Proceedings of the Twentieth International AAAI Conference on Web and Social Media (ICWSM 2026)

How Growing Toxicity Manifests: A Topic Trajectory Analysis of U.S. Immigration Discourse on Social Media

Una Joh[1], Yiqi Li[1], Jeff Hemsley[1]

[1]School of Information Studies, Syracuse University
sjoh01@syr.edu, yli360@syr.edu, jjhemsle@syr.edu

Abstract

In the online public sphere, discussions about immigration often become increasingly fractious, marked by toxic language and polarization. Drawing on 4 million X posts over six months, we combine a user- and topic-centric approach to study how shifts in toxicity manifest as topical shifts. Our topic discovery method, which leverages instruction-based embeddings and recursive HDBSCAN, uncovers 157 fine-grained subtopics within the U.S. immigration discourse. We focus on users in four groups: those with increasing toxicity, those with decreasing toxicity, and two reference groups with no significant toxicity trend but matched toxicity levels. Treating each posting history as a trajectory through a five-dimensional topic space, we compare average group trajectories using permutational MANOVA. Our findings show that users with increasing toxicity drift toward alarmist, fear-based frames, whereas those with decreasing toxicity pivot toward legal and policy-focused themes. Both patterns diverge statistically significantly from their reference groups. This pipeline, which combines hierarchical topic discovery with trajectory analysis, offers a replicable method for studying dynamic conversations around social issues at scale.

Introduction

Social media is an increasingly indispensable public sphere, offering a democratized space for diverse discussion. However, alongside the problems such as fragmentation, addiction, misinformation, toxicity is another critical issue that may emerge on social media, potentially leading to harms such as opinion distortion, biases, hostility, polarization, and even real-world violence (Bruns and Highfield 2015; Kim et al. 2021; Gallacher, Heerdink, and Hewstone 2021; Klein and Majdoubi 2024). Immigration, one of the most contentious topics discussed everyday on social media, exemplifies the diversities and complexities of social media discussion (Mittos et al. 2020). For example, as per Pew Research's recent report, immigration was a major point of contention in the 2024 presidential election, deepening the divide between Trump and Harris supporters (Mukherjee and Krogstad 2024). Immigration has been raised by existing scholarship to be a space contaminated by toxic content and divisive debates (Santana 2015). This study draws from a rich, longitudinal, and comprehensive immigration-related social media data, and explores the dynamics of toxicity in the online public sphere of X (formerly known as Twitter).

Conceptually, this research is novel in a few ways. Firstly, it bridges the isolation between social media toxicity research that are centered on topics (e.g., (Klein and Majdoubi 2024; Rossini and Maia 2021; Stromer-Galley, Bryant, and Bimber 2015)), and those that revolve around toxic users or behaviors (e.g., (Rajadesingan, Resnick, and Budak 2020; Coe, Kenski, and Rains 2014)). We found that users with increasing toxicity and those with decreasing toxicity engage with distinctive sets of topics, differing from the topic trajectories of users with same average toxicity level with each groups. This highlights that toxicity in communication can follow complex and varied topical pathways rather than simply representing a static trait of individual users.

Secondly, under the backdrop of this rich and contentious immigration issue public sphere, seldom is there longitudinal observations on the co-evolutionary patterns of users' dynamic engagement across topics, and toxicity communication patterns. Lastly, while most existing research studies extreme toxic behaviors or toxic accounts (Qayyum et al. 2023; Kumar et al. 2022), this current research sheds light on users of different toxic tendencies (specifically, escalating or reducing toxicity over time). Such a focus not only provides unique angle on a group of shifting and impacted individuals in the issue community, but also drives inquiries of associated factors linked with toxicity changes. For example, we found that as users' toxicity rises, they gravitate toward alarmist, threat-framed narratives, whereas users whose toxicity declines increasingly engage with procedural or policy-oriented themes.

In addition, this study makes methodological contributions. We refine topic discovery models by incorporating instruction-based embedding and a recursive HDBSCAN clustering framework with hierarchical merging based on topic coherence. This technique mitigates the challenge of over-segmentation, yielding interpretable subtopics that capture both the semantic and stance-related nuances of immigration discourse. Moreover, we pair this topic-discovery framework with a trajectory-analysis pipeline that models each user's posting history as a continuous path through embedding space and statistically compares group trajectories, offering a scalable way to link topical movement with shifts in toxic behavior.

Related Work & Research Questions

Toxicity on Social Media

Social media platforms like X reflect Habermas' concept of the public sphere, increasingly playing a critical role in facilitating democratized, open, real-time, and diverse discussions. However, the social media public sphere is also highly fragmented, characterized by toxic conversations, biased viewpoints, and increasing polarization (Bruns and Highfield 2015). Catalyzed by factors such as the social media algorithms, social influences, and emotional contagion, toxicity on these platforms is significantly impacting participants, distorting public opinion, fueling negativity, and exacerbating polarization (Kim et al. 2021). Social media toxicity harms public discourse by triggering irrational discussions, obstructing productive conversation, and engendering shallow deliberation (Klein and Majdoubi 2024). As the emotional foundation of toxic language further spreads through social media networks, adverse outcomes may also emerge, including the spread of negativity, hostility, polarization, or even real-world violence across groups (Gallacher, Heerdink, and Hewstone 2021; Klein and Majdoubi 2024).

