Original Paper

Topic and Sentiment Trends in Semaglutide Discussions on X: Subpopulation-Based Longitudinal Analysis

Parisa Momeni[1*], MSCS; Gabriel Laverghetta[1*], MSCS; Jay Ligatti[1], PhD; Lingyao Li[2], PhD

[1]Bellini College of Artificial Intelligence, Cybersecurity and Computing, University of South Florida, Tampa, FL, United States
[2]School of Information, University of South Florida, Tampa, FL, United States
[*]these authors contributed equally

Corresponding Author:
Parisa Momeni, MSCS
Bellini College of Artificial Intelligence, Cybersecurity and Computing
University of South Florida
4202 E. Fowler Avenue
Tampa, FL, 33620
United States
Phone: 1 8137090539
Email: parisamomeni@usf.edu

Abstract

Background: User experience has a significant impact on pharmaceutical drug effectiveness. Social media platforms like X (formerly Twitter) have become prominent spaces where individuals share their medication-related experiences, especially with widely marketed drugs such as semaglutide. Despite the large volume of conversation, a comprehensive understanding of how various user subpopulations engage with semaglutide-related discussions remains underdeveloped.

Objective: This study aims to explore how semaglutide is perceived and discussed across different X user groups. Within these user groups, we investigate (1) the evolution of sentiment patterns toward semaglutide and (2) the evolution and prevalence of semaglutide-related discussion topics.

Methods: We prepared a dataset consisting of 859,751 X posts (tweets) pertaining to semaglutide, along with related metadata, that were posted between July 2021 and April 2024. We apply sentiment analysis and topic modeling to the collected posts and analyze the sentiment patterns and topics within specific user subpopulations and time periods.

Results: Our analysis reveals a mean sentiment score of –0.24 (SD 0.669) across all posts, with all user subpopulations experiencing a decline in sentiment during the study period. User discussions focus on semaglutide's applications in weight loss and potential side effects, along with economic factors and celebrity/political influence. We also uncover differences in sentiment and discussion topics across user subpopulations. Notably, organizational accounts consistently express less negative sentiment (mean –0.04, SD 0.542) than individuals (mean –0.28, SD 0.605), with a statistically significant difference ($P<.001$), particularly in discussions related to drug efficacy and regulatory concerns. Interrupted time-series analysis shows a marked decrease in sentiment during the November 2022-January 2023 period, coinciding with regulatory announcements about potential adverse effects. In addition, we observe gender-based variations, such as a greater prevalence of discussions involving celebrities and politicians within female user posts (8368/39,786, 21%) compared to male user posts (8087/46,133, 17.5%), and male users expressing more positive sentiment.

Conclusions: This study helps advance the understanding of how diverse user groups perceive and discuss widely marketed drugs like semaglutide. Although we observe a general negativity, there are nuanced differences among the subpopulations. Our results offer valuable implications for health communication strategies and pharmacovigilance.

KEYWORDS

Introduction

Background

Semaglutide, also known by brand names such as Ozempic and Wegovy, has surged in popularity in recent years [1]. Originally developed as a diabetes medication, semaglutide has recently shown effectiveness in off-label use for weight-loss treatment [2]. Semaglutide was the fourth-highest drug expenditure in the United States in 2021, with US $10.8 billion spent on the drug [3]. Social media platforms like X (formerly Twitter) have become key venues for the public to share experiences and express opinions about semaglutide [4]. Celebrities or public figures, such as Elon Musk, have shared personal weight-loss stories and endorsed the drug, further amplifying conversations [1,5]. The widespread advertising, endorsements by high-profile figures, and increased consumer interest have made semaglutide a trending topic in medications [6]. As such, social media offers a unique lens to examine how it is perceived and discussed, shedding light on public sentiment, misconceptions, and concerns [4,7].

Understanding public perceptions is essential, as user experiences significantly influence the evaluation of pharmaceutical drugs' effectiveness [8,9]. Positive experiences not only enhance user satisfaction but also contribute to improved adherence and overall well-being [10]. Mining social media data allows policymakers and pharmaceutical providers to tap into a vast repository of real-time data pertaining to user experiences [11]. In addition, by analyzing user-generated content, researchers can uncover nuanced insights into the concerns, preferences, and challenges faced by specific subpopulations by gender, location, or other demographic attributes. This granular analysis provides an opportunity to identify unmet needs, tailor interventions, and ensure more equitable health care outcomes.

In addition to obtaining information from social media, recent research has focused on applying natural language processing (NLP) techniques due to their ability to quickly process large-scale and unstructured information [12,13]. NLP has been leveraged to extract chemical-disease relations [14], build health knowledge graphs [15], and ease the process of documentation in electronic health records [16] (which often use unstructured and nonstandardized formats) [17]. Text analytics approaches such as named entity recognition, topic modeling, and sentiment analysis have been applied within the health context [18-20].

Although mining drug-related user experiences on social media has been widely explored [21,22], few studies have focused specifically on semaglutide-related discourse. Our study combines large-scale sentiment analysis and topic modeling with user subgroup analysis, offering a granular view of public engagement with semaglutide. Prior work has primarily leveraged social media to identify adverse reactions to semaglutide that were not detected during clinical trials [7,23]. The study by Alvarez-Mon [4] includes a manual analysis of 2045 posts to determine user interests, beliefs, and experiences pertaining to semaglutide and other antiobesity drugs. However, the public discourse including the sentiments and prevalent topics within specific user groups has been underexplored. To uncover patterns in how different user subpopulations experience and discuss semaglutide on X, we investigate the following research questions (RQs).

- RQ1 (Sentiment Analysis): What underlying factors explain how sentiment toward semaglutide evolves over time and across different user subpopulations, and what insights can be drawn about public concerns and motivations?
- RQ2 (Topic Modeling): Which topics of discussion are most prevalent in positive and negative semaglutide-related posts across different user subpopulations, how do their prevalence patterns change over time, and what do these patterns reveal about group-specific attitudes, priorities, and health communication needs?

The first research question aims to explore differences in engagement patterns and sentiment expressions across user subpopulations and over time. The second question aims to identify the various discussion topics emphasized by distinct user subpopulations, exploring the prevalence of these topics among the subpopulations. Addressing these two research questions, our study provides a comprehensive discourse analysis across user subgroups, along with an identification of external events that influence the evolution of these patterns. This insight into the real-world user experience is crucial for tailoring public health communication and improving medicine support strategies for diverse communities.

Related Work

Exploring User Experiences via Social Media

Crowdsourcing, originally defined as the act of a company or institution taking a function once performed by employees and outsourcing it to an undefined (and generally large) network of people in the form of an open call, has revolutionized how researchers gather and analyze public opinion [24]. Similar methodological approaches have been applied in other domains to understand public perceptions of urban accessibility and inclusion through crowdsourced online reviews [25]. While conventional methods like polls and surveys remain valuable, crowdsourcing through social media platforms enables researchers to collect and analyze large-scale, near-real-time data about user experiences and perspectives [26].

In the health care domain, crowdsourcing via social media has become particularly valuable for understanding public opinions about medical treatments and pharmaceutical drugs [27,28]. First, social media data have shown invaluable potential in pharmacovigilance due to the tendency of users to share their opinions or experiences such as adverse drug reactions [29]. Researchers have leveraged social media data to track topic trends [30], estimate disease prevalence [31], and analyze public response to health policies [32]. Second, crowdsourcing has proven effectiveness in capturing user experiences that might not be readily available through traditional clinical studies or surveys [33]. In particular, social media platforms provide researchers with access to diverse user populations and their real-world experiences with pharmaceutical drugs, such as off-label use [9] and adverse reactions [34].

Analyzing User Experiences via NLP Techniques

Our work makes use of sentiment analysis and topic modeling to uncover patterns in user experiences with semaglutide. Sentiment analysis, also called opinion mining, is a branch of NLP that focuses on classifying people's opinions into positive, negative, or neutral associated with data [35]. Research in this domain spans various levels of granularity, from assigning a single sentiment to an entire document or individual sentences to analyzing distinct aspects linked to specific entities [36]. After the COVID-19 pandemic, there has been increasing interest in using sentiment analysis to evaluate the attitudes, perceptions, and emotions expressed by social media users [37-39]. Numerous studies have focused on platforms such as X, Reddit, and Facebook, which have become prominent spaces for sharing public opinions related to COVID-19 [40].

Topic modeling techniques, such as latent Dirichlet allocation [41] and BERTopic [42], seek to discover the key themes present in a corpus of documents [43]. Topic models can summarize large datasets by capturing the topics (ie, the major discourse) that appear most commonly in the text. Numerous works have used topic models to study health-related discussions on social media. For example, Asghari [44] identifies trending topics pertaining to health care on X. Topic analysis has provided insights into news reports surrounding COVID-19 [45] and public opinion regarding blood donation [46]. Another study trains an aspect-based topic model to characterize the health topics and then estimates the prevalence of influenza and allergies over time by observing the number of mentions of each topic during different time periods [47]. In addition, prior work has analyzed semaglutide-related discussions on Reddit via topic modeling [48-50]. Our study focuses on exploring how public discourse and sentiment differ across user subpopulations. Analyzing specific subpopulations enables a nuanced understanding of user concerns, such as accessibility, side effects, and insurance coverage, guiding targeted strategies for addressing subpopulation-specific needs.

Methods

Overview

Figure 1 provides an overview of our research design. We compile a dataset consisting of semaglutide-related X posts (stage 1) and perform sentiment analysis and topic modeling on these posts (stage 2). We analyze the results to address our two research questions (stage 3).

Figure 1. Overview of the designed framework to implement the research. BERT: bidirectional encoder representations from transformers; c-TF-IDF: class-based term frequency-inverse document frequency; ITS: interrupted time series; RQ: research question; UMAP: uniform manifold approximation and projection.
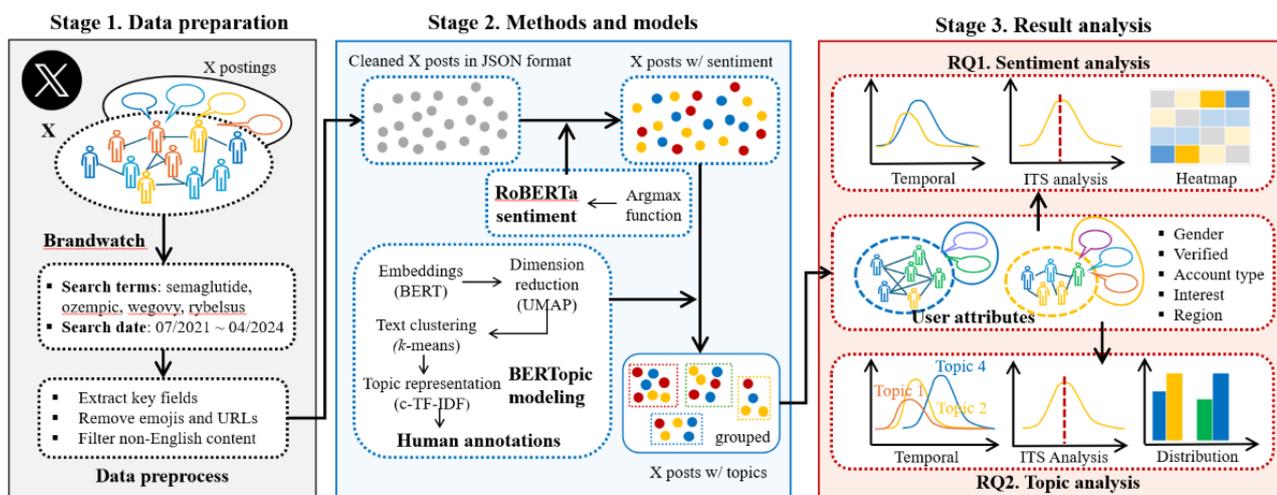
Data Preparation and User Attributes

We use Brandwatch [51], a social media analytics platform, to collect data from X. Brandwatch uses the X application programming interface to obtain posts from prior periods, offering a representative sample of X's entire dataset. Our data collection targets X posts posted between July 1, 2021—when the US Food and Drug Administration (FDA) approved semaglutide for chronic weight management [52]—and April 30, 2024. This nearly 3-year timeframe enables our longitudinal discourse analysis surrounding semaglutide within different user communities.

