



MedShort: Medicine Shortage Detection on Social Media with Privacy Preservation

Xueping Peng*
xueping.peng@uts.edu.au
University of Technology Sydney
Ultimo, NSW, Australia

Peng Yan
peng.yan-1@uts.edu.au
University of Technology Sydney
Ultimo, NSW, Australia

Tao Shen
tao.shen@uts.edu.au
University of Technology Sydney
Ultimo, NSW, Australia

Ming Xie
ming.xie@uts.edu.au
University of Technology Sydney
Ultimo, NSW, Australia

Guodong Long*
guodong.long@uts.edu.au
University of Technology Sydney
Ultimo, NSW, Australia

Abstract

Medicine shortages pose a significant threat to healthcare systems worldwide, disrupting patient care and increasing costs. Social media platforms, with their real-time user-posted content, offer a valuable opportunity to detect early signs of such shortages. This paper introduces **MedShort**, a novel framework that leverages large language models (LLMs) to monitor, analyze, and detect **Medicine Shortages** from social media data while preserving user privacy. Key features include filtering relevant posts, extracting specific medicines and geographic insights, and providing actionable analytics for stakeholders. Experimental results show MedShort's effectiveness, with refined LLMs achieving high accuracy, precision, and recall across several medicines. By integrating privacy-preserving mechanisms, MedShort ensures ethical data handling, enabling near real-time and scalable monitoring of medicine shortages.

CCS Concepts

• **Computing methodologies** → **Natural language processing; Information extraction.**

Keywords

Medicine Shortage, Social Media, Privacy Preservation, Large Language Models, LLMs

ACM Reference Format:

Xueping Peng, Peng Yan, Tao Shen, Ming Xie, and Guodong Long. 2025. MedShort: Medicine Shortage Detection on Social Media with Privacy Preservation. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3701716.3715187>

1 Introduction

Medicine shortages have become a critical issue, significantly affecting patient outcomes, healthcare providers, and pharmaceutical supply chains [10]. Early identification of shortages is essential so

that health professionals, medicine sponsors and the health sectors can take action to assist patients [6]. The widespread use of social media platforms such as Reddit, and X provides a promising opportunity for timely and dynamic monitoring. These platforms serve as hubs of real-time user-generated content, where individuals share experiences, complaints, and concerns, including those related to medicine availability [9]. However, leveraging such data comes with challenges: extracting meaningful signals from noisy data, ensuring the scalability of analysis, and addressing privacy concerns [1] due to the sensitive nature of shared health information.

This paper presents **MedShort**, a comprehensive framework designed to detect and analyze **Medicine Shortages** using advanced artificial intelligence techniques. MedShort employs state-of-the-art large language models (LLMs) to filter, classify, and analyze social media posts, enabling accurate identification of shortage-related discussions. To optimize the balance between accuracy and efficiency, MedShort consists of key modules like real-time keyword-based filtering, extraction of medicines and locations, and visualization of shortage trends and geographic disparities. Furthermore, privacy preservation is also a cornerstone of the framework – we propose to incorporate robust de-identification processes to ensure the ethical handling of sensitive data.

As such, MedShort holds great potential to provide complementary and actionable insights to stakeholders, including policymakers, healthcare professionals, and pharmaceutical companies. By integrating scalable AI-driven analytics with privacy preservation mechanisms, MedShort addresses the dual challenges of processing vast amounts of user-generated data and maintaining privacy preservation. The tool's capability to identify emerging shortage patterns and geographical impacts can support decision-making and resource targeting. This paper demonstrates MedShort's effectiveness through experiments on real-world social media datasets, showcasing its potential to further medicine shortage detection and communication in healthcare systems.

2 Overview of MedShort

2.1 Existing Tools

2.1.1 Social Media Monitoring Tools. Tools like Meltwater [8], Sprout Social [12], and Hootsuite [4] are widely used for managing and monitoring social media accounts. These tools provide customizable dashboards for tracking conversations, understanding community sentiment, and identifying valuable mentions. They

*Corresponding authors.