Toxicity is defined as communication "manifest(ed) in the tone and style with which a speaker attacks their addressee's 'face,' or public self-image" (Sydnor 2019, p. 5). A plethora of social media toxicity research falls into topic-driven or user-centered realms. Topic-wise, toxicity is often linked with contextual factors such as information sources or topics (Klein and Majdoubi 2024). Existing research identifies that "hard news" tends to generate high incivility or toxicity, while lighter topics such as lifestyle or technology are linked with reduced toxicity (Klein and Majdoubi 2024). When theme of the toxic content is targeting different groups (especially LGBTQ population), audiences are likely to seek for content moderation from the platform, although the general motivation of moderation-seeking is limited (Pradel et al. 2024). Political topics are also especially likely to provoke toxic discussion (Chen and Wang 2022; Rossini and Maia 2021; Stromer-Galley, Bryant, and Bimber 2015). Immigration is among a few topics (e.g., climate change, genetic testing) that are especially toxicity-prone (Mittos et al. 2020; Salminen et al. 2020; Santana 2015).

From a user-centered perspective, researchers mainly study what user-level factors contribute to toxic communication and community behaviors. For example, Rajadesingan and colleagues found that pre-entry learning allows newcomers into the Reddit communities to conform to the community's preexisting toxicity norms (Rajadesingan, Resnick, and Budak 2020). Coe and colleagues found that frequent commenters on newspaper websites tend to be less toxic in comments than less frequent commenters (Coe, Kenski, and Rains 2014). Importantly, many studies focus on extreme behaviors by tracking toxic profiles (Qayyum et al. 2023; Kumar et al. 2022), this study takes a novel approach of identifying types of invested discussants from a longitudinal angle, exploring users who are naturally engaged in an issue discussion, and examining over-time engagement patterns—specifically those whose toxicity increases and decreases within an issue space (Yang 2020).

The topic-centered and user-centered perspectives of toxicity are often approached in isolation. Such isolation risks oversimplifying toxicity by missing the potential interplay between user behaviors and topics. To address this gap, this study examines longitudinal patterns of users' over-time toxic communication patterns, and explore the interplay between different types of users and their topic-engagement. This way, we are able to map out how the topical and user-level toxicity level co-evolve, thus providing valuable empirical observations on where the toxicity may emerge, and where cross-group dialogue and shared values that reduce toxicity may develop.

This study situates against the backdrop of longitudinal immigration-related discussion on X. As discussed, immigration is one of the most polarized and emotionally charged issues (Mittos et al. 2020; Salminen et al. 2020), strongly associated with many related issues of discussion such as racial and ethnicity, unemployment, crime, border safety, economic outlook, social justice, and more (Santana 2015). Empirical analysis on the discourse around immigration informs understanding of how the changing and controversial policies (e.g., DACA, refugee cities, border security) are linked with audiences' evolving communication and engagement patterns. Immigration can be conceptualized as a social issue space and a bounded issue ecology because it "channels public attention and provides a space for the communication of identities and ideologies" (Yang 2020, p. 9).

Research Questions

As outlined in the previous subsection, research on social media toxicity typically adopts one of two primary perspectives: topic-centric or user-centric. While each approach has yielded valuable insights, they are often treated in isolation. To address this gap, this study integrates both perspectives by identifying user groups with increasing or decreasing toxicity levels and analyzing the trajectories of the subtopics these groups engage with. We statistically compare these groups to corresponding reference groups (i.e., users with the same average toxicity as each group, but without a statistically significant trend of increasing or decreasing toxicity). This comparison aims to determine whether differences in subtopic trajectories are associated with changes in toxicity over time, independent of overall toxicity levels, thereby uncovering the thematic contexts that accompany escalation or de-escalation. Accordingly, the research questions of this study are:

1. **RQ1 (Increasing Toxicity Users):**
   Is topic engagement different for users whose toxicity *increases* compared to the users whose toxicity does not?

2. **RQ2 (Decreasing Toxicity Users):**
   Is topic engagement different for users whose toxicity *decreases* compared to the users whose toxicity does not?

Data

To answer our research questions, we collected data from X using U.S. immigration-related keywords. X was chosen because it is one of the major social media platforms in the U.S. context (Gottfried 2024) and is also known for its abundance of toxicity, especially in recent times (Hickey et al. 2023). Due to the limitations of X's official API, we used Apify, which is a web scraping and automation platform, to collect the data.

The collection time frame was from April 17, 2023, to October 27, 2023. This timeframe was selected because, according to a poll (Jones 2024), this period coincided with a noticeable worsening of sentiment toward immigration issues. We concluded that using this timeframe would allow us to capture a substantial portion of the dynamics of user toxicity change.