To collect the data, we use several key search terms, including "semaglutide" and its branded names "Ozempic," "Wegovy," and "Rybelsus." We chose these search terms to ensure comprehensive coverage of discussions related to semaglutide and its marketed variants. To maintain consistency, we limit the dataset to English-language X posts containing these terms. The final dataset consists of 859,751 posts, including original posts, replies, reposts (retweets), and quotes.

The user attributes we study are gender, US region, interests, account type, and verification status. These attributes are inferred by the data provider, Brandwatch. According to their documentation, the methodology for this inference varies by attribute; for instance, location is primarily determined from explicit, user-provided information in their X public profiles, whereas attributes like gender and interests are classified using machine learning models that analyze public data such as first names and biographical text [53]. We acknowledge that the specific algorithms used for this demographic inference are proprietary to Brandwatch, and as such, the details of their methods are not publicly available [53].

We choose these attributes for their relevance in capturing diverse user perspectives and behaviors. For example, gender and region can reveal variations in health care access and cultural attitudes [54,55] and have been identified as factors associated with semaglutide initiation [56]. Account type and verification status help differentiate individual users from organizations. Interests reflect personal activities. It should be noted that users may have multiple interests or none, placing them in zero or more interest-based subpopulations.

Since individual users may post multiple posts, we measure the size of each subpopulation by the number of users and by the number of posts. When grouping each post by its respective user, we assign aggregated values for attributes (gender, region, interests, account type, and verification) using the most frequently observed values across that user's posts. This aggregation ensures that each user is represented by a single record while preserving their attributes. The dataset contains 436,551 unique users. Tables S1 and S2 in Multimedia Appendix 1 summarize the user subpopulations.

Ethical Considerations

The University of South Florida Institutional Review Board reviewed this study and determined it to be exempt (STUDY009222). To remain compliant with X's policies, we adhere to all application programming interface rate limits and do not collect deleted posts. All the code used to generate our results is available from our GitHub repository [57]. To protect user privacy, our full dataset containing identifiers such as usernames is not available for public access [58]. An anonymized version can be accessed from Multimedia Appendix 1 or requested from the corresponding author. Lastly, we follow standard guidelines [59] to mitigate potential harms from sensitive content contained in our dataset (eg, suicide mentions).

RoBERTa Sentiment Analysis

To classify the sentiment of X posts, we use the *cardiffnlp/twitter-roberta-base-sentiment-latest* model from Hugging Face [60], which is pretrained on X data and widely recognized for its state-of-the-art performance in sentiment classification tasks [61], particularly in handling X posts [62]. We apply the RoBERTa (Robustly Optimized BERT Pretraining Approach) model to each post, obtaining a set of values representing the likelihood of the post having negative, neutral, or positive sentiment. The model outputs these values in a list format, with the elements at indices 0, 1, and 2 corresponding to the negative, neutral, and positive probabilities, respectively [60]. The final sentiment label is determined by the *argmax* function; no additional calibration or thresholding is performed on the sentiment scores. After that, we assign each post a sentiment label, that is, –1 for negative, 0 for neutral, and 1 for positive. As a result, of the 859,751 total posts, 116,091, 429,074, and 314,586 are classified as positive, neutral, and negative sentiment, respectively.

We calculate mean sentiment scores within each of the user subpopulations (gender, US region, account type, verification status, and interests). To measure sentiment per user, we first group each post by its respective user. We then calculate the mean sentiment score for each user by averaging their per-post sentiment scores, resulting in a continuous value between –1 and 1. This grouping allows us to analyze the aggregated averages of user-level sentiments across the subpopulations. Since the volume of posts varies by subpopulation (eg, individual user accounts create 1.9 posts on average, compared to 4.4 from organizational users), we also calculate average sentiment per post. In addition, to assess the robustness of our sentiment findings, we conduct a repost-excluding sensitivity analysis, in which we measure average sentiment within each subpopulation while excluding reposts (resulting in 411,747 posts for sensitivity check).

Interrupted Time-Series Regression

To examine how sentiment evolves over time, we group the dataset and calculate the average sentiment bimonthly spanning from July 2021 to April 2024. This procedure yields a time-series of average sentiment scores. Our longitudinal analysis is performed per post, as a user can publish multiple posts at different times. To determine whether the observed sentiment shifts exceed baseline trends, we apply an interrupted time-series (ITS) regression. Using bimonthly sentiment averages, we model (1) the baseline time trend, (2) an immediate level change at the intervention point, and (3) a slope change after the event. This approach allows us to separate long-term temporal patterns from abrupt discontinuities. Given the large dataset and multiple comparisons, we treat results as exploratory and emphasize effect sizes and CIs over strict hypothesis testing. Our ITS regression is modeled as follows:

$$Y_t = \beta_0 + \beta_1 Time_t + \beta_2 Event_t + \beta_3(Time_t Event_t) + \epsilon$$

where Y_t represents the mean sentiment score during the bimonthly period, $Time_t$ is the continuous time index, $Event_t$ is a binary indicator coded 0 before the intervention and 1 afterward, and $Time_t Event_t$ captures the postevent slope. Here, β_0 estimates the baseline level at the beginning of the series, β_1 captures the pre-event trend, β_2 reflects the immediate level change at the intervention, and β_3 represents the postevent slope change.

BERTopic Modeling

After performing sentiment analysis, we use the BERTopic model [42] to discover the commonly discussed topics in the dataset. We divide the dataset into posts with a positive RoBERTa sentiment label and posts with a negative sentiment label; posts with neutral sentiment are excluded to focus on identifying the topics that contribute to positivity and negativity. We perform topic modeling for the positive and negative tweets separately. To create the positive and negative document corpora from our dataset, we first clean the text of each post. This cleaning process involves steps such as removing emojis, punctuation, and stop words, normalizing whitespace, and converting all text to lowercase. Cleaning the text is an important step due to the unstructured nature of social media posts [63]. Note that these cleaning steps are not performed prior to obtaining the sentiment of each post, as they may have affected the sentiment results. For example, emojis [64] and punctuation [65] can impact sentiment scores. Our topic modeling exercise

focuses solely on the themes present in the text, as opposed to the sentiment of the text. After cleaning each post's text, we remove duplicated reposts to avoid situations in which distinct reposts of a given post are mapped to different topics. The positive and negative document corpora are lists consisting of the snippets of each of the positive and negative cleaned X posts, respectively.

After performing the cleaning steps, we initialize a BERTopic model with default values for all hyperparameters. We run the BERTopic model on our document corpora to generate a list of topics and then extract the sentence/document embeddings for those topics. We use *k*-means clustering [66] with the Elbow method to determine the optimal number of topics. Figures S1 and S2 in Multimedia Appendix 1 show the *k*-means clustering results for the positive and negative document corpora, respectively. We then run the BERTopic model on our document corpora using 100 clusters.

BERTopic outputs a topic representation and a document representation. In our case, the documents are cleaned X posts. For each topic, the topic representation lists its representative keywords, representative documents, and document count. The BERTopic document representation maps each document in the corpus to its topic number. After reviewing the 100 positive and 100 negative topics produced by BERTopic, we observe that many of the topics shared similar topics. We therefore manually annotate the 200 topics into 10 umbrella topics. These topic groupings consist of all the topics that share a common theme. For instance, umbrella topic 0 consists of topics relevant to weight loss.

We perform the manual annotation by reviewing the representative keywords and documents. For example, one of the most common positive topics is represented by the following list of keywords: [*semaglutide*, *semaglutides*, *weightloss*, *diet*, *appetite*, *medication*, *eat*, *treatment*, *craving*, *fda*]. A representative document (X post after cleaning) for this topic is: "ready lose weight gain confidence say hello semaglutide gamechanging prescription medication help achieve significant weight loss with semaglutide take control craving appetite finally reach weight loss goal." These keywords and the document are associated with themes of weight loss, so we map this topic to umbrella topic 0 (weight loss). As another example,

one of the most common negative topics is represented by these keywords: [*nausea*, *diarrhea*, *vomiting*, *nauseous*, *constipation*, *vomit*, *nauseate*, *constipate*, *diarrhoea*, *stomach*]. One of the representative documents for this topic is: "nausea diarrhea stomach abdominal pain vomiting constipation side affect ozempic." This topic pertains to adverse reactions experienced after taking semaglutide; therefore, we map it to umbrella topic 8 (acute harm/adverse drug reactions). The mapping of all 200 initial topics to the 10 umbrella topics is available in Multimedia Appendix 2. Additionally, to verify that the umbrella topic shares do not change significantly under different cluster counts, we conduct stability checks using 50 and 25 clusters. The results are described in the Umbrella Topic Stability Checks section in Multimedia Appendix 1.

Two of the authors (PM and GL) independently map the initial 200 topics to the 10 umbrella topics. We calculate the intercoder agreement between the annotators using Krippendorff α [67], as described by the following equation, where D_o is the observed disagreement between the annotators and D_e is the expected random disagreement:

$$\alpha = 1 - \frac{D_o}{D_e}$$

The intercoder agreement is 0.806, indicating a satisfactory level of agreement between the annotators. Each X post is mapped to one of the clustered 200 topics using BERTopic, and each of the clustered 200 topics is mapped to one of the 10 umbrella topics based on manual annotation. We can therefore map each post to its umbrella topic. Table 1 shows the 10 umbrella topics, examples of their representative keywords, and the number of posts mapped to each topic. In addition, Table S3 in Multimedia Appendix 1 provides example representative documents (posts) for the 10 umbrella topics. During the annotation, the topics are mapped to umbrella topic T9 ("other") if the topic does not clearly match any of the other 9 umbrella topics. We observe various subtopics that appear within T9, such as drug marketing/news, other nonsemaglutide drugs, and health/beauty. With the labeled umbrella topics, we group the data by user attributes to discover the prevalence of each umbrella topic among user subpopulations.

Table 1. Number, name, and representative keywords for each of the 10 umbrella topics.

Topic no.	Topic name	Posts, n	Example representative keywords
T0	Weight loss	50,542	exercise, workout, eat, diet, appetite, skinny, obesity, craving
T1	Celebrities/politicians	40,120	nikkifried, erikajayne, oliviawilde, rhianna, tuckercarlson, oprah, elonmusk, trump
T2	Obtaining the drug	22,724	prescription, medication, walgreens, walmart, coupon, insurance, coverage, supply, ordered, appointment, affordable, shot, injection
T3	Drug indicators	16,948	diabetes, inflammation, treatment
T4	Drug authorities	23,783	pharma, novo, nordisk, doctor, physician, fda, goldman, economy, gdp, market
T5	General and profane negativity	9488	[swear words]
T6	Suicide risk	1486	suicide, autopsy, death, overdose
T7	Chronic harm	3075	addiction, cancer, tumor, aloceptia, hair, hairline, dialysis
T8	Acute harm/adverse drug reactions	10,208	nausea, diarrhea, constipation, pain, effect
T9	Other	29,475	[anything that does not fit in with the other topics]

Results

RQ1 (Sentiment Across User Subpopulations)

Addressing RQ1, we find that the overall sentiment toward semaglutide during the study period (July 2021 to April 2024) is slightly negative, with a mean sentiment score of –0.24 (SD 0.669) across all posts and –0.28 (SD 0.605) across all users. Within all user subpopulations, we observe a decline in sentiment over time, but the trend varies across different subpopulations, as discussed in the following subsections.