This work is licensed under a Creative Commons Attribution 4.0 International License. *WWW Companion '25, Sydney, NSW, Australia*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1331-6/2025/04
<https://doi.org/10.1145/3701716.3715187>

are primarily designed to enhance brand engagement and public relations. However, they operate as online services, where users have limited control over data and models. Unlike these tools, the proposed MedShort offers local deployment, enabling users to store data securely, eliminate biases, and trigger real-time alerts for specific events like medicine shortages.

2.1.2 Social Network Analysis (SNA) Tools. Commercial tools such as Meltwater [8] and Isentia [5] and open-source software like Gephi [2] and Socnetv [11] focus on analyzing relationships and structures within social networks using nodes and links. These tools primarily provide network visualization and structural analysis of social interactions. While effective for visualizing connections, they lack the capability for detailed conversation analysis and event detection. The MedShort combines SNA capabilities with robust conversation analysis, making it suitable for identifying critical events like medicine shortages and displaying relevant demographic insights.

2.2 Overview of Architecture

Figure 1 illustrates the overall architecture and the processing flow of MedShort. Each module can be featured as follows.

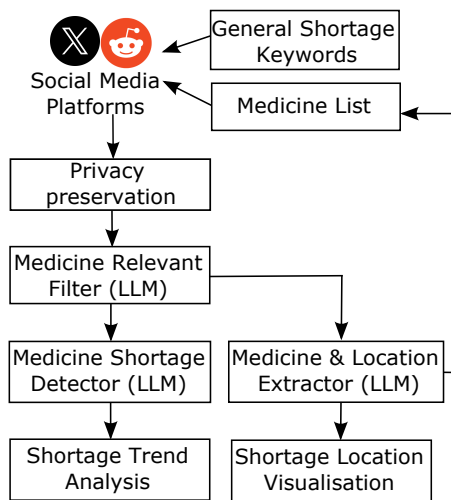


Figure 1: Overall architecture.

- **General Shortage Keywords:** This module identifies and processes a predefined set of keywords related to shortages as a recall stage (a.k.a first-stage retrieval in an efficient way). These keywords serve as the foundation for detecting relevant discussions on social media, ensuring the tool captures diverse contexts of medicine shortages.
- **Medicine List:** This module maintains a comprehensive database of medicines, enabling the system to match posts or discussions with specific medicines. This ensures precise identification of medicine-related content.
- **Privacy Preservation:** This critical module ensures that all data collected and analyzed adheres to privacy standards by de-identifying user information. It helps maintain compliance with data protection regulations while enabling data compatibility

with federated learning frameworks for private-preserving model training.

- **Medicine Relevant Filter (LLM):** This module, built upon an LLM, filters out irrelevant content from the dataset to focus solely on posts that are contextually relevant to medicines. It performs as a re-ranker and improves the precision of the system.
- **Medicine Shortage Detector (LLM):** Powered by LLMs, this module identifies posts or conversations that discuss potential medicine shortages. It leverages natural language understanding to pinpoint contextually relevant information.
- **Medicine & Location Extractor (LLM):** This module extracts the mentioned medicines and locations from the contents of social media posts. It uses advanced language processing techniques to identify and link medicines with specific regions accurately.
- **Shortage Trend Analysis:** This module analyzes the temporal patterns of detected shortages to identify trends. It helps in understanding the dynamics of medicine availability over time and predicting potential future shortages.
- **Shortage Location Visualization:** This module presents a geographical visualization of shortage data, allowing users to easily identify regions most affected by medicine shortages. The visualization is only possible if the geographical data is available in the post. It provides actionable insights for stakeholders to address location-specific issues.

3 Demo Scenario

The following demo scenarios illustrate the capabilities of the MedShort Tool in addressing key challenges in identifying and monitoring medicine shortages. By leveraging advanced data analysis and visualization features, the tool empowers policymakers, healthcare professionals, industry stakeholders, and analysts to gain actionable insights from social media conversations that complement other reporting mechanisms.

Demo Scenario 1: Identifying Trends in Medicine Shortage Posts. The MedShort Tool enables users to track and analyze trends in medicine shortage discussions over a specific period. In this scenario shown in the left hand of Figure 2, a healthcare analyst uses the tool to monitor social media platforms for posts related to medicine shortages. By selecting “All” platforms and monitoring the date range from November 1, 2024, to December 3, 2024, the tool generates a detailed bar chart showing shortage-relevant posts to non-relevant posts over time. The analyst observes more discussions related to Vyvanse in mid-November compared to late November and early December, which may suggest shortages eased in December. However, these trends could be explained by many factors which require further investigation. This feature allows stakeholders to monitor patterns over time and detect signals that may be indicative of shortages, potentially prompting further analysis and timely responsiveness to emerging shortages.