Data collection was performed using keywords. According to the API provider, they extracted posts from the web search feature of X. This feature offered two options, "latest" and "top," but the distinction between these options was not clearly explained in X's official documentation. After consulting with the API provider, we chose the "latest" option to collect the most comprehensive data, following their advice. Additionally, we restricted the language option to English because it was an economical way to filter out a substantial number of posts regarding non-U.S. immigration issues, despite the potential limitation this poses for the scope of discourse we can collect.

The search query used for the collection was as follows:

> *(immigrant OR immigrants OR immigration OR migrant OR migrants OR migration OR illegals OR undocumented OR refugee OR refugees OR "guest worker" OR "guest workers" OR "asylum seeker" OR "asylum seekers" OR "illegal alien" OR "illegal aliens") AND (USA OR "U.S." OR "United States" OR "the US" OR America OR American OR Americans OR Biden OR Trump)*

We opted for broad terms rather than attempting to construct an exhaustive list of keywords because X's web UI does not function properly when too many keywords are provided as input. Additionally, X's documentation does not clearly specify the maximum number of keywords allowed.

In total, we collected 8,995,234 posts, including original posts, quotes, and replies. The API did not have a feature to collect reposts (retweets), which was less important for our research since we aimed to trace changes in the toxicity of individual users' posts.

Methodology

Classification for Filtering Relevant Posts

Although the search query was designed to retrieve relevant posts about U.S. immigration issues, it was inevitable to collect some irrelevant posts due to the limitations of keyword-based retrieval. To address this, we filtered out irrelevant posts from the dataset.

Given the strong performance of decoder-based large language models (LLMs) in classification tasks within social science contexts (Ziems et al. 2024), we employed a decoder-based LLM to filter out irrelevant posts. Considering the simplicity of the task, we performed zero-shot classification and evaluated the results using the F1 score.

Each post was processed independently as input, with the classification task divided into two steps. The prompt for Task 1 was: "Is this tweet about immigration? Answer with either 'Yes' or 'No.' " Only posts classified as "Yes" in Task 1 proceeded to Task 2, which used the following prompt: "Is this tweet in a U.S. context? Answer with either 'Yes' or 'No.' "

Due to the simplicity of the task, we used a small-sized LLM, the unquantized Gemma2-9B-instruct[1] model, which has a context window size of 8,192 tokens. This was sufficient for our use case, as each input consisted of only one question and one post. To ensure reproducibility, we set the temperature to 0 and the top-p value to 0.9. The model was accessed via the Deep Infra[2] API, with a total processing cost of $14.66.

To validate the classification performance, we conducted human coding on 300 posts prior to running the zero-shot classification with the Gemma2 model. The human coders included one of the authors and a fourth-year Ph.D. student at a U.S. institution. Inter-coder reliability was measured using Cohen's Kappa (Cohen 1960), yielding scores of 0.86 for Task 1 and 0.79 for Task 2. Discrepancies were resolved through discussions among the coders to establish the gold label.

Gemma 2 achieved an F1 score of 0.92 on Task 1 and 0.81 on Task 2 when evaluated against the gold labels. Most errors were false positives, indicating that the model is more permissive than the human annotators. In Task 1 (immigration vs. other), for example, Gemma 2 tagged a post reminiscing about photographing the annual migration of birds as "immigration-related," presumably because of the lexical overlap with migration. Likewise, a discussion of Canadian elections that mentioned the author's immigrant background was flagged as immigration content even though human immigration policy was not the focus. Typical misclassifications in Task 2 (U.S. vs. non-U.S.) include a post criticizing government spending on illegal immigrants in hotels that Gemma2 labeled as U.S.-specific despite wording that is also common in the U.K.

Fortunately, the issue of false positives is less of a problem than false negatives, as those false positives are likely mostly filtered out during the clustering process as outliers or small clusters. After filtering, a total of 4,651,275 posts were used for the subsequent analysis.

Toxicity Assessment

To assess the toxicity of posts on X, we once again leveraged large language models (LLMs) instead of using Google's Perspective API, which is the most prevalent tool for evaluating toxicity in user-generated online content (Gervais, Dye, and Chin 2025). The primary reason for not using the Perspective API is that its outputs represent the likelihood

[1]https://huggingface.co/google/gemma-2-9b-it
[2]https://deepinfra.com

of toxicity, rather than providing a true measure of severity, as Gervais, Dye, and Chin (2025) has noted. Another widely used alternative, the Detoxify library (Hanu and Unitary team 2020), shares this limitation, as it is trained on the same dataset as the Perspective API.

We explored the potential of LLMs to offer an alternative approach by prompting the models to evaluate toxicity in a more human-like manner. In a recent study, de Wynter et al. (2025) investigated whether LLMs can serve as reliable toxicity evaluators across multiple languages, introducing RTP-LX, a human-annotated benchmark. Their findings showed that toxicity ratings generated by GPT-4 Turbo aligned closely with human annotations on the English subset of the RTP-LX dataset.