Overall Sentiment Declines

Our longitudinal sentiment analysis results are displayed in Figure 2. To illustrate uncertainty, we include 95% CIs for key categorical attributes as shaded bands. The temporal progression of sentiment can be divided into four phases: (1) initial positive sentiment across categories (2021 to mid-2022), (2) universal decline (November 2022 to January 2023), (3) variable recovery rates through 2023, and (4) eventual stabilization at slightly negative levels by early 2024. Three notable periods of universal decline are present: the largest from November 2022 to January 2023, the second from September to November 2023, and the final from January to March 2024. Our repost-excluding sensitivity analysis yields similar trends, supporting the robustness of these temporal trends.

Figure 2. Time series plots showing bimonthly sentiment analysis categorized by gender (A), verified status (B), account type (C), interest (D), and region (E). To improve readability, we limit the interest visualization to the top 5 most popular of the 21 user interests. Shaded areas represent 95% CIs.
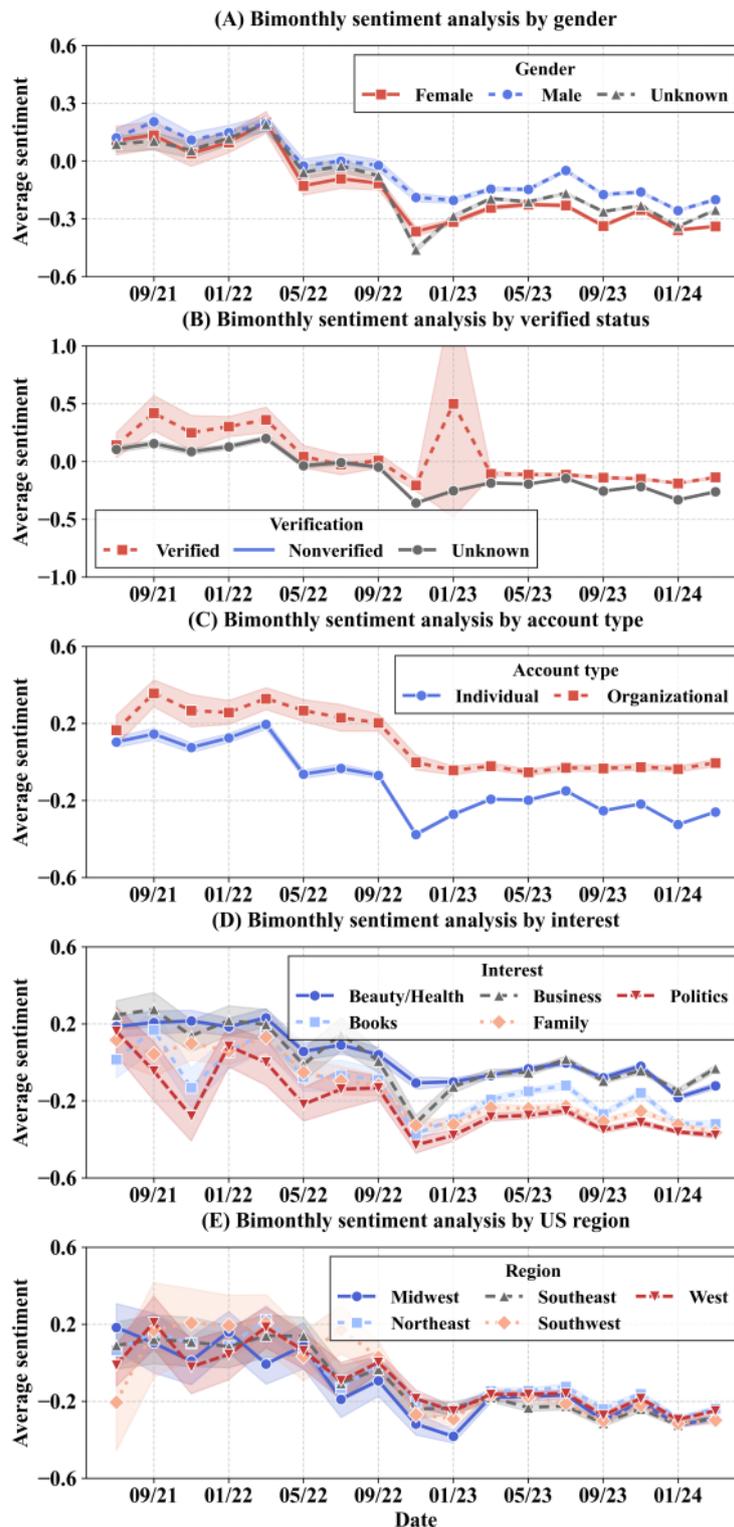

Connection With External Events

The sentiment changes are associated with external events (some key events are highlighted in Figure 3). The first clear sentiment decline starting in mid-2022 coincides with the national shortage of glucagon-like peptide 1 medications, which led to increased reliance on compounded semaglutide formulations and raised concerns about access, safety, and regulatory oversight [68].

During this same period, the FDA approved Wegovy for adolescent patients, intensifying public debate around expanded clinical use [69]. Moreover, there is a pronounced decline in sentiment within all user subpopulations during November 2022–January 2023, which coincides with reports of adverse gastrointestinal reaction, suggesting a public concern about potential side effects [70]. Following this decline, a temporary

spike in positive sentiment, particularly among verified users, emerges between January 2023 and March 2023. However, the wide CI for verified users indicates a high standard error likely resulting from the limited number of verified-user posts within that bimonthly window. This increase in sentiment coincides with heightened media attention and celebrity endorsements of semaglutide as a "Hollywood weight loss drug" [71], suggesting that promotional activity and influencer-driven narratives could temporarily reverse prevailing sentiment patterns.

Figure 3. Interrupted time series analysis of average sentiment over time with external events. The dashed red line indicates the start of the intervention window. CROI: Conference on Retroviruses and Opportunistic Infections; CV: cardiovascular; FDA: US Food and Drug Administration; SELECT: Semaglutide Effects on Cardiovascular Outcomes in People With Overweight or Obesity.
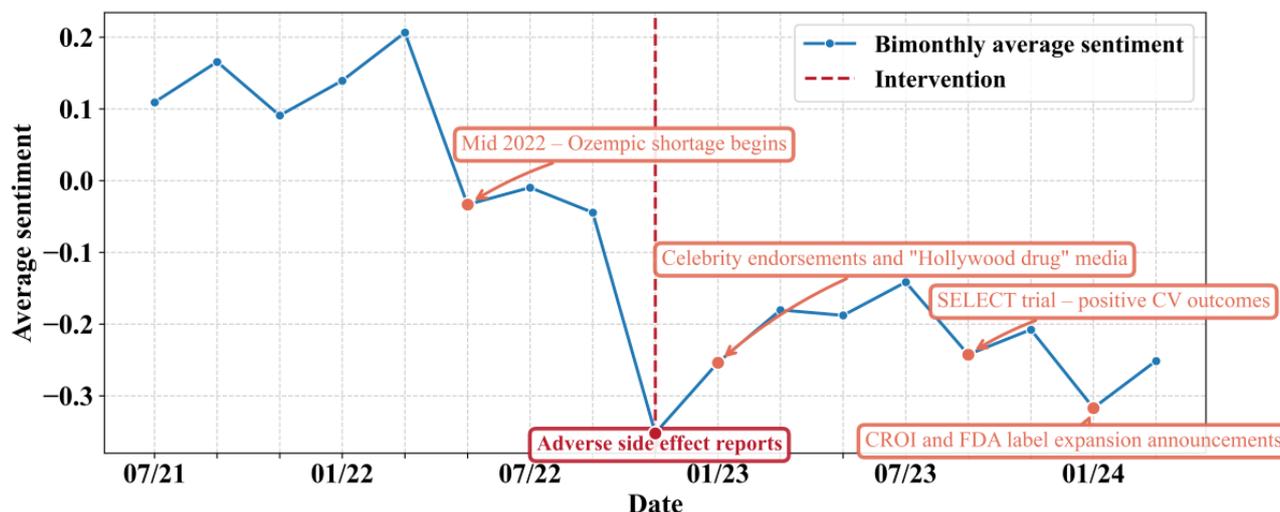
Following the second dip between September and November 2023, sentiment increases across all user attributes. This recovery aligns with the release of positive cardiovascular outcomes from the SELECT (Semaglutide Effects on Cardiovascular Outcomes in People With Overweight or Obesity) trial, which was presented at the American Heart Association Scientific Sessions and published in the *New England Journal of Medicine* in November 2023 [72]. The results show that Wegovy (semaglutide 2.4 mg) can significantly reduce major cardiovascular events in adults with overweight or obesity and established cardiovascular disease, potentially improving public perceptions of the drug's efficacy and safety. Following the third dip in sentiment between January and March 2024, we observe a rebound across nearly all subpopulations. This recovery coincides with two major announcements in early March 2024: (1) a National Institutes of Health–sponsored study presented at the 2024 Conference on Retroviruses and Opportunistic Infections showed semaglutide significantly reduces liver fat in people with HIV and MASLD (metabolic dysfunction-associated steatotic liver disease) [73] and (2) the FDA approved a label expansion for Wegovy to include cardiovascular risk reduction based on long-term SELECT trial data [74]. These developments likely contributed to renewed optimism about semaglutide's broader therapeutic value.

However, the sentiment patterns among user groups vary. For example, we observe that organizational accounts and verified users have consistently more positive sentiment scores than individual accounts and unverified users. Users who are interested in "business" exhibit moderate sentiment scores, reflecting the professional nature of corporate communications. Users interested in "politics" exhibit the most negative and volatile sentiment, whereas the "beauty/health" category maintains the most positive sentiment among all groups throughout the timeline. Users with an interest in "books" maintained the most stable sentiment pattern. These patterns highlight how different user communities process health-related information through their respective contextual frameworks, with users valuing family and personal health showing the highest sentiment variation.

The Impact of Semaglutide Shortage and Side Effect Reports

In addition, we observe a significant downturn in sentiment across all user subpopulations during November 2022–January 2023. This universal decline aligns temporally with reports of semaglutide shortages [75-77] and reports of adverse drug reactions [70], suggesting a public concern regarding the drug's availability and potential side effects. We apply an ITS regression to assess the impact of these shortages and determine whether the sharp decline in sentiment exceeds baseline trends, as displayed in Table 2. The baseline sentiment is significantly positive (β_0=0.172, P=.003), with a modest but significant negative trend prior to the event (β_1=−0.027, P=.03). At the intervention point, an immediate and statistically significant drop occurs (β_2=−0.430, P=.004). While the postintervention slope shows a slight positive trend (β_3=0.029, P=.07), this effect does not reach conventional significance thresholds.

Table 2. Interrupted time-series regression of average sentiment. The intervention point is set at November 2022–January 2023; post time indicates the slope change thereafter.

Variable	β coefficient (95% CI)	SE	P value
Constant	0.172 (0.070 to 0.274)	0.047	.003
Time	–0.027 (–0.051 to –0.003)	0.011	.03
Event	–0.430 (–0.700 to 0.160)	0.125	.004
Post time	0.029 (–0.00 to 0.06)	0.015	.07

Figure 3 visualizes the ITS regression, displaying observed mean sentiment across bimonthly intervals. A vertical dashed line marks the November 2022 intervention, illustrating changes in level and slope. As shown in Figure 3, sentiment remains relatively stable and slightly positive throughout 2021 and most of 2022. A clear decline emerges in the November 2022-January 2023 period, after which sentiment consistently remains below zero, indicating a sustained downturn in public discourse. The ITS results suggest that the intervention period coincides with a significant immediate downturn in sentiment, followed by a slight, nonsignificant recovery trend thereafter.