Scenario 2: Geographical Visualization of Shortage Data. In this scenario shown on the right of Figure 2, a public health officer employs the MedShort Tool to visualize regional differences in medicine shortage content. After selecting the “Map” view and specifying the same date range, the tool produces a heatmap showcasing areas with the highest number of shortage-relevant posts. For

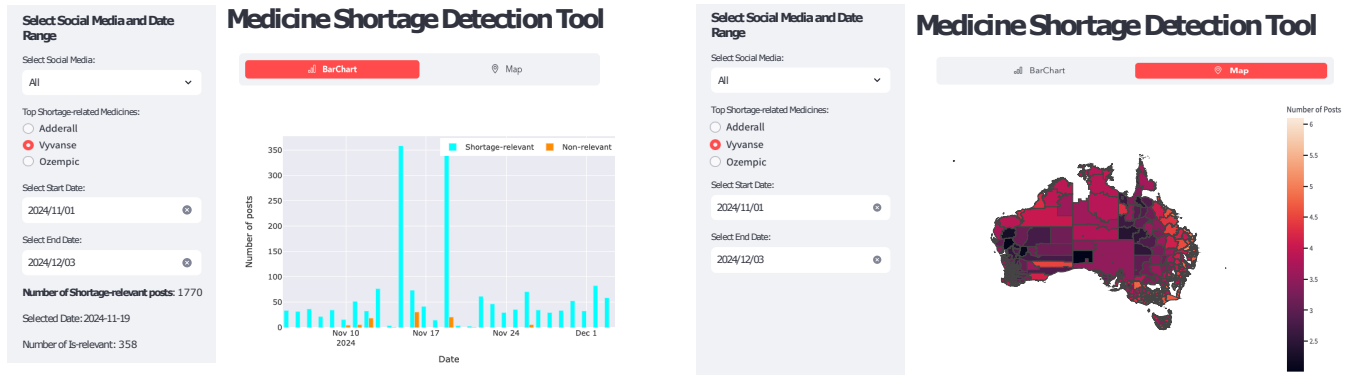


Figure 2: Medicine shortage detection tool interface.

Table 1: Statistics of data collection for three medicines

Medicine	# of Twitter	# of Reddit	Time Period
Vyvanse	1,079	498	Feb-Aug 2024
Adderall	835	197	Aug 2024
Ozempic	3,559	1,736	Mar-Apr 2024

Table 2: Data labeling for three medicines

Medicine	# Shortage-related	# Non-related	Total No.
Vyvanse	108	11	119
Adderall	143	14	157
Ozempic	33	26	59

instance, the visualization may highlight dense clusters of posts in specific regions, which may signal localized supply issues. This geographic representation provides the officer with actionable insights into which areas are most vocal about the shortage, potentially signalling the need for medical resources or communication strategies. Note that this scenario is a potential outcome rather than what has already been established. At this stage, the heatmap only shows the population by LGA and is not related to posted content.

4 Experiments & Results

The experiments aim to analyze an LLM’s baseline performance on human-labeled social media data to classify text as either shortage-related or non-related, and subsequently refine the model through in-context learning, which lets LLMs adapt to tasks by observing examples given in the prompt, without retraining.

4.1 Dataset

4.1.1 Data Collection. We conducted data collection using the official APIs of X and Reddit to demonstrate the potential of MedShort, posts were gathered on three specific medicines: Vyvanse, Adderall and Ozempic, which are not limited to posts from Australia. Vyvanse and Ozempic are candidates given by the Therapeutic Goods Administration (TGA) for model development, and Adderall is an alternative to Vyvanse. While Adderall has not been approved by the TGA for sale in Australia, we included the collected posts mentioning the medicine to evaluate the performance of the LLM. The keywords used for this collection were combinations of each medicine’s name with the term "shortage," such as "Vyvanse shortage." Table 1 summarizes the statistics of the data collected for each medicine across both social media platforms.

