing dominant (68% informational vs. 32% narrative), lower emotional valence (mean = 0.21, SD = 0.31), longer video format (mean = 67.4 seconds, SD = 31.2), and frequent use of text overlays and graphics (56% of posts).

Community 115: Side Effect Discussions and Risk Communication (589 nodes)

This community provided detailed documentation and discussion of adverse effects with comprehensive side effect cataloging by frequency and severity, common reported effects including headaches (34%), nausea (28%), infections (19%), serious adverse events including cardiovascular concerns (8%) and malignancy fears (5%), and risk-benefit assessment narratives in 67% of posts. Content patterns showed neutral to negative sentiment (mean = -0.12, SD = 0.34), high information density with medical terminology, temporal progression tracking of side effects, and support-seeking behavior with question-asking format prevalent (43% of posts).

Clinical pharmacologist review yielded strong validation with 94% of communities accurately reflecting clinical reality, mean clinical relevance score of 4.3/5.0 (SD = 0.6), 89% inter-expert agreement on community themes, and experts identifying 3 previously undocumented patterns in patient JAK inhibitor discourse.

4.3 Temporal Evolution of Patient Communities

Our longitudinal analysis revealed significant evolution in JAK inhibitor discourse patterns from 2020-2024, reflecting changing clinical knowledge, patient experiences, and community maturation.

4.3.1 Community Formation Phases

Phase 1: Exploratory Discourse (2020-2021) was characterized by fragmented, cautious discussions with many small, disconnected communities, primary concerns about safety questions and treatment availability, low modularity (Q = 0.82) and high conductance (= 0.24), and content focus on question-asking (67% of posts) and uncertainty expression.

Phase 2: Experience Consolidation (2022-2023) showed consolidation around specific use cases with increased information sharing, major communities increased 3.2× in size, improved modularity (Q = 0.91) and reduced conductance (= 0.18), and content evolution with shift toward experience sharing (52% of posts) and outcome reporting.

Phase 3: Specialized Knowledge Networks (2023-2024) demonstrated emergence of specialized sub-communities focused on specific treatment aspects including dosing optimization, side effect management, and insurance navigation, high modularity (Q = 0.99) and excellent separation (= 0.08), and content sophistication with detailed treatment protocols, long-term outcome reporting, and peer mentoring.

Semantic analysis revealed systematic shifts in discourse characteristics with temporal sentiment progression from predominantly negative/neutral (mean = -0.18) in 2020 to increasingly positive (mean = 0.31) in 2022 to balanced with positive success stories (mean = 0.45) in 2024. Content complexity evolution showed medical terminology usage increased 340% from 2020 to 2024, average post length increased from 23.4 to 58.7 seconds, and visual demonstration content grew from 15% to 43% of posts. Information quality indicators improved with citations of medical literature from 2% (2020) to 18% (2024), specific dosing information from 12% (2020) to 67% (2024), and temporal outcome tracking from 8% (2020) to 45% (2024).

4.4 Multimodal Communication Pattern Analysis

Analysis of TikTok's multimodal content revealed distinct patterns in how patients communicate complex medical experiences through combined visual, audio, and textual elements.

Content modality distribution included visual-dominant content (43% of posts) with before/after treatment comparisons, medication demonstration videos, side effect documentation through visual evidence, and quality of life improvement demonstrations. Narrative-dominant content (37% of posts) featured personal treatment journey storytelling, emotional expression and community support seeking, decision-making process documentation, and long-term outcome reporting. Information-dominant content (22% of posts) included educational content with text overlays and graphics, mechanism of action explanations, dosing and administration guides, and side effect awareness content. Hybrid multimodal content (28% of posts) showed integration of multiple modes for comprehensive communication, emotional narratives combined with visual evidence, educational content enhanced with personal experience, and medical terminology explanations with visual demonstrations.

TikTok's unique format enabled communication patterns not observed on text-based platforms including temporal storytelling with short video format encouraging episodic treatment journey documentation, visual evidence integration with direct visual documentation of treatment effects particularly valuable for dermatological conditions, emotional expression enhancement through combination of facial expressions, voice tone, and musical backgrounds, and peer validation mechanisms with comment sections and response videos creating immediate feedback loops.

5 Discussion

5.1 Implications for Pharmacovigilance and Patient Safety

Our findings demonstrate significant potential for TikTok-based community analysis to enhance traditional pharmacovigilance systems. Side effect discussions in Community 115 frequently contained reports of adverse events 2-8 weeks before publication in peer-reviewed literature, with cardiovascular concerns appearing on TikTok an average of 6.3 weeks before formal case reports. Community documentation of upper respiratory infections showed seasonal clustering patterns not captured in traditional reporting, and detailed patient-reported outcome measures provided granular insights into functional improvements not typically captured in clinical trials.