Sentiment Patterns Across Subpopulations

We observe that the per-user and per-post sentiment scores are similar; therefore, to avoid duplicated explanation and focus on user experience, we solely discuss per-user scores in this section (per-post scores and repost-excluded scores are available in Table S5 in Multimedia Appendix 1). The sentiment distribution among user subpopulations varies significantly, as illustrated in Figure 4. In addition, Table 3 presents 95% CIs for the estimated mean sentiment differences across key user groups. All comparisons yield statistically significant results (*P*<.001), providing strong evidence that the observed differences are not due to random variation.

Figure 4. Sentiment patterns across user subpopulations, measured per user and per post (the sentiment of a post can only be the value from –1, 0, and 1).
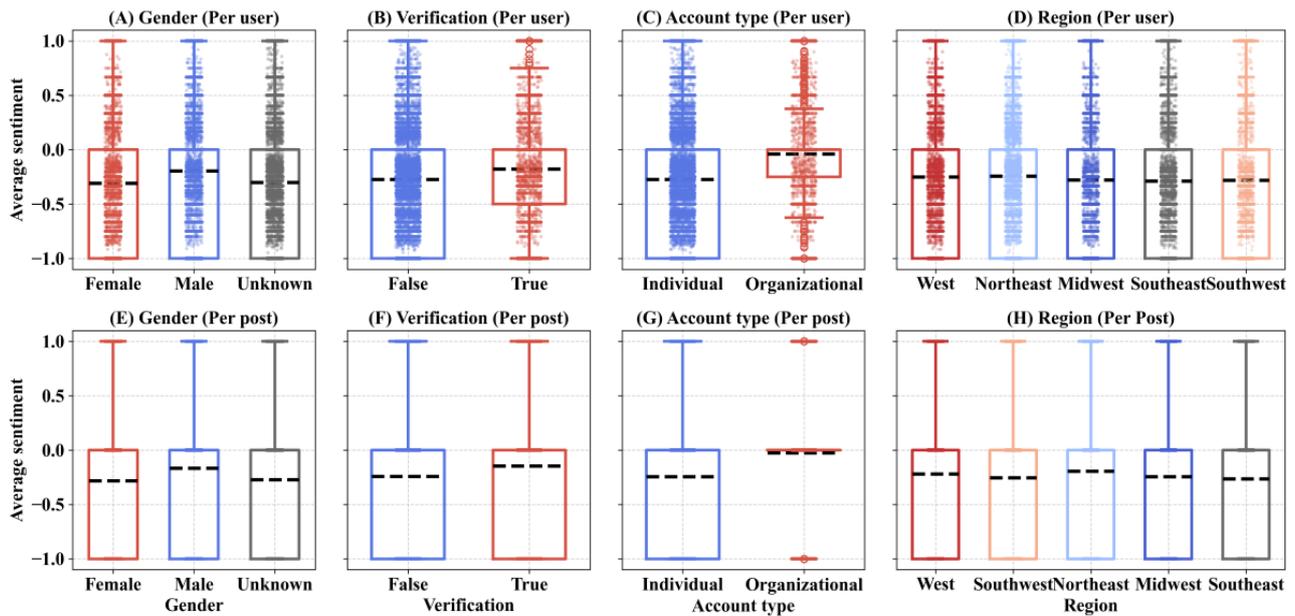
Table 3. Estimated differences in sentiment across key subgroups, with 95% CIs.

Comparison and levels	Mean 1	Mean 2	Mean difference (95% CI)	P value
Male vs female				
Per post	–0.17	–0.29	0.12 (0.121-0.129)	<.001
Per user	–0.20	–0.32	0.12 (0.117-0.128)	<.001
Verified vs nonverified				
Per post	–0.14	–0.24	0.10 (0.092-0.102)	<.001
Per user	–0.17	–0.28	0.11 (0.101-0.116)	<.001
Organizational vs individual				
Per post	–0.01	–0.24	0.23 (0.222-0.235)	<.001
Per user	–0.04	–0.28	0.24 (0.232-0.256)	<.001

Based on Figure 4, male users exhibit more positive sentiment (averaged sentiment –0.20) than female users (averaged sentiment –0.32). Verified users, typically public figures or organizations with confirmed identities, expressed less negative sentiment (averaged sentiment –0.17) toward semaglutide than nonverified users (averaged sentiment –0.28). This contrast highlights how identity and accountability influence sentiment expression online. Organizational accounts expressed less negative sentiment compared to individual users. These findings suggest that organizations tend to frame their discussions about semaglutide in a more positive manner, possibly due to pharmaceutical marketing, professional communication standards, or endorsement practices. Within the United States, regional variations in sentiment toward semaglutide are evident in the analysis. Users in the Southeast region express the most negativity (averaged sentiment –0.29), while users from the Northeast are less negative.

Lastly, Figure 5 presents the average sentiment scores for each combination of account type and user interest, measured per user and per post. Overall, individual users tend to express more negative sentiment than organizational accounts. For example, in the "travel" category, individual sentiment is clearly negative (–0.24), while organizational sentiment is positive (0.12), yielding one of the largest absolute gaps between the two account types. This figure highlights that user sentiment varies widely across domains of interest. Health-related categories, particularly those tied to "beauty/health and fitness," generate a higher level of positivity, while "business" discussions are closer to neutral, reflecting a more corporate and less personal orientation. These findings underscore the importance of considering both account type and user interest when analyzing public sentiment.

Figure 5. Combined heatmap of mean sentiment scores by account type and interest.
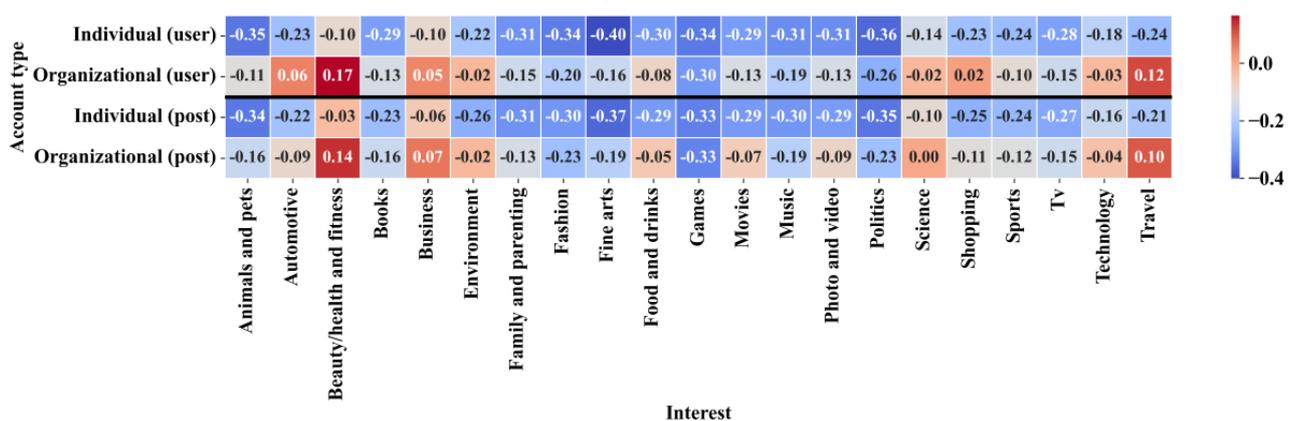

RQ2 (Topic Results Across User Subpopulations)

For the topic analysis, we focus on original posts by successfully assigning umbrella topics to 207,849 out of 859,751 posts. Not all posts in the dataset receive topic assignments because we exclude duplicated reposts (to avoid counting the same content multiple times) and neutral-sentiment posts from the topic modeling process. We conduct topic modeling analysis by post rather than aggregating by user, which preserves the complete topic distribution for prolific users.

Topic Prevalence Across Subpopulations

Figure 6 shows the prevalence of each topic within each user subpopulation, highlighting the variations in each topic's prevalence across subpopulations. The most common topic among subpopulations was T0 (weight loss). However, the popularity of the other topics is less consistent across the subpopulations. We present a descriptive analysis of the topic prevalence results in the following paragraphs; CIs and effect sizes (calculated via Cramér V [78]) for these results are given in Table S7 in Multimedia Appendix 1.

Figure 6. Prevalence of each topic within each user subpopulation. The bars and left axes measure the number of posts pertaining to each topic, and the lines and right axes measure the percentage of each subpopulation's posts that pertain to each topic.
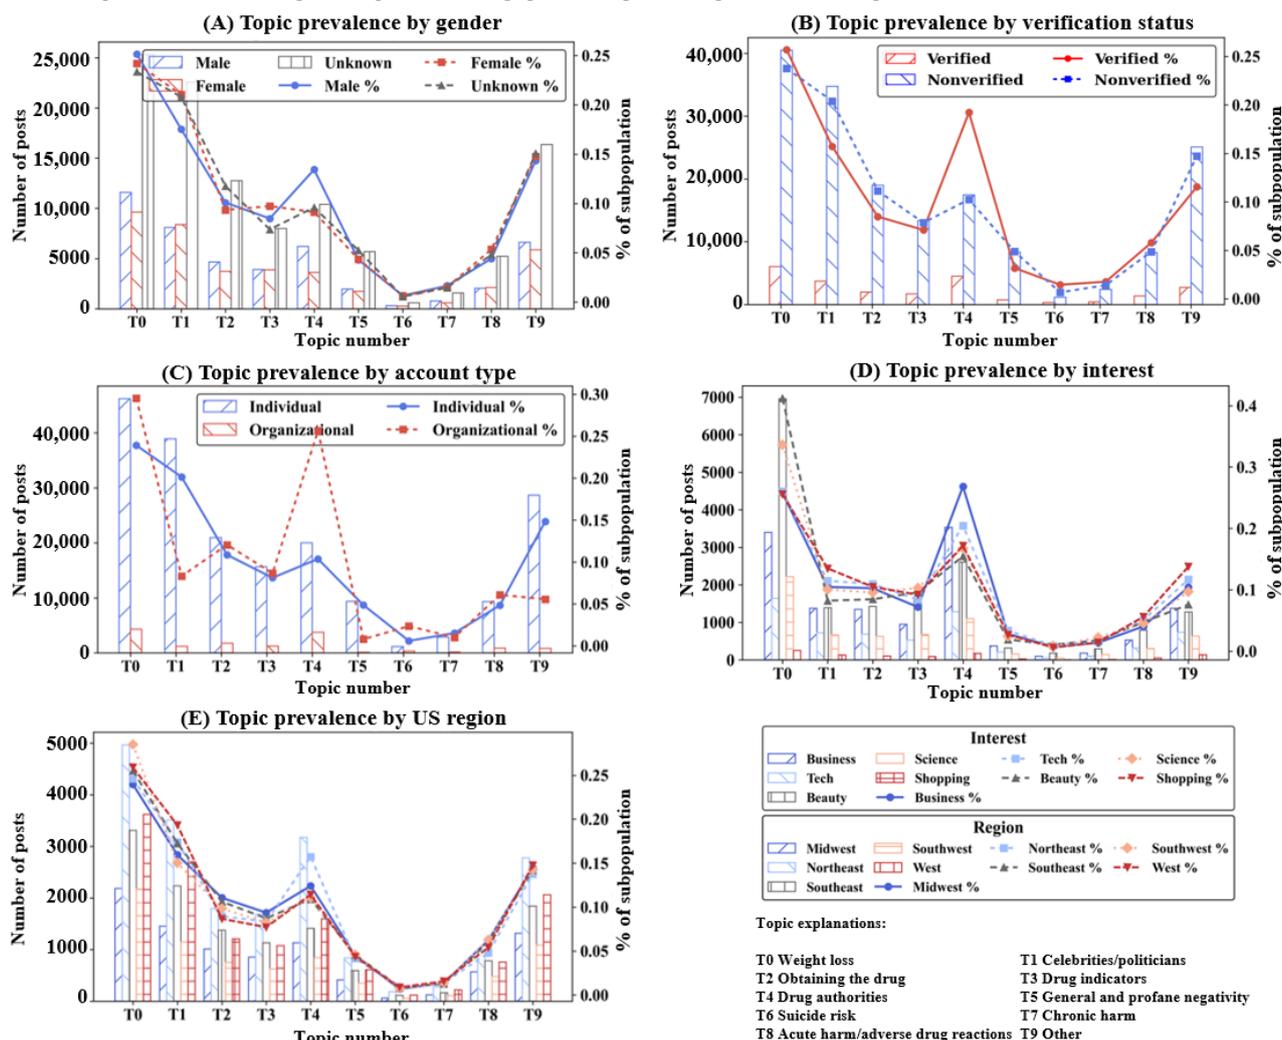
T1 (celebrities/politicians) is noticeably more popular among female users compared to male users, comprising 21% (8368/39,786) of posts from female users and 17.5% (8087/46,133) of posts from male users. Although there are more overall posts originating from male users, female users post more T1 posts. In addition, 13.4% (6200/46,133) of male user posts pertained to T4 (drug authorities), compared to 9.1% (3630/39,786) of posts from female users.