4.1.2 Data Labeling. Following data collection, the team undertook a labelling process to categorize the posts based on their relevance to medicine shortages. This involved manually reviewing a subset of the collected posts for each medicine to determine whether they were "Shortage-related" or "non-related". Data labelling criteria for a specific medicine included: 1) Shortage-related: a post had to include the specific medicine name and referencing or discussing the shortage of that medicine. This included posts relating to "lived experienced" content, and promotional/advertising material; 2) Non-related: this included all content that was not classified as shortage-related, such as reference to the shortage of a different medicine to the one in the given name, non-specified medicine shortage, and side effects of some medicines. Table 2 presents the results of this labelling effort.

4.2 Demo Setting

We deployed a powerful open-source LLM, Llama3.1 8B-q8 [7] (llama3.1:8b-instruct-q8_0 with Model size 8.03 billion parameters and 8.5GB Memory) on a system equipped with an Intel Core i9 CPU, utilizing 20 CPU cores and two TITAN XP 12GB GPUs within the secure environment. The primary objective was to evaluate the model’s capabilities in identifying shortage-related posts concerning three medicines: Vyvanse, Adderall, and Ozempic.

We also select GPT-4o [3] (roughly 1.8 trillion parameters) as a state-of-the-art benchmark, offering performance metrics closely aligned with human evaluations. De-identified posts are summarised by Llama 3.1 8B-q8 before being sent to GPT-4o for analysis. While their meaning is preserved, the summarised posts differ from the original de-identified content, ensuring data privacy and adherence to ethical standards for handling social media data.

Table 3: Performance on Llama3.1

Medicine	LLMs	Acc	Prec	Recall	F1	AUC
Vyvanse	Llama3.1-Ori	0.90	0.97	0.92	0.95	0.75
	Llama3.1-Ref	0.94	0.97	0.96	0.97	0.77
Adderall	Llama3.1-Ori	0.88	0.95	0.92	0.93	0.69
	Llama3.1-Ref	0.93	0.97	0.95	0.96	0.80
Ozempic	Llama3.1-Ori	0.54	0.55	0.94	0.70	0.49
	Llama3.1-Ref	0.76	0.76	0.83	0.71	0.77

Table 4: Performance on GPT-4o

Medicine	LLMs	Acc	Prec	Recall	F1	AUC
Vyvanse	GPT-4o-Ori	0.90	1.0	0.89	0.94	0.94
	GPT-4o-Ref	0.94	1.0	0.93	0.97	0.97
Adderall	GPT-4o-Ori	0.91	0.98	0.92	0.95	0.89
	GPT-4o-Ref	0.95	1.0	0.95	0.97	0.97
Ozempic	GPT-4o-Ori	0.76	0.78	0.78	0.78	0.76
	GPT-4o-Ref	0.89	0.93	0.87	0.90	0.89

4.3 Performance

4.3.1 Comparison on Llama3.1. The comparison shown in Table 3 of Llama3.1 between the original (Llama3.1-Ori) and refined (Llama3.1-Ref) highlights that Llama3.1 has substantial improvements in its ability to classify shortage-related social media posts for Vyvanse, Adderall, and Ozempic. The improved metrics across accuracy, precision, and F1-score suggest that the model is now more effective in capturing relevant posts while minimizing errors. The increase in AUC values for Vyvanse, Adderall, and Ozempic further indicates an enhanced discriminatory ability, making the model more robust in distinguishing between shortage-related and non-shortage-related content.

4.3.2 Comparison on GPT-4o. The results in Table 4 show a clear improvement in the refined version of GPT-4o compared to the original model across all three medicines: Vyvanse, Adderall, and Ozempic. For Vyvanse, the refined model achieved an increase in accuracy and recall, while maintaining perfect precision (1.0) and improving the AUC score. For Adderall, the refined model showed similar gains, with increased accuracy, perfect precision (1.0), and an improved AUC score. For Ozempic, the improvements were even more pronounced, with an increase in accuracy, precision, and an improved F1 score, highlighting a substantial overall performance boost.