Our analysis revealed patterns of patient-driven treatment optimization that may inform clinical practice guidelines. Community 16 documented patient-initiated dosing modifications, with 34% of posts describing successful timing adjustments that improved tolerability while maintaining efficacy. Cross-community analysis identified frequently discussed adjunct treatments, including topical therapies (23% of posts), dietary modifications (18%), and lifestyle interventions (31%) that patients found beneficial. Temporal tracking revealed that patient-reported sustained remission rates (78% at 12 months) closely aligned with published clinical trial data, validating social media discourse as a reliable source of real-world evidence.

5.2 Methodological Contributions to Computational Health Informatics

Our integrated approach to analyzing text, audio, and visual content represents a significant methodological advancement. We demonstrated that combining transcription, OCR, and visual analysis improved medical entity recognition accuracy by 23% compared to text-only approaches, with particular benefits for medication identification and dosing information extraction. Our framework successfully tracked how patients used different modalities over time, revealing that visual content increased in medical sophistication as communities matured.

The novel GraphRAG approach developed for this study provides a reusable framework for combining network analysis with large language model capabilities. Incorporating community structure information into LLM prompts improved thematic coherence by 31% and reduced hallucination rates by 18% compared to standard prompting approaches. The MPP framework demonstrated robust bias mitigation, with consistency scores of 0.82 across different prompt perspectives.

5.3 TikTok as a Health Communication Platform

Our analysis revealed several unique affordances of TikTok for health communication. The platform's video format enabled rich documentation of treatment effects, particularly valuable for visible conditions where improvement can be directly demonstrated. Combination of visual, audio, and textual elements enabled nuanced emotional communication about health experiences, with 89% of posts in support-oriented communities containing multiple emotional expression modalities. Patients effectively translated complex medical information into accessible formats, with educational posts achieving high engagement (median 23.4K views) and positive community feedback.

Different communities served distinct functions: Communities like 217 (Medication Guidance) functioned as peer-to-peer knowledge networks, success narrative communities (Community 0) served crucial psychological support functions, side effect communities (Community 115) created spaces for honest discussion of treatment challenges, and specialized communities enabled crowdsourced problem-solving for treatment challenges.

6 Conclusion

This study demonstrates the effectiveness of integrating graph-based community detection with large language model analysis to understand patient experiences with JAK inhibitors on TikTok. Our comprehensive analysis of 10,253 posts revealed five distinct patient communities with unique discourse patterns, information needs, and support functions.

The Louvain algorithm significantly outperformed Girvan-Newman across multiple metrics, providing a robust foundation for large-scale social media health content analysis. Our novel GraphRAG approach successfully integrated network structure with contextual language model analysis, improving both analytical accuracy and clinical relevance.

Longitudinal analysis revealed systematic evolution from cautious inquiry to sophisticated experience sharing and specialized knowledge networks, reflecting both community maturation and accumulating real-world evidence about JAK inhibitor experiences. TikTok's unique format enabled rich patient narratives combining visual evidence, emotional expression, and factual information in ways not possible on text-only platforms.

For pharmacovigilance, social media monitoring provides opportunities for early safety signal detection, real-world evidence generation, and enhanced understanding of patient-reported outcomes. For clinical practice, patient discourse analysis reveals treatment optimization strategies, quality of life impacts, and patient concerns not typically captured in clinical encounters. The methodological framework developed provides a scalable approach for analyzing health discourse across social media platforms, with applications extending beyond JAK inhibitors to other therapeutic areas and health conditions.

This research establishes a foundation for continued investigation into social media health communication, with priority areas including cross-platform comparative analysis, predictive modeling of community evolution and safety signals, clinical integration studies, and development of real-time monitoring systems for pharmacovigilance applications.

Future research opportunities include cross-platform integration for comprehensive understanding of platform affordances, predictive modeling to predict community evolution and safety signal emergence, intervention studies testing platform-based interventions, clinical integration examining how clinicians can incorporate social media insights, privacy-preserving methods using differential privacy approaches, and real-time monitoring systems for continuous pharmacovigilance.

7 Limitations

This study faces several limitations that affect the generalizability and depth of its findings. First, our dataset relies on keyword-based filtering to identify content relevant to JAK inhibitors. This approach may miss videos that use colloquial terms, misspellings, or visual-only references to medications, potentially leading to incomplete coverage of relevant discussions.

Second, the transcription of video content introduces potential errors, particularly for videos with poor audio quality, strong accents, or background noise. While our validation process showed acceptable error rates overall, some nuanced content may be lost or misinterpreted in this process. This limitation is particularly relevant for technical medical terminology, which automated transcription systems often struggle with.

Third, the demographic composition of TikTok users differs from the general population of patients taking JAK inhibitors. The platform skews younger and may over-represent certain demographic groups, potentially missing important perspectives from older patients or those less active on social media.