Verified users are less likely to post profane posts, with T5 (general and profane negativity) comprising 3.2% (748/23,549) of their posts, compared to 4.9% (8318/170,279) of posts from unverified users. However, verified users are about twice as likely to post about T6 (suicide risk) than unverified users. They are also more likely to create posts pertaining to T4 (drug authorities); 19.2% (4525/23,549) of verified user posts belong to T4, compared to 10.3% (17,460/170,279) from unverified users.

Examining the most prevalent topics among individual accounts and organizational accounts, the most striking difference is the very low number of profane posts within the organizational account subpopulation. T5 comprises just 0.8% (117/14,711) of organizational posts, compared to 4.9% (9371/193,138) of posts from individual users. As companies and organizations likely do not want to damage their reputation by posting profane content, this result is in line with our expectations. On the other hand, organizations are about 4 times more likely than individuals to create posts pertaining to T6 (suicide risk), perhaps due to medical organizations posting warnings about potential side effects of semaglutide. T4 is far more common among organizational accounts, with 25.7% (3774/14,711) of organizational posts belonging to T4 compared to 10.4% (20,009/193,138) from individual users.

Dividing the users by interest reveals several differences in topics of discussion. Notably, the subpopulation consisting of users interested in "business" is the only subpopulation in which T0 is not the most prevalent topic. T4 is the most common topic among users interested in "business." Users in the "business" subpopulation appear to be more interested in the economic impact of semaglutide, as opposed to its usage in weight loss treatment. As expected, T0 is by far the most popular topic among users interested in "beauty/health." Lastly, the prevalence of each topic is mostly consistent across different US geographic regions. However, there are some variations; for example, T4 is noticeably more popular in the Northeast compared to other regions.

Topic Prevalence Over Time

Figure 7 shows the number of posts pertaining to each umbrella topic posted during each bimonthly period from July 2021 to April 2024. To assess the evolution of topic prevalence over time, we first present an exploratory analysis of external events that may have influenced the topic trends. All topics rose in popularity from July 2021 to April 2024. This result is consistent with the general increase in popularity of semaglutide. T3 surged in popularity during November 2022-January 2023. This trend aligns with the initial FDA approval of Wegovy for adolescents [69], which occurred on December 23, 2022. During September 2023-November 2023, T0 (weight loss) increases in prevalence, while T8 (acute harm/adverse drug reactions) and T6 (suicide risk) decline. These changes align with the release of positive cardiovascular outcomes for Wegovy from the SELECT trial [72].

Figure 7. The number of posts pertaining to each topic over time. T0: weight loss; T1: celebrities/politicians; T2: obtaining the drug; T3: drug indicators; T4: drug authorities; T5: general and profane negativity; T6: suicide risk; T7: chronic harm; T8: acute harm/adverse drug reactions; T9: other.
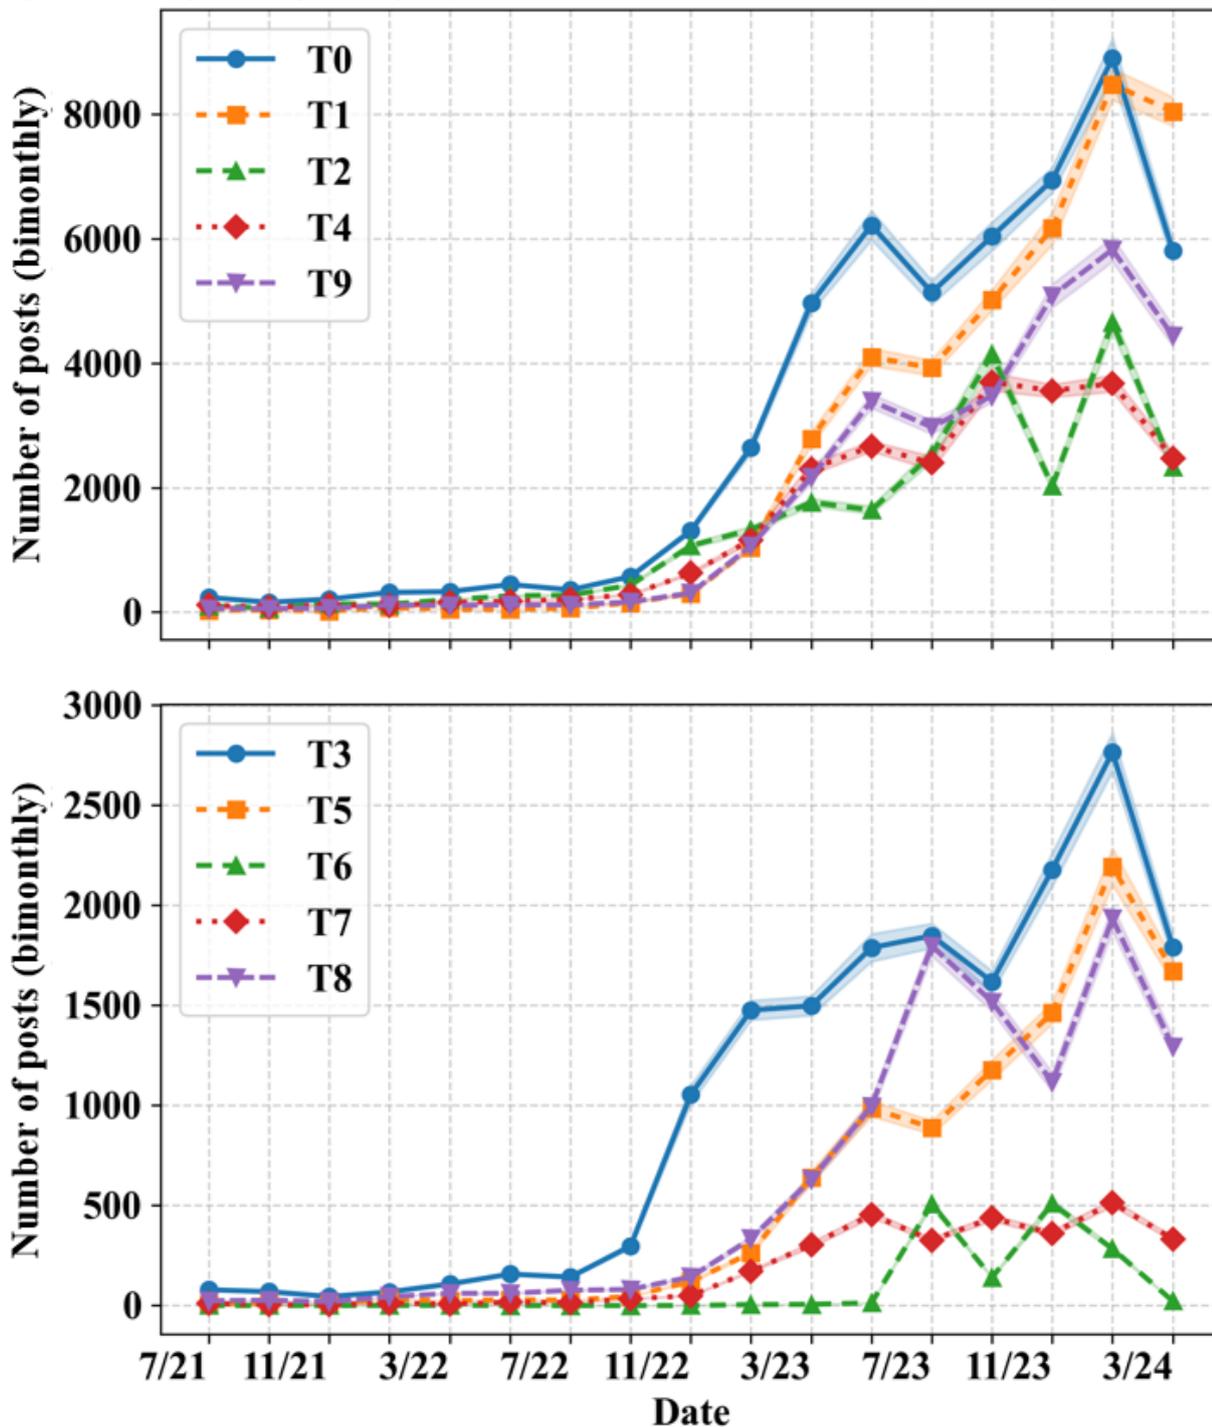
Starting in July 2023–September 2023, there is a noticeable uptick in the number of posts pertaining to T6 (suicide risk) and T8 (acute harm/adverse drug reactions). To determine whether this change exceeds baseline trends, we apply an ITS regression. The ITS results are visualized in Figure 8 (full results are available in Table S6 in Multimedia Appendix 1). A statistically

significant (β_2=449.51, P<.001) increase occurs at the intervention point, followed by a significant downward trend (β_3=–83.10, P=.02). The ITS results suggest that the intervention window coincides with a sharp increase in the prevalence of T6. The surge in the popularity of T6 may have been caused by a statement released by the European Medicines Agency on July 11, 2023, acknowledging "about 150 reports of possible cases of self-injury and suicidal thoughts" from "people using liraglutide and semaglutide medicines" [79].