4.3.3 Comparison between Llama3.1 and GPT-4o. The comparison between Llama3.1 and GPT-4o reveals that while both models improved with refinement, GPT-4o consistently outperforms Llama3.1 in precision and AUC across all three medicines. It achieves near-perfect precision for Vyvanse and Adderall and outperforms all metrics compared to Llama3.1 for Ozempic. Overall, focusing on F1 and AUC scores, GPT-4o is the more robust and high-performing model for identifying medicine shortage-related content.

5 Conclusion & Future Work

This paper presents MedShort, a framework that leverages social media for near real-time detection of medicine shortages while ensuring privacy preservation. By employing advanced LLMs for data filtering and shortage detection, MedShort demonstrates significant potential to enhance insights for healthcare stakeholders. The experimental results highlight the efficacy of refined LLMs in accurately classifying shortage-related posts. Additionally, the integration of privacy preservation mechanisms ensures the ethical handling of sensitive data.

The project is still ongoing and further work is underway to improve the tool’s capability. Future work will focus on expanding the scope of MedShort by incorporating additional data sources to further improve detection accuracy. Enhancing geographic analysis with real-time mapping of supply chain disruptions and exploring multilingual capabilities to support global applications are also key priorities. In addition, efforts will be directed toward improving scalability, ensuring MedShort can handle larger datasets and provide faster, more comprehensive insights to mitigate the impact of medicine shortages effectively.

Acknowledgments

This paper relates to a project supported by the Digital Health CRC Limited (“DHCRC”). DHCRC is funded under the Australian Commonwealth’s Cooperative Research Centres (CRC) Program. The authors acknowledge the Australian Government Department of Health and Aged Care, including the Therapeutic Goods Administration, for project collaboration and support. The authors also acknowledge Leah Gerrard from the Australian Government Department of Health and Aged Care and the University of Technology Sydney for her contributions to data labelling and manuscript editing.

References

- [1] Ghazaleh Beigi and Huan Liu. 2020. A survey on privacy in social media: Identification, mitigation, and applications. *ACM Transactions on Data Science* 1, 1 (2020), 1–38.
- [2] Gephi 2024. *Gephi: The Open Graph Viz Platform*. Retrieved September 14, 2024 from <https://gephi.org/>
- [3] GPT-4o 2024. *GPT-4o*. Retrieved December 01, 2024 from <https://platform.openai.com/docs/models/gpt-4o>
- [4] Hootsuite 2024. *Hootsuite: Social Media Marketing and Management Tool*. Retrieved September 14, 2024 from <https://www.hootsuite.com/>
- [5] Isentia 2024. *Isentia: Media Monitoring, Intelligence & Insights Solution*. Retrieved September 14, 2024 from <https://www.isentia.com/>
- [6] Karket Research 2024. *Understanding the impact of medicine shortages in Australia*. Retrieved June 10, 2024 from <https://www.tga.gov.au/sites/default/files/2024-11/market-research-understanding-impact-medicine-shortages-australia.pdf>
- [7] Llama-3.1-8B-Instruct 2024. *Meta Llama 3.1 8B Instruct*. Retrieved July 22, 2024 from <https://huggingface.co/bartowski/Meta-Llama-3.1-8B-Instruct-GGUF>
- [8] Meltwater 2024. *Meltwater: Social Media Publishing Platform*. Retrieved September 14, 2024 from <https://www.meltwater.com/en/products/social-media-publishing>
- [9] Doerine J Postma, Magali LA Heijkoop, Peter AGM De Smet, Kim Notenboom, Hubert GM Leufkens, and Aukje K Mantel-Teeuwisse. 2024. Identifying Medicine Shortages With the Twitter Social Network: Retrospective Observational Study. *Journal of Medical Internet Research* 26 (2024), e51317.
- [10] Tom Simpson and Jerry Yik. 2024. Managing medicine shortages. *Australian Prescriber* 47, 5 (October 2024), 47:148–47:152. <https://doi.org/10.18773/austprescr.2024.045>
- [11] SocNetV 2024. *SocNetV: Social Network Visualizer*. Retrieved September 14, 2024 from <https://socnetv.org/>
- [12] Sproutsocial 2024. *Sprout Social: Social Media Management Solutions*. Retrieved September 14, 2024 from <https://sproutsocial.com/>