In our study, we experimented with two different LLMs to rate toxicity on a Likert-type scale ranging from 1 (benign) to 5 (highly toxic): OpenAI's GPT-4.1-nano-2025-04-14 and Google's Gemma3-4B-Instruct[3]. Given the size of our dataset, we limited our evaluation to smaller models to reduce computational costs. Both models were configured with a temperature of 0 and a top-p value of 0.9.

For each text sample in the English subset of RTP-LX, we prompted the models to output a toxicity score using a slightly modified version of the prompt provided in the original RTP-LX paper (de Wynter et al. 2025), tailored to our context (see Appendix B). We then compared the LLM-generated scores to ground-truth human ratings. The OpenAI model achieved a Pearson correlation of 0.7701 with the human ratings, while the Gemma3-4B model yielded a correlation of 0.7204. Our results suggest that LLMs can approximate human judgments of toxicity with a reasonably strong degree of correlation, especially considering the inherently subjective nature of the task.

Despite the superior performance of GPT-4.1-nano, we opted for the Gemma3-4B-Instruct model via the Deep Infra API, primarily for cost-efficiency, given that the OpenAI API was nearly ten times more expensive. The total cost using the Gemma model was $20.52. For our subsequent analysis, we normalized toxicity scores from the 1–5 Likert scale to a 0–100 scale.

Topic Discovery

Identifying subtopics in U.S. immigration-related discourse is a crucial task for our study. Since latent Dirichlet allocation (LDA) was first introduced by Blei et al. (2003), numerous topic modeling techniques have been proposed (Vayansky and Kumar 2020). In particular, with the advancement of natural language processing and deep neural networks, a variety of topic models have been developed, broadly categorized as neural topic models (Wu, Nguyen, and Luu 2024).

A characteristic of these models is their probabilistic nature. Specifically, probabilistic topic models assume that topics are defined as probability distributions over keywords. Documents are generated as sequences of words sampled from these topics. The sampling follows the probability distribution of topics assigned to the document and the probability distribution of keywords within each topic.

Criticism of these assumptions in probabilistic topic models has led to the exploration of simpler frameworks, such as embedding-clustering-based topic discovery models (Thompson and Mimno 2020; Angelov 2020; Zhang et al. 2022). These models avoid the assumption that documents are generated based on probabilistic distributions, often categorized as "topic discovery" models rather than traditional "topic modeling" approaches (Wu, Nguyen, and Luu 2024). One such topic discovery model that has gained popularity among social science researchers is BERTopic (Grootendorst 2022). BERTopic combines document embedding using Sentence-BERT (Reimers and Gurevych 2019) with clustering using HDBSCAN (Campello, Moulavi, and Sander 2013).

Our topic discovery method builds upon BERTopic but differs in two significant ways:

1. Instruction-based document embedding, and
2. Hierarchical topic merging based on topic coherence.

These modifications are designed to address specific limitations in the original BERTopic approach and better align with the needs of our study.

Instruction-based Document Embedding In this work, we leverage an instruction-based document embedding model, first introduced by INSTRUCTOR (Su et al. 2023). Unlike S-BERT (Reimers and Gurevych 2019) and SimCSE (Gao, Yao, and Chen 2021), which were designed for generating general-purpose document embeddings, instruction-based embedding models allow more task-specific embeddings by incorporating instructions with the input text (Su et al. 2023).

While general-purpose embedding models are useful for many applications, they have limitations in distinguishing nuanced differences in semantic meaning, especially for tasks requiring an understanding of context or stance. For example, as noted by Introne (2023), general-purpose embedding models often yield high cosine similarity for semantically opposite sentences, such as "Illegal immigrants are causing problems" and "Illegal immigrants are not causing problems." This limitation arises because these models are not explicitly optimized to capture task-specific distinctions, such as differences in sentiment or stance within a given topic.

Instruction-based document embedding models address this issue by jointly taking the input text and an instruction describing the downstream task (Su et al. 2023). This approach enables the model to produce embeddings that align better with the task's requirements.

To investigate this, we experimented with several sentence pairs to evaluate the performance of different models. One example pair was "Illegal immigration helps the U.S. economy by filling jobs and contributing to growth" and "Illegal immigration hurts the U.S. economy by taking jobs and draining resources." The all-mpnet-base-v2 model[4], which was one of the embedding models used in the original BERTopic paper, produced a cosine similarity score of 0.864 for these two sentences. While these sentences are

[3]https://huggingface.co/google/gemma-3-4b-it

[4]huggingface.co/sentence-transformers/all-mpnet-base-v2

similar in terms of topic (illegal immigration), they express opposing stances, highlighting the inability of task-agnostic embedding models to capture differences in stance.

In contrast, an instruction-based embedding model, NV-Embed-v2[5], yielded more nuanced results when provided with different instructions. For instance, when instructed with "What topic is this tweet addressing?", the model produced a cosine similarity score of 95.90 for the two example sentences. Meanwhile, when instructed with "What is this tweet's view on illegal immigration?", the cosine similarity dropped to 71.87. This demonstrates that instruction-based embedding models can effectively distinguish between subtopics and stance when guided by appropriate instructions.