Figure 8. Interrupted time series analysis of T6 (suicide risk). The intervention window is July 2023–September 2023. EMA: European Medicines Agency.
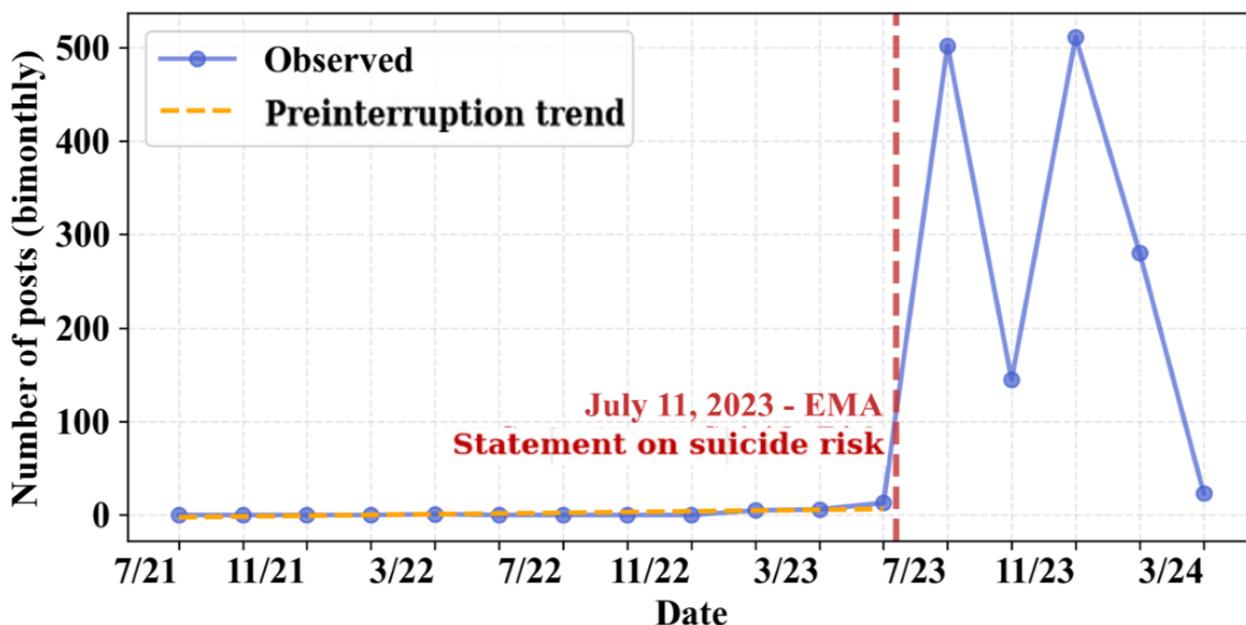
Breakdown of Topics by Sentiment

Figure 9 shows the breakdown of sentiment scores within each topic, within the gender, verification status, and account type subpopulations. This figure connects our sentiment and topic analyses and provides insights into potential sources of user negativity and positivity. For example, topics T5 (general and profane negativity), T6 (suicide risk), T7 (chronic harm), and T8 (acute harm/adverse drug reactions) are almost entirely negative.

Figure 9. Breakdown of sentiment scores by topic number, within various user subpopulations. T0: weight loss; T1: celebrities/politicians; T2: obtaining the drug; T3: drug indicators; T4: drug authorities; T5: general and profane negativity; T6: suicide risk; T7: chronic harm; T8: acute harm/adverse drug reactions; T9: other.
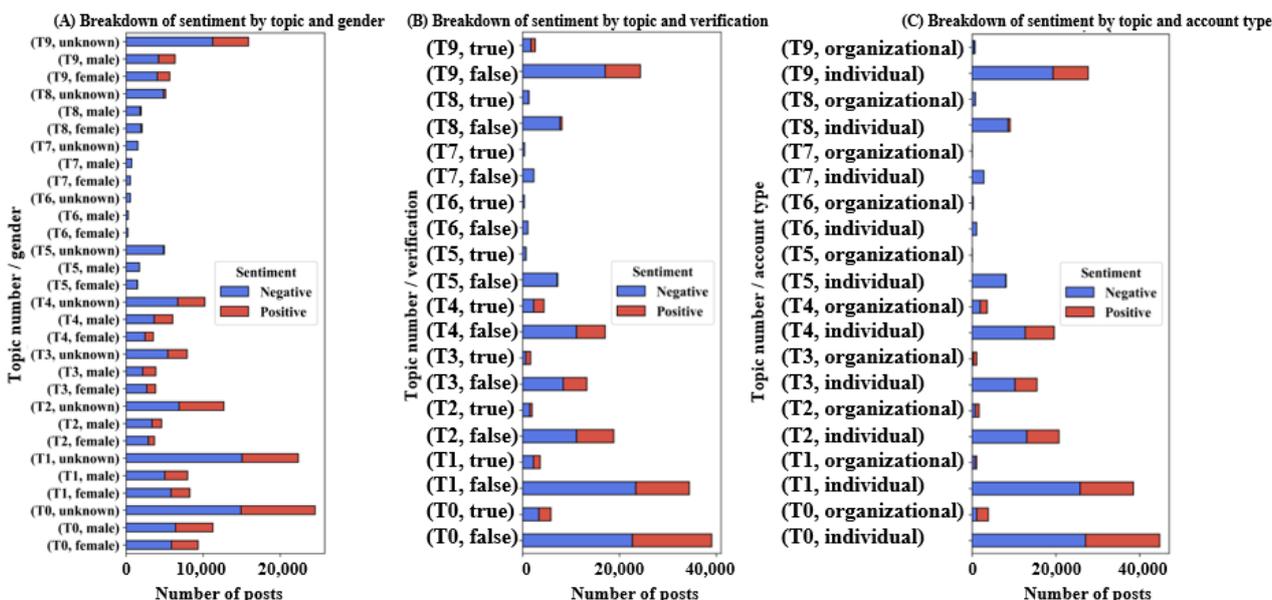
The breakdown also reveals notable differences in sentiment within topics. In T4, for instance, male-identified accounts contribute a significant volume of positive posts, whereas female-identified accounts are almost exclusively negative (Figure 9A). A similar, though less pronounced, pattern is visible in T9, where male accounts again show a larger share of positive

sentiment. Furthermore, in T0, organizational accounts exhibit a much higher *proportion* of positive sentiment compared to individual accounts, whose posts in that topic are predominantly negative (Figure 9C).

Discussion

Key Findings

The primary takeaway from our results is that, although semaglutide has generated considerable attention, its reception from individual X users is generally negative. This negativity suggests an overall public skepticism regarding the drug's efficacy, accessibility, and potential side effects, which are consistently highlighted in user discussions. Such negative sentiment is particularly notable in individual accounts as compared to the positive or neutral sentiment found more often in organizational accounts.

The temporal analysis demonstrates a general decrease in sentiment over time. In addition, we observe a notable shift in sentiment during late 2022, when regulatory announcements related to adverse effects and safety warnings surfaced [70]. This period sees a sharp increase in negative sentiment, which aligns with concerns raised about the drug's safety. The general throughline of negativity beginning in mid-2022 aligns temporally with FDA confirmation that Wegovy (semaglutide) is in shortage as well as heightened media coverage of semaglutide's off-label use for weight loss. The observed negativity likely reflects public concern about limited access and fairness, as shortages reported by the FDA were simultaneously amplified in professional and mainstream outlets. These shortages and their subsequent resolution are noted in multiple sources, including the FDA Drug Shortages Database and trade publications documenting semaglutide's removal from the shortage list and its ongoing legal and ethical implications [76,77].

On a similar note, discussions concerning a serious impact spike following the European Medicines Agency's report on suicides on July 11, 2023 [79]. These findings underscore the impact that regulatory decisions and public health announcements can have on shaping public perceptions, particularly when safety and efficacy concerns are at the forefront of the discussion. The ITS analyses reinforce our interpretation that the drop in sentiment and uptick in T6 exceed baseline trends, likely having been influenced by external events. Given the temporal alignment and the magnitude of the changes, it is likely that public sentiment is shaped in response to reports of adverse effects and that the European Medicines Agency statement on self-harm reports influenced public discussions.

In addition, user subpopulations show variance in sentiment and topic discussion. Users who are interested in "beauty/health" have the most positive sentiment and the highest prevalence of T0 (weight loss). Male users are slightly more positive than female users, though they appear less interested in T1 (celebrities/politicians). The observed gender differences may be influenced by the increased media attention surrounding celebrities endorsing the drug, which is more prominently featured in female-driven narratives about weight loss and

beauty [80,81]. On the whole, verified users and organizations are more positive than their counterparts. Additionally, T4 (drug authorities) is more prevalent among verified users and organizations, while T1 (celebrities/politicians) and T5 (general and profane negativity) are more prevalent among unverified users. These differences suggest that verified users (often public figures or organizations) tend to use a more conservative tone. On the other hand, nonverified users (often individuals) tend to emphasize personal concerns, particularly regarding side effects and affordability. This difference highlights how the identity and motivations of the speaker can influence sentiments, with verified accounts potentially downplaying issues for commercial or reputational reasons, while individual users are more candid about their negative experiences.

Practical Implications

Our findings have several practical implications for stakeholders, particularly health care providers and pharmaceutical companies. Overall, our results underscore the need for transparent communication strategies. These strategies should prioritize addressing concerns raised by users, including issues related to accessibility, side effects, and the drug's overall safety. Clear communication can help bridge the gap between public perception and medical realities while fostering trustworthy and informed decision-making.

For pharmaceutical companies, transparency in messaging about the use and effects of semaglutide is critical. Our observations highlight a gap in sentiment between individual and organizational accounts. One potential contributing factor is the societal emphasis on beauty and weight loss, which can be amplified by advertising and endorsements from influential figures [6]. This phenomenon often skews public perception and drives expectations. Given that regulatory announcements and safety warnings significantly shape public sentiment [79], pharmaceutical companies should adopt proactive approaches to build credibility. This includes consistent and timely communication that addresses misconceptions and reinforces the drug's benefits and limitations.

Limitations

While this study offers comprehensive insights into public perceptions of semaglutide on X, some limitations must be acknowledged. First, our dataset is limited to English-language posts, and our regional sentiment analysis is limited to the United States. While this constraint is necessary for linguistic consistency in sentiment and topic modeling, it may exclude important perspectives from non-English-speaking users, introducing a potential source of bias. Second, gender classification in our dataset is limited to 3 categories: male, female, and unknown. We recognize that gender is not binary and includes a spectrum of identities. However, the Brandwatch data source only provides these limited groupings. As a result, our analysis may not capture the experiences of gender-diverse users, representing a gap in inclusivity. Third, our study only makes use of data from the X platform, potentially limiting the generalizability of the results to different platforms (eg, Reddit and clinical forums). Future work may study additional platforms to build a broader understanding of user experience. Lastly, although our work identifies numerous trends in public

attitudes toward semaglutide, we do not interpret it as a causal impact of real-world reports (eg, semaglutide usage and adverse events) on online discourse.

Conclusions

Public interest in semaglutide has greatly increased in recent years. This study explores X users' experiences with semaglutide via an analysis of 859,751 X posts created between July 2021 and April 2024. We observe a general decrease in sentiment across most user subpopulations over time, with a particularly noteworthy decrease occurring in November 2022. Our research highlights the complex dynamics of user experiences with semaglutide, driven by a combination of user demographics, regional factors, and external events such as regulatory announcements. The practical implications of these findings are crucial for health care communicators and pharmaceutical companies seeking to engage with the public in a more informed, responsive, and regionally targeted manner. Future research should focus on further unraveling the role of side effects in shaping public opinion and exploring how sentiment changes in response to evolving health-related information.

Data Availability

Information about accessing an anonymized version of the dataset generated and analyzed during this study is available in Multimedia Appendix 1.

Authors' Contributions

Conceptualization: LL, JL
Data curation: LL, PM, GL
Formal analysis: PM, GL
Investigation: PM, GL
Methodology: PM, GL, LL
Project administration: LL, JL
Software: PM, GL, LL
Supervision: LL, JL
Visualization: PM, GL, LL
Writing – original draft: PM, GL, LL
Writing – review & editing: PM, GL, LL

Conflicts of Interest

None declared.