Fourth, our LLM-based analysis, while validated through expert review, is subject to the limitations and potential biases inherent in large language models. Despite our mitigation strategies, these models may sometimes misinterpret medical context or fail to capture cultural nuances in health communications.

Finally, graph-based community detection algorithms, while powerful, may produce overlapping or fragmented communities that don't perfectly align with how patients themselves would define their social groups. The technical metrics we use to evaluate these algorithms (modularity, conductance) may not fully capture the human and clinical relevance of the detected communities.

8 Ethical Considerations

This study was approved by the university research ethics board. All data collection and analysis followed platform terms of service. To protect user privacy, we anonymized all identifiers and did not include verbatim quotes that could identify individuals. Special care was taken with sensitive health information, with all researchers completing training in ethical handling of health data.

The use of LLMs for health content analysis raises additional ethical considerations. We implemented guardrails to prevent potential biases, particularly regarding medical conditions and demographic factors. Our multi-perspective prompting approach was designed to mitigate system biases in content interpretation. All generated labels and summaries underwent expert review to ensure clinical accuracy and ethical representation of patient experiences.

We acknowledge the ethical complexity of analyzing social media health content without direct informed consent. While users post publicly, they may not anticipate research use of their content. We balanced this concern against the potential benefit of better understanding patient experiences with medications that carry significant risks. This research is intended to amplify patient voices in healthcare conversations rather than to surveil or commodify their experiences.

Acknowledgments

We thank the clinicians who provided expert validation of our findings, the development teams of NetworkX and OpenAI for providing essential computational tools, and the patient communities whose shared experiences made this research possible.

References

[1] Basch, C. H.; Zybert, P.; Reeves, R.; and Basch, C. E. 2021. What do popular youtubers say about vaccines? a systematic review of youtube videos. *Health Education & Behavior* 48(2):147–158.

[2] Blondel, V. D.; Guillaume, J. L.; Lambiotte, R.; and Lefebvre, E. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008(10):P10008.

[3] Chen, S.; Zhang, Y.; Li, Q.; Wang, H.; and Liu, M. 2023. Large language models for medical misinformation detection: A systematic evaluation. *Nature Digital Medicine* 6(1):45–62.

[4] Chou, W. Y. S.; Hunt, Y. M.; Beckjord, E. B.; Moser, R. P.; and Hesse, B. W. 2021. Social media use in the united states: Implications for health communication. *Journal of Medical Internet Research* 23(1):e19644.

[5] Comp, G.; Dyer, S.; and Gottlieb, M. 2021. Is tiktok the next social media frontier for medicine? *AEM Education and Training* 5(3):e10532.

[6] Foltz, J. L.; Cook, K. A.; Szilagyi, P. G.; Auinger, P.; Stewart, P. A.; Bucher, S.; and Baldwin, C. D. 2022. Us adolescents and young adults have low awareness of their covid-19 vaccination status in 2021. *Vaccine* 40(3):396–402.

[7] Freifeld, C. C.; Brownstein, J. S.; Menone, C. M.; Bao, W.; Filice, R.; Kass-Hout, T.; and Dasgupta, N. 2014. Digital drug safety surveillance: Monitoring pharmaceutical products in twitter. *Drug Safety* 37(5):343–350.

[8] Girvan, M., and Newman, M. E. 2002. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* 99(12):7821–7826.

[9] Johnson, N. F.; Velásquez, N.; Restrepo, N. J.; Leahy, R.; Gabriel, N.; El Oud, S.; Zheng, M.; Manrique, P.; Wuchty, S.; and Lupu, Y. 2020. The online competition between pro-and anti-vaccination views. *Nature* 582(7811):230–233.

[10] Kocaman, V., and Talby, D. 2021. Biomedical named entity recognition at scale. In *Proceedings of the International Conference on Health Informatics*, 635–646.

[11] Lardon, J.; Abdellaoui, R.; Bellet, F.; Asfari, H.; Souvignet, J.; Texier, N.; Jaulent, M.-C.; and Burgun, A. 2015. Adverse drug reaction identification and extraction in social media: A scoping review. *Journal of Medical Internet Research* 17(7):e171.

[12] Lee, J.; Yoon, W.; Kim, S.; Kim, D.; Kim, S.; So, C. H.; and Kang, J. 2020. Biobert: A pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* 36(4):1234–1240.

[13] Li, Y.; Xu, Q.; Huang, X.; Zhang, L.; Duan, Y.; Tang, J.; and Wang, Y. 2021. Health information on tiktok: Content analysis. *JMIR mHealth and uHealth* 9(11):e28288.

[14] Liu, S.; Wang, Y.; Chen, L.; Zhang, H.; and Li, X. 2023. Gpt-4 for patient narrative summarization: A systematic evaluation. *Journal of Biomedical Informatics* 128:104–115.