To simultaneously capture both the subtopics of discourse on U.S. illegal immigration and the stance on those subtopics, we constructed the following instruction for embedding: "This is one of the tweets about U.S. immigration issues. What is this user's stance on immigration, and which specific subtopic of immigration does this tweet address?" We chose the NV-Embed-v2 model due to its superior performance on the MTEB benchmark (Muennighoff et al. 2023) as of November 28, 2024, according to the Hugging Face leaderboard.[6] The embedding process took a total of 107 hours and 28 minutes using an NVIDIA RTX A6000 GPU.

Recursive Clustering To cluster the embedded posts, dimensionality reduction was necessary to mitigate the curse of dimensionality (Aggarwal, Hinneburg, and Keim 2001). Following prior literature (Grootendorst 2022), we reduced the dimensionality of the embeddings from 4,096 to 5 dimensions using UMAP (McInnes, Healy, and Melville 2020). UMAP was chosen because it performs well in preserving the global and local structure of data compared to alternatives like t-SNE (Maaten and Hinton 2008) or PCA (Maćkiewicz and Ratajczak 1993), while also being computationally efficient.

Given our dataset of over 4 million vectors, each with 4096 dimensions, it was not computationally feasible to apply UMAP to the entire dataset. Instead, we randomly sampled 10% of the dataset, as this subset was sufficient to capture the global structure of the data and train the UMAP model effectively. The trained model was then used to project the remaining 90% of the dataset, reducing all embeddings to 5-dimensional vectors.

As the final step, we employed the HDBSCAN (Campello, Moulavi, and Sander 2013) algorithm for clustering. HDBSCAN, a density-based clustering algorithm, was selected for its ability to discover clusters of arbitrary shapes and handle noise effectively by not forcing all data points into clusters. Additionally, it offers advantages over DBSCAN (Ester et al. 1996) by handling clusters with varying densities through its hierarchical clustering approach.

We experimented with six parameter sets for `min_cluster_size` and `min_samples`: (100, 200, 300, 1000, 2000, 3000). Parameters of 100 and 200 resulted in over 500 clusters, introducing excessive granularity. A value of 300 caused memory issues, leading to time-out errors. Thus, we proceeded with the parameter sets of 1000, 2000, and 3000. The Euclidean distance metric was used for clustering, as the reduced 5-dimensional space was compact enough for this metric to perform effectively.

Despite these optimizations, initial results from a single pass of HDBSCAN revealed a significant imbalance in clustering. Approximately 90% of the data was assigned to a single dominant cluster, with only a few small clusters capturing the remaining points. This imbalance likely stemmed from HDBSCAN's method of constructing density hierarchies and selecting clusters that are most stable across a range of density thresholds. When a dominant cluster spans a broad range of densities, it tends to absorb points that could otherwise form subclusters, particularly when low-density regions act as bridges between subclusters.

To address this issue, we adopted an iterative clustering strategy. In this approach, we identified all clusters larger than the `min_cluster_size` parameter and ran additional passes of HDBSCAN on each large cluster independently. This recursive clustering allowed the algorithm to focus on narrower density ranges within each large cluster, revealing subclusters that were initially masked. While HDBSCAN is inherently a hierarchical clustering algorithm and its tree structure theoretically captures all possible clusters, it was impractical to directly use the full hierarchy due to the overwhelming number of transient clusters across millions of data points.

Hierarchical Topic Merging Based on Topic Coherence Using the three sets of parameters (`min_cluster_size` = `min_samples` = 1000, 2000, 3000), we obtained different numbers of unique clusters. The maximum subcluster levels for each hyperparameter set were 5, 6, and 5, respectively.

Retaining all subclusters could lead to over-segmentation, as some clusters may arise from minor density fluctuations or noise within the parent cluster. To address this, we aimed to retain only those subclusters that demonstrated statistically significantly higher coherence compared to their parent cluster.

Although several methods exist for evaluating the coherence of topics generated by a topic modeling algorithm, most rely on keywords, which are not directly applicable to our approach. Instead, we adopted a qualitative evaluation method inspired by Newman et al. (2009), wherein human annotators rate the quality of topics based on their coherence. To scale this method, we used the unquantized Llama-3.3-70B-instruct model[7] to evaluate topic coherence automatically.

For the evaluation of each subcluster, we randomly sampled 30 posts from within the subcluster and 30 posts from outside it. These 60 sample posts were provided to Llama-3.3 using a standardized prompt (details in Appendix A). The model returned a single integer score on a 5-point Likert scale, representing the coherence of the topic. Symbolically,

[5]https://huggingface.co/nvidia/NV-Embed-v2

[6]https://huggingface.co/leaderboard

[7]https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct

this classification can be represented as:

$$\text{Coherence} = f\big(x_1, x_2, \ldots, x_{30},\ y_1, \ldots, y_{30}\big),$$

where x_n are in-cluster samples, y_n are out-of-cluster samples, and f represents the classifier constructed by the language model and prompt. To improve statistical reliability, we repeated this sampling and scoring process 30 times for each subcluster, generating a distribution of coherence scores. On average, the 95% confidence interval for the coherence scores across all subclusters was ± 0.08, with the smallest CI being ± 0.00 and the largest ± 0.62.