Multimedia Appendix 1

Additional dataset results.
[DOCX File , 1080 KB-Multimedia Appendix 1]

Multimedia Appendix 2

The complete mapping of all 200 initial clusters to the 10 umbrella topics.
[XLSX File (Microsoft Excel File), 41 KB-Multimedia Appendix 2]

References

1. Han SH, Safeek R, Ockerman K, Trieu N, Mars P, Klenke A, et al. Public interest in the off-label use of glucagon-like peptide 1 agonists (Ozempic) for cosmetic weight loss: a Google Trends analysis. Aesthet Surg J. 2023;44(1):60-67. [doi: 10.1093/asj/sjad211] [Medline: 37402640]
2. Mailhac A, Pedersen L, Pottegård A, Søndergaard J, Mogensen T, Sørensen HT, et al. Semaglutide (Ozempic) use in Denmark 2018 through 2023　user trends and off-label prescribing for weight loss. Clin Epidemiol. 2024;16:307-318. [FREE Full text] [doi: 10.2147/CLEP.S456170] [Medline: 38685990]
3. Tichy EM, Hoffman JM, Suda KJ, Rim MH, Tadrous M, Cuellar S, et al. National trends in prescription drug expenditures and projections for 2022. Am J Health Syst Pharm. 2022;79(14):1158-1172. [FREE Full text] [doi: 10.1093/ajhp/zxac102] [Medline: 35385103]
4. Alvarez-Mon MA, Llavero-Valero M, Asunsolo Del Barco A, Zaragozá C, Ortega MA, Lahera G, et al. Areas of interest and attitudes toward antiobesity drugs: hematic and quantitative analysis using Twitter. J Med Internet Res. 2021;23(10):e24336. [FREE Full text] [doi: 10.2196/24336] [Medline: 34698653]
5. Raubenheimer JE, Myburgh PH, Bhagavathula AS. Sweetening the deal: an infodemiological study of worldwide interest in semaglutide using Google Trends extended for health application programming interface. BMC Glob Public Health. 2024;2(1):63. [doi: 10.1186/s44263-024-00095-w] [Medline: 39681910]

6. Rad J, Melendez-Torres GJ. Critical discourse analysis of social media advertisements for GLP-1 receptor agonist weight loss drugs: implications for public perceptions and health communication. BMC Public Health. 2025;25(1):2996. [FREE Full text] [doi: 10.1186/s12889-025-24197-8] [Medline: 40890660]

7. Bremmer MP, Hendershot CS. Social media as pharmacovigilance: the potential for patient reports to inform clinical research on glucagon-like peptide 1 (GLP-1) receptor agonists for substance use disorders. J Stud Alcohol Drugs. 2024;85(1):5-11. [doi: 10.15288/jsad.23-00318] [Medline: 37917019]

8. Olsen AK, Whalen MD. Public perceptions of the pharmaceutical industry and drug safety: implications for the pharmacovigilance professional and the culture of safety. Drug Saf. 2009;32(10):805-810. [doi: 10.2165/11316620-000000000-00000] [Medline: 19722724]

9. Hua Y, Jiang H, Lin S, Yang J, Plasek JM, Bates DW, et al. Using Twitter data to understand public perceptions of approved versus off-label use for COVID-19-related medications. J Am Med Inform Assoc. 2022;29(10):1668-1678. [FREE Full text] [doi: 10.1093/jamia/ocac114] [Medline: 35775946]

10. Rathert C, Wyrwich MD, Boren SA. Patient-centered care and outcomes: a systematic review of the literature. Med Care Res Rev. 2013;70(4):351-379. [doi: 10.1177/1077558712465774] [Medline: 23169897]

11. Farsi D. Social media and health care, Part I: Literature review of social media use by health care providers. J Med Internet Res. 2021;23(4):e23205. [FREE Full text] [doi: 10.2196/23205] [Medline: 33664014]

12. Hao T, Huang Z, Liang L, Weng H, Tang B. Health natural language processing: methodology development and applications. JMIR Med Inform. 2021;9(10):e23898. [FREE Full text] [doi: 10.2196/23898] [Medline: 34673533]

13. Zhou B, Yang G, Shi Z, Ma S. Natural language processing for smart healthcare. IEEE Rev Biomed Eng. 2024;17:4-18. [doi: 10.1109/RBME.2022.3210270] [Medline: 36170385]

14. Wang E, Wang F, Yang Z, Wang L, Zhang Y, Lin H, et al. A graph convolutional network-based method for chemical-protein interaction extraction: algorithm development. JMIR Med Inform. 2020;8(5):e17643. [FREE Full text] [doi: 10.2196/17643] [Medline: 32348257]

15. Li L, Wang P, Wang Y, Wang S, Yan J, Jiang J, et al. A method to learn embedding of a probabilistic medical knowledge graph: algorithm development. JMIR Med Inform. 2020;8(5):e17645. [FREE Full text] [doi: 10.2196/17645] [Medline: 32436854]

16. Kaufman DR, Sheehan B, Stetson P, Bhatt AR, Field AI, Patel C, et al. Natural language processing–enabled and conventional data capture methods for input to electronic health records: a comparative usability study. JMIR Med Inform. 2016;4(4):e35. [FREE Full text] [doi: 10.2196/medinform.5544] [Medline: 27793791]

17. Iroju OG, Olaleke JO. A systematic review of natural language processing in healthcare. J Inf Technol Comput Sci. 2015;7(8):44-50. [doi: 10.5815/ijitcs.2015.08.07]

18. Elbattah M, Arnaud É, Gignon M, Dequen G. The role of text analytics in healthcare: a review of recent developments and applications. In: Proceedings of the 14th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2021). Setúbal, Portugal. SciTePress; 2021:825-832.

19. Yu Y, Guan Y, Hu Y. Natural language processing applications in social network analysis: a data mining approach. J Phys Conf Ser. 2024;2813(1):012009. [doi: 10.1088/1742-6596/2813/1/012009]

20. Sandu A, Cotfas L, Stănescu A, Delcea C. A bibliometric analysis of text mining: exploring the use of natural language processing in social media research. Appl Sci. 2024;14(8):3144. [doi: 10.3390/app14083144]

21. Sarker A, DeRoos A, Perrone J. Mining social media for prescription medication abuse monitoring: a review and proposal for a data-centric framework. J Am Med Inform Assoc. 2020;27(2):315-329. [FREE Full text] [doi: 10.1093/jamia/ocz162] [Medline: 31584645]

22. Nasralah T, El-Gayar O, Wang Y. Social media text mining framework for drug abuse: development and validation study with an opioid crisis case analysis. J Med Internet Res. 2020;22(8):e18350. [FREE Full text] [doi: 10.2196/18350] [Medline: 32788147]

23. Zhang J, Wang X, Zhou Y. Comparative analysis of semaglutide induced adverse reactions: insights from FAERS database and social media reviews with a focus on oral vs subcutaneous administration. Front Pharmacol. 2024;15:1471615. [FREE Full text] [doi: 10.3389/fphar.2024.1471615] [Medline: 39502525]

24. Cricelli L, Grimaldi M, Vermicelli S. Crowdsourcing and open innovation: a systematic literature review, an integrated framework and a research agenda. Rev Manag Sci. 2021;16(5):1269-1310. [doi: 10.1007/s11846-021-00482-9]

25. Li L, Hu S, Dai Y, Deng M, Momeni P, Laverghetta G, et al. Toward satisfactory public accessibility: a crowdsourcing approach through online reviews to inclusive urban design. Comput Environ Urban Syst. 2025;122:102329. [doi: 10.1016/j.compenvurbsys.2025.102329]

26. Certomà C, Corsini F, Rizzi F. Crowdsourcing urban sustainability: data, people and technologies in participatory governance. Futures. 2015;74:93-106. [doi: 10.1016/j.futures.2014.11.006]

27. Alvaro N, Conway M, Doan S, Lofi C, Overington J, Collier N. Crowdsourcing Twitter annotations to identify first-hand experiences of prescription drug use. J Biomed Inform. 2015;58:280-287. [FREE Full text] [doi: 10.1016/j.jbi.2015.11.004] [Medline: 26556646]

28. Cascini F, Pantovic A, Al-Ajlouni YA, Failla G, Puleo V, Melnyk A, et al. Social media and attitudes towards a COVID-19 vaccination: a systematic review of the literature. eClinicalMedicine. 2022;48:101454. [FREE Full text] [doi: 10.1016/j.eclinm.2022.101454] [Medline: 35611343]

29. Sarker A, Ginn R, Nikfarjam A, O'Connor K, Smith K, Jayaraman S, et al. Utilizing social media data for pharmacovigilance: a review. J Biomed Inform. 2015;54:202-212. [FREE Full text] [doi: 10.1016/j.jbi.2015.02.004] [Medline: 25720841]

30. Müller MM, Salathé M. Crowdbreaks: tracking health trends using public social media data and crowdsourcing. Front Public Health. 2019;7:81. [FREE Full text] [doi: 10.3389/fpubh.2019.00081] [Medline: 31037238]

31. Signorini A, Segre AM, Polgreen PM. The use of Twitter to track levels of disease activity and public concern in the U.S. during the influenza A H1N1 pandemic. PLoS One. 2011;6(5):e19467. [FREE Full text] [doi: 10.1371/journal.pone.0019467] [Medline: 21573238]

32. Boucher JC, Cornelson K, Benham JL, Fullerton MM, Tang T, Constantinescu C, et al. Analyzing social media to explore the attitudes and behaviors following the announcement of successful COVID-19 vaccine trials: infodemiology study. JMIR Infodemiology. 2021;1(1):e28800. [FREE Full text] [doi: 10.2196/28800] [Medline: 34447924]

33. Buntain C, McGrath E, Golbeck J, LaFree G. Comparing social media and traditional surveys around the Boston Marathon Bombing. In: Proceedings of the 6th Workshop on 'Making Sense of Microposts' co-located with the 25th International World Wide Web Conference (WWW 2016). Montréal, Canada. CEUR-WS.org; 2016:34-41.

34. Lardon J, Abdellaoui R, Bellet F, Asfari H, Souvignet J, Texier N, et al. Adverse drug reaction identification and extraction in social media: a scoping review. J Med Internet Res. 2015;17(7):e171. [FREE Full text] [doi: 10.2196/jmir.4304] [Medline: 26163365]

35. Birjali M, Kasri M, Beni-Hssane A. A comprehensive survey on sentiment analysis: approaches, challenges and trends. Knowl Based Syst. 2021;226:107134. [doi: 10.1016/j.knosys.2021.107134]

36. Feldman R. Techniques and applications for sentiment analysis. Commun ACM. 2013;56(4):82-89. [doi: 10.1145/2436256.2436274]

37. Tsao SF, Chen H, Tisseverasinghe T, Yang Y, Li L, Butt ZA. What social media told us in the time of COVID-19: a scoping review. Lancet Digit Health. 2021;3(3):e175-e194. [FREE Full text] [doi: 10.1016/S2589-7500(20)30315-0] [Medline: 33518503]

38. Alamoodi AH, Zaidan BB, Zaidan AA, Albahri OS, Mohammed KI, Malik RQ, et al. Sentiment analysis and its applications in fighting COVID-19 and infectious diseases: a systematic review. Expert Syst Appl. 2021;167:114155. [FREE Full text] [doi: 10.1016/j.eswa.2020.114155] [Medline: 33139966]

39. Andhale S, Mane P, Vaingankar M, Karia D, Talele KT. Twitter sentiment analysis for COVID-19. In: 2021 International Conference on Communication information and Computing Technology (ICCICT). New York. IEEE; 2021:1-12.