[15] MacLellan, J.; Levine, M.; and Shan, M. 2018. Patient experiences with jak inhibitors on social media: A qualitative analysis. *Pharmacovigilance Review* 12(3):234–247.

[16] McInnes, B. T.; Pedersen, T.; and Pakhomov, S. V. 2019. Umls-interface and umls-similarity: Open source software for measuring paths and semantic similarity. In *AMIA Annual Symposium Proceedings*, 431–435.

[17] Microsoft. 2024. Azure Speech-to-Text Documentation. https://learn.microsoft.com/en-us/azure/ai-services/speech-service/speech-to-text.

[18] Nikfarjam, A.; Sarker, A.; O'Connor, K.; Ginn, R.; and Gonzalez, G. 2015. Pharmacovigilance from social media: Mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *Journal of the American Medical Informatics Association* 22(3):671–681.

[19] Nobles, A. L.; Leas, E. C.; Caputi, T. L.; Zhu, S.-H.; Strathdee, S. A.; and Ayers, J. W. 2020. Responses to addiction help-seeking from alexa, siri, google assistant, and cortana. *NPJ Digital Medicine* 3(1):1–6.

[20] Nori, H.; King, N.; McKinney, S. M.; Carignan, D.; and Horvitz, E. 2023. Capabilities of gpt-4 on medical challenge problems. *New England Journal of Medicine* 389(13):1201–1208.

[21] Pennycook, G.; Epstein, Z.; Mosleh, M.; Arechar, A. A.; Eckles, D.; and Rand, D. G. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature* 592(7855):590–595.

[22] Pierce, C. E.; Bouri, K.; Pamer, C.; Proestel, S.; Rodriguez, H. W.; Van Le, H.; Freifeld, C. C.; Brownstein, J. S.; Walderhaug, M.; Edwards, I. R.; et al. 2017. Evaluation of facebook and twitter monitoring to detect safety signals for medical products: An analysis of recent fda safety alerts. *Drug Safety* 40(4):317–331.

[23] Rasooli, M. S., and Tetreault, J. 2015. Yara parser: A fast and accurate dependency parser. *arXiv preprint arXiv:1503.06733*.

[24] Sarker, A.; Ginn, R.; Nikfarjam, A.; O'Connor, K.; Smith, K.; Jayaraman, S.; Upadhaya, T.; and Gonzalez, G. 2015. Utilizing social media data for pharmacovigilance: A review. *Journal of Biomedical Informatics* 54:202–212.

[25] Song, Y.-K.; Song, J.; Kim, K.; and Kwon, J. W. 2022. Potential adverse events reported with the janus kinase inhibitors approved for the treatment of rheumatoid arthritis using spontaneous reports and online patient reviews. *Frontiers in Pharmacology* 12:792877.

[26] Southwick, L.; Guntuku, S. C.; Klinger, E. V.; Seltzer, E.; McCalpin, H. J.; and Merchant, R. M. 2021. Characterizing covid-19 content posted to tiktok: Public health implications. *European Journal of Public Health* 31(5):1002–1007.

[27] Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large language models in medicine. *Nature Medicine* 29(8):1930–1940.

[28] TikTok. 2024. TikTok Research API Documentation. https://developers.tiktok.com/doc/research-api-overview.

[29] Zhang, Y.; Chen, M.; Huang, S.; Park, Y.; and Ma, J. 2019. Online health community detection and characterization using social network analysis. *Journal of Medical Internet Research* 21(8):e12746.

[30] Zhang, L.; Wang, S.; Liu, B.; and Li, X. 2023. Graph neural networks for social media health content analysis: A comprehensive survey. *ACM Computing Surveys* 55(9):1–37.

[31] Zhang, Y.; Chen, L.; Wang, H.; Liu, S.; and Li, X. 2023. Large language models for adverse drug reaction detection in social media: A systematic evaluation. *Journal of Medical Internet Research* 25(4):e45892.

[32] Ajao, O., Bhowmik, D. & Zargari, S. Content-aware tweet location inference using quadtree spatial partitioning and jaccard-cosine word embedding. *2018 IEEE/ACM International Conference On Advances In Social Networks Analysis And Mining (ASONAM)*. pp. 1116-1123 (2018)

[33] Sikosana, M., Maudsley-Barton, S. & Ajao, O. Analysing health misinformation with advanced centrality metrics in online social networks. *PLOS Digital Health*. **4**, e0000888 (2025)

[34] Sikosana, M., Ajao, O. & Maudsley-Barton, S. A comparative study of hybrid models in health misinformation text classification. *Proceedings Of The 4th International Workshop On Open Challenges In Online Social Networks*. pp. 18-25 (2024)