Once we obtained coherence scores for each subcluster and its parent cluster, we compared the distributions using the Mann-Whitney U test (Nachar 2008). This non-parametric test evaluates differences in median values and allowed us to determine whether the subcluster's coherence was significantly higher than its parent cluster. Subclusters that failed to demonstrate significant improvement or had significantly lower coherence scores were merged back into their parent cluster.

Table 1 summarizes the clustering results after subcluster merging for different minimum cluster sizes.

Minimum Cluster Size	**1000**	**2000**	**3000**
Level 1 Clusters	11	9	9
Level 2 Clusters	157	4	1
Level 3 Clusters	9	–	–

Table 1: Cluster Counts by Level After Subcluster Merging for Different Minimum Cluster Sizes

After qualitative probing of the final clusters, we found that using a minimum cluster size of 1000 resulted in slight over-segmentation. However, this was acceptable compared to the results for 2000 and 3000, which produced overly coarse clusters. Therefore, we proceeded with the results from the `min_cluster_size` of 1000.

Finally, using the Llama-3.3-70B model again, we generated labels for each topic. Labels consisted of a 3–7 word noun phrase summarizing the topic and an approximately 100-word concise description, derived from 30 in-cluster and 30 out-of-cluster sample posts. A comprehensive list of topic labels, full topic descriptions, post counts, and toxicity levels can be found on this interactive visualization page[8] (https://topic-immigration.onrender.com). Appendix C provides the brief topic descriptions instead of full descriptions, also generated by Llama-3.3-70B.

Topic Trajectory Analysis

Based on the discovered topics, we treated them as semantically meaningful regions in the embedding space (as illustrated in Figure 1), where users "visit" topics by posting content associated with them over time. In this sense, each user's visiting history in topic space can be conceptualized as a *topic trajectory*.

[8]The page takes approximately 5 seconds to fully load.

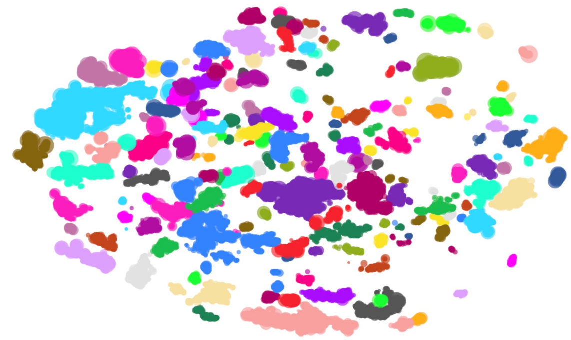

Figure 1: Two-Dimensional Projection of Post Embeddings Belonging to 157 U.S. Level 2 Immigration Subtopics (Color-Coded by Topic)

We focus our analysis on a set of 157 topics identified as Level 2 clusters, which offer a reasonable balance between comprehensiveness and granularity. Although this level of clustering is not perfect, as will be discussed in the Results section, it provides a useful resolution for evaluating which regions of topic space users' trajectories are most closely associated with.

User Grouping We began by identifying *active users* as those who posted at least 50 times on X about U.S. immigration issues within our six-month study window. This yielded 8,180 active users (just 0.68% of the 1,206,512 unique users in our dataset) who together produced 17.86% of all US immigration-related posts.

For each active user i, we modeled their toxicity score over time by fitting a simple linear regression:

$$\text{toxicity}_{it} = \beta_{0,i} + \beta_{1,i}\left(\text{timestamp}_{it}\right) + \varepsilon_{it},$$

where toxicity_{it} is the toxicity of user i's tth post, timestamp_{it} is the post time in seconds, $\beta_{1,i}$ captures the rate of change in toxicity for user i, and ε_{it} is the residual error. We extracted those users for whom the time coefficient $\beta_{1,i}$ was statistically significant ($p < 0.05$), yielding 1,124 users with a clear temporal trend in toxicity. Of these, 718 had $\beta_{1,i} > 0$ (the *Increasing Toxicity Group*) and 406 had $\beta_{1,i} < 0$ (the *Decreasing Toxicity Group*).

To isolate trajectory effects of increasing or decreasing toxicity from baseline toxicity level effects, we created matched *reference groups*. First, we computed each user's average post toxicity, $\bar{T}_i$. The mean of these averages was $\bar{T}_{\text{inc}} = 0.5762$ for the Increasing Toxicity Group and $\bar{T}_{\text{dec}} = 0.5450$ for the Decreasing Toxicity Group. Then, among the 7,056 users without a significant trend, we selected:

- 718 users whose $\bar{T}_i$ was closest to $\bar{T}_{\text{inc}}$, forming the *Reference Group for the Increasing Toxicity Group*, and
- 406 users whose $\bar{T}_i$ was closest to $\bar{T}_{\text{dec}}$, forming the *Reference Group for the Decreasing Toxicity Group*.