40. He L, He C, Reynolds TL, Bai Q, Huang Y, Li C, et al. Why do people oppose mask wearing? A comprehensive analysis of U.S. tweets during the COVID-19 pandemic. J Am Med Inform Assoc. 2021;28(7):1564-1573. [FREE Full text] [doi: 10.1093/jamia/ocab047] [Medline: 33690794]

41. Jelodar H, Wang Y, Yuan C, Feng X, Jiang X, Li Y, et al. Latent dirichlet allocation (LDA) and topic modeling: models, applications, a survey. Multimed Tools Appl. 2018;78(11):15169-15211. [doi: 10.1007/s11042-018-6894-4]

42. Grootendorst M. BERTopic: neural topic modeling with a class-based TF-IDF procedure. arXiv. Preprint posted online. Mar 11, 2022. [FREE Full text] [doi: 10.48550/arXiv.2203.05794]

43. Churchill R, Singh L. The evolution of topic modeling. ACM Comput Surv. 2022;54(10s):1-35. [doi: 10.1145/3507900]

44. Asghari M, Sierra-Sosa D, Elmaghraby A. Trends on health in social media: analysis using Twitter topic modeling. In: 2018 IEEE International Symposium on Signal Processing and Information Technology (ISSPIT). New York. IEEE; 2018:558-563.

45. Liu Q, Zheng Z, Zheng J, Chen Q, Liu G, Chen S, et al. Health communication through news media during the early stage of the COVID-19 outbreak in China: digital topic modeling approach. J Med Internet Res. 2020;22(4):e19118. [FREE Full text] [doi: 10.2196/19118] [Medline: 32302966]

46. Ramondt S, Kerkhof P, Merz E. Blood donation narratives on social media: a topic modeling study. Transfus Med Rev. 2022;36(1):58-65. [FREE Full text] [doi: 10.1016/j.tmrv.2021.10.001] [Medline: 34810071]

47. Paul MJ, Dredze M. Discovering health topics in social media using topic models. PLoS One. 2014;9(8):e103408. [FREE Full text] [doi: 10.1371/journal.pone.0103408] [Medline: 25084530]

48. Fong S, Carollo A, Lazuras L, Corazza O, Esposito G. Ozempic (glucagon-like peptide 1 receptor agonist) in social media posts: unveiling user perspectives through Reddit topic modeling. Emerg Trends Drugs Addict Health. 2024;4:100157. [doi: 10.1016/j.etdah.2024.100157]

49. Javaid A, Baviriseaty S, Javaid R, Zirikly A, Kukreja H, Kim CH, et al. Trends in glucagon-like peptide-1 receptor agonist social media posts using artificial intelligence. JACC Adv. 2024;3(9):101182. [FREE Full text] [doi: 10.1016/j.jacadv.2024.101182] [Medline: 39372460]

50. Somani S, Jain SS, Sarraju A, Sandhu AT, Hernandez-Boussard T, Rodriguez F. Using large language models to assess public perceptions around glucagon-like peptide-1 receptor agonists on social media. Commun Med (Lond). 2024;4(1):137. [doi: 10.1038/s43856-024-00566-z] [Medline: 38987347]

51. The most intelligent social suite. Brandwatch. URL: https://www.brandwatch.com/ [accessed 2026-01-29]

52. Amaro A, Sugimoto D, Wharton S. Efficacy and safety of semaglutide for weight management: evidence from the STEP program. Postgrad Med. 2022;134(Suppl 1):5-17. [FREE Full text] [doi: 10.1080/00325481.2022.2147326] [Medline: 36691309]

53. Jaume J. Understanding your market with new demographic insights. Brandwatch. Feb 13, 2014. URL: https://www.brandwatch.com/blog/demographic-insights/ [accessed 2026-01-29]

54. Brabete AC, Greaves L, Maximos M, Huber E, Li A, Lê M-L. A sex- and gender-based analysis of adverse drug reactions: a scoping review of pharmacovigilance databases. Pharmaceuticals (Basel). 2022;15(3):298. [FREE Full text] [doi: 10.3390/ph15030298] [Medline: 35337096]

55. Furman JL. Location and organizing strategy: exploring the influence of location on the organization of pharmaceutical research. In: Baum JAC, Sorenson O, editors. Geography and Strategy. Australia, Malaysia. Emerald Group Publishing Limited; 2003:49-87.

56. Podolsky MI, Raquib R, Shafer PR, Hempstead K, Ellis RP, Stokes AC. Factors associated with semaglutide initiation among adults with obesity. JAMA Netw Open. 2025;8(1):e2455222. [FREE Full text] [doi: 10.1001/jamanetworkopen.2024.55222] [Medline: 39836425]

57. Momeni P, Laverghetta G. OJPHI__twitter_semaglutide_clean. GitHub. URL: https://github.com/ParisaMomeni/OJPHI_twitter_semaglutide_clean [accessed 2026-01-29]

58. Williams ML, Burnap P, Sloan L. Towards an ethical framework for publishing Twitter data in social research: taking into account users' views, online context and algorithmic estimation. Sociology. 2017;51(6):1149-1168. [FREE Full text] [doi: 10.1177/0038038517708140] [Medline: 29276313]

59. Skinner T, Brance K, Halligan S, Tsang E, Girling H. Coping with emotionally challenging research: developing a strategic approach to researcher wellbeing. J Acad Ethics. 2025;23(4):2559-2583. [doi: 10.1007/s10805-025-09665-5] [Medline: 41001435]

60. Twitter-roBERTa-base for sentiment analysis. Hugging Face. URL: https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest [accessed 2026-01-29]

61. Loureiro D, Barbieri F, Neves L, Espinosa AL, Camacho-collados J. TimeLMs: diachronic language models from Twitter. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations. Dublin, Ireland. Association for Computational Linguistics; 2022:251-260.

62. Liao W, Zeng B, Yin X, Wei P. An improved aspect-category sentiment analysis model for text sentiment analysis based on RoBERTa. Appl Intell. 2020;51(6):3522-3533. [doi: 10.1007/s10489-020-01964-1]

63. Ma Z, Li L, Mao Y, Wang Y, Patsy OG, Bensi MT, et al. Surveying the use of social media data and natural language processing techniques to investigate natural disasters. Nat Hazards Rev. 2024;25(4):03124003. [doi: 10.1061/NHREFO.NHENG-2047]

64. Lou Y, Zhang Y, Li F, Qian T, Ji D. Emoji-based sentiment analysis using attention networks. ACM Trans Asian Low-Resour Lang Inf Process. 2020;19(5):1-13. [doi: 10.1145/3389035]

65. Cureg MQ, De La Cruz JAD, Solomon JCA, Saharkhiz AT, Balan AKD, Samonte MJC. Sentiment analysis on tweets with punctuations, emoticons, and negations. In: ICISS '19: Proceedings of the 2nd International Conference on Information Science and Systems. New York. Association for Computing Machinery; 2019:266-270.

66. Ikotun AM, Ezugwu AE, Abualigah L, Abuhaija B, Heming J. K-means clustering algorithms: a comprehensive review, variants analysis, and advances in the era of big data. Inf Sci. 2023;622:178-210. [doi: 10.1016/j.ins.2022.11.139]

67. Marzi G, Balzano M, Marchiori D. K-Alpha Calculator-Krippendorff's Alpha Calculator: a user-friendly tool for computing Krippendorff's Alpha inter-rater reliability coefficient. MethodsX. 2024;12:102545. [doi: 10.1016/j.mex.2023.102545] [Medline: 39669968]

68. FDA clarifies policies for compounders as national GLP-1 supply begins to stabilize. U.S. Food and Drug Administration. URL: https://www.fda.gov/drugs/drug-safety-and-availability/fda-clarifies-policies-compounders-national-glp-1-supply-begins-stabilize [accessed 2026-01-29]

69. FDA approves once-weekly Wegovy injection for the treatment of obesity in teens aged 12 years and older. Drugs.com. Dec 2022. URL: https://www.drugs.com/newdrugs/fda-approves-once-weekly-wegovy-obesity-teens-aged-12-years-older-5949.html [accessed 2026-01-29]

70. Shu Y, He X, Wu P, Liu Y, Ding Y, Zhang Q. Gastrointestinal adverse events associated with semaglutide: a pharmacovigilance study based on FDA adverse event reporting system. Front Public Health. 2022;10:996179. [FREE Full text] [doi: 10.3389/fpubh.2022.996179] [Medline: 36339230]

71. Jayanthi R. The demedicalization of GLP-1 receptor agonists: an analysis of TikTok influencer endorsement of Ozempic-like medications for cosmetic weight loss [Thesis]. University of North Carolina at Chapel Hill. Mar 26, 2025. URL: https://cdr.lib.unc.edu/concern/honors_theses/8c97m561s [accessed 2026-01-29]

72. Wegovy (semaglutide) injection 2.4 mg cardiovascular outcomes data presented at the American Heart Association Scientific Sessions and simultaneously published in the New England Journal of Medicine. Drugs.com. Nov 23, 2023. URL: https://www.drugs.com/clinical_trials/wegovy-semaglutide-2-4-mg-cardiovascular-outcomes-data-presented-american-heart-association-21170.html [accessed 2026-01-29]

73.   Semaglutide reduces severity of common liver disease in people with HIV. Drugs.com. Mar 5, 2024. URL: https://www.drugs.com/clinical_trials/semaglutide-reduces-severity-common-liver-hiv-21374.html [accessed 2026-01-29]

74.   Wegovy approved in the US for cardiovascular risk reduction in people with overweight or obesity and established cardiovascular disease. Drugs.com. Mar 8, 2024. URL: https://www.drugs.com/newdrugs/wegovy-approved-us-cardiovascular-risk-reduction-overweight-obesity-established-cardiovascular-6212.html [accessed 2026-01-29]

75.   FDA drug shortages: semaglutide injection (Wegovy). U.S. Food and Drug Administration. URL: https://www.accessdata.fda.gov/scripts/drugshortages/dsp_ActiveIngredientDetails.cfm?AI=Semaglutide%20Injection&st=c [accessed 2026-01-29]

76.   Ghiam A, Heba D. Semaglutide's removal from the FDA shortages list sets the stage for more Novo Nordisk lawsuits. Medical Economics. Aug 18, 2025. URL: https://www.medicaleconomics.com/view/semaglutide-s-removal-from-the-fda-shortages-list-sets-the-stage-for-more-novo-nordisk-lawsuits [accessed 2026-01-29]

77.   Ferruggia K. FDA ends semaglutide shortage listing, contributing to ongoing legal challenges. Pharmacy Times. Feb 27, 2025. URL: https://www.pharmacytimes.com/view/fda-ends-semaglutide-shortage-listing-contributing-to-ongoing-legal-challenges [accessed 2026-01-29]

78.   Cramér H. Mathematical Methods of Statistics. Princeton, NJ. Princeton University Press; 1946.

79.   EMA statement on ongoing review of GLP-1 receptor agonists. European Medicines Agency. Jul 11, 2023. URL: https://www.ema.europa.eu/en/news/ema-statement-ongoing-review-glp-1-receptor-agonists [accessed 2026-01-29]

80.   Faw MH, Davidson K, Hogan L, Thomas K. Corumination, diet culture, intuitive eating, and body dissatisfaction among young adult women. Pers Relat. 2020;28(2):406-426. [doi: 10.1111/pere.12364]

81.   Mills J, Fuller-Tyszkiewicz M. Fat talk and body image disturbance: a systematic review and meta-analysis. Psychol Women Q. 2016;41(1):114-129. [doi: 10.1177/0361684316675317]

Abbreviations

FDA: US Food and Drug Administration
ITS: interrupted time-series
MASLD: metabolic dysfunction-associated steatotic liver disease
NLP: natural language processing
RoBERTa: Robustly Optimized BERT Pretraining Approach
RQ: research question
SELECT: Semaglutide Effects on Cardiovascular Outcomes in People With Overweight or Obesity