This yields four groups—(1) Increasing Toxicity Group, (2) Reference Group for the Increasing Toxicity Group, (3) Decreasing Toxicity Group, and (4) Reference Group for the

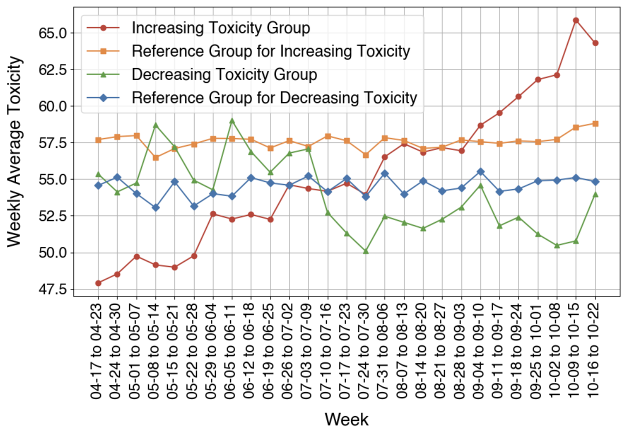

Figure 2: Weekly Average Toxicity of Increasing and Decreasing Toxicity Groups and Their Reference Groups

Decreasing Toxicity Group—allowing us to test whether the topic trajectories of the Increasing and Decreasing Toxicity Groups differ significantly from those of their respective reference groups. Figure 2 shows each group's weekly average toxicity.

Linear Interpolation of Topic Trajectories Each user's trajectory is represented as a temporal sequence of 5-dimensional embedding vectors derived from their posts. Because post timestamps are recorded to the nearest second, we first align all trajectories to a common 194-day grid (the full span of the study) by linearly interpolating between successive post embeddings.

For every post by user u with timestamp s_i $(i = 1, \ldots, N)$ and embedding $\mathbf{e}_u(s_i) \in \mathbb{R}^5$, we compute a *normalised time*

$$\tau_i = \frac{s_i - t_0}{t_{\text{end}} - t_0} \in [0, 1],$$

where t_0 is 00:00 UTC on 17 Apr 2023 and t_{end} is 23:59 UTC on 27 Oct 2023. A linear interpolant $\mathbf{e}_u(\tau)$ is then fitted through the ordered pairs $(\tau_i, \mathbf{e}_u(s_i))$.

Finally, we evaluate this interpolant at the 194 equally spaced grid points

$$\tau_g = \frac{g}{193} \qquad (g = 0, \ldots, 193),$$

which correspond to *noon of each calendar day*. The resulting sequence

$$\{\mathbf{e}_u(\tau_0), \mathbf{e}_u(\tau_1), \ldots, \mathbf{e}_u(\tau_{193})\}$$

is the daily 5-dimensional trajectory for user u. If a user has no post before a particular grid point, we carry the earliest observed embedding backward; similarly, we carry the final embedding forward beyond the user's last post.

A weekly version of each trajectory is obtained by coordinate-wise averaging over non-overlapping 7-day windows; the final five days (23 – 27 Oct) are omitted because they do not complete a full week.

PERMANOVA of Trajectory Pairs To answer RQ1 (whether the Increasing Toxicity Group follows a different topical path from its reference group) and RQ2 (likewise for the Decreasing Toxicity Group), we ran two distance-based permutational MANOVAs (PERMANOVA; Anderson 2001). Our test treats each user's entire interpolated path as one multivariate observation. Formally, we set up the following hypothesis test for each trajectory pair:

1. **H0:** The two group trajectories are *not* significantly different in their sequence of positions in embedding space.
2. **H1:** The two group trajectories are significantly different in their sequence of positions in embedding space.

For user u we flatten the 5-dimensional trajectory obtained in the previous section into a single vector containing the embeddings from either the daily grid ($T = 194$) or the weekly grid ($T = 27$):

$$\mathbf{x}_u = \left[\mathbf{e}_u(\tau_0)^\top, \mathbf{e}_u(\tau_1)^\top, \ldots, \mathbf{e}_u(\tau_{T-1})^\top\right]^\top \in \mathbb{R}^{5T}.$$

Thus each comparison involves two sets of points in $\mathbb{R}^{5T}$: Group A has n_A vectors $\{\mathbf{x}_1, \ldots, \mathbf{x}_{n_A}\}$, Group B has n_B vectors $\{\mathbf{x}_{n_A+1}, \ldots, \mathbf{x}_N\}$ with $N = n_A + n_B$. In our model $n_A = n_B$ (718 vs. 718 for the Increasing Toxicity pair, 406 vs. 406 for the Decreasing Toxicity pair).

We use the ordinary Euclidean distance to form two sums of squares:

$$\text{SS}_{\text{between}} = n_A \left\|\bar{\mathbf{x}}_A - \bar{\mathbf{x}}\right\|_2^2 + n_B \left\|\bar{\mathbf{x}}_B - \bar{\mathbf{x}}\right\|_2^2,$$

$$\text{SS}_{\text{within}} = \sum_{u \in A} \|\mathbf{x}_u - \bar{\mathbf{x}}_A\|_2^2 + \sum_{u \in B} \|\mathbf{x}_u - \bar{\mathbf{x}}_B\|_2^2,$$

where the overall mean is $\bar{\mathbf{x}} = \dfrac{n_A \bar{\mathbf{x}}_A + n_B \bar{\mathbf{x}}_B}{N}$, and the group-mean vectors are $\bar{\mathbf{x}}_A = \frac{1}{n_A} \sum_{u \in A} \mathbf{x}_u$ and $\bar{\mathbf{x}}_B = \frac{1}{n_B} \sum_{u \in B} \mathbf{x}_u$. $\text{SS}_{\text{between}}$ measures the separation of the two group centroids, whereas $\text{SS}_{\text{within}}$ measures the dispersion of trajectories around their own group mean.

The PERMANOVA statistic, *pseudo-F* is

$$F = \frac{\text{SS}_{\text{between}}}{\text{SS}_{\text{within}}/(N-2)},$$

with 1 and $N - 2$ nominal degrees of freedom.

To obtain a p-value at $\alpha = 0.01$ we generated a null distribution of F by randomly reshuffling the user-level group labels 4,999 times, following Anderson (2001).

Average Trajectory by Group Building on the daily and weekly user 5D trajectories obtained through linear interpolation, we computed each group's average trajectory by taking the Euclidean mean of the group's user vectors at each time step. Formally, for a given group G and day d, the group-average embedding is:

$$\bar{\mathbf{e}}_G(d) = \frac{1}{|G|} \sum_{u \in G} \mathbf{e}_u(d),$$

the coordinate-wise centroid of all users in G on day d. In the similar fashion, for a given group G and week w, the group-average embedding is:

$$\bar{\mathbf{e}}_G(w) = \frac{1}{|G|}\sum_{u \in G} \mathbf{e}_u(w),$$

the coordinate-wise centroid of all users in G on week w. Because the embeddings have been projected into a 5-dimensional UMAP space, which preserves important distance relationships, we judged that using Euclidean averages to compute each group's average trajectory is conceptually valid.

Topic Classification of Average Trajectories To convert each group's average trajectory as a sequence of interpretable topics, we applied a nearest-neighbor classifier to every timestep of the trajectory. In particular, we trained a 15-nearest-neighbor (KNN) classifier in the 5-dimensional UMAP embedding space, using cosine similarity as the distance metric for finding neighbors. The classifier's training data consisted solely of posts that were assigned to one of the 157 discovered topics from our Topic Discovery step.

This KNN model proved highly effective, achieving near-perfect performance on held-out test data (macro F1 = 99.93 and micro F1 = 99.95). We used this classifier to label each point along the average trajectories. As a result, each group's average daily and weekly trajectory in embedding space was transformed into a sequence of interpretable topics.

Results

Topic Discovery

Our recursive HDBSCAN procedure uncovered 157 Level-2 subtopics that span the policy, humanitarian, economic, cultural, and conspiratorial aspects of the U.S. immigration conversation (Appendix C).

- *Policy, law and enforcement.* Classical policy arguments appear in topics such as *"Supreme Court Immigration Rulings"* (Topic 29), *"DACA"* (Topic 8), or *"E-Verify"* (Topic 112). Toxicity in this bloc is generally moderate (40–50), indicating relatively civil, though partisan, legal argumentation.
- *Economic and labor frames.* Topics that portray migrants as either vital labor (*"Immigrant Work Ethic"*, Topic 87; *"Labor Shortage and Immigration"*, Topic 115) or unwelcome competitors (*"Hiring of Illegal Immigrants"*, Topic 117) show different levels of toxicity (mean scores 45 vs. 57), illustrating how the same economic lens can lead to opposing narratives.
- *Humanitarian Support.* Humanitarian narratives cluster at the low-toxicity extreme, such as *"Support for Migrant Communities"* (Topic 59, toxicity 30.5) and *"Immigration Support and Resources"* (Topic 78, toxicity 20.3), frequently co-occurring with resource coordination such as housing, legal aid or language access.
- *Hostile nationalism.* The highest toxicity averages are concentrated in "threat" frames that depict immigration as criminal invasion or demographic subversion: *"Replacement Migration Conspiracy"* (Topic 89), *"Anti-Immigration White Nationalism"* (Topic 132), and *"Trump's Anti-Immigrant Rhetoric"* (Topic 94). Keywords here such as *"invasion"* and *"poisoning the blood"* reveal that these clusters are shaped by fear-based narratives and strong anti-immigrant hostility.
- *Identity-focused debates.* Several Level-2 clusters reveal narratives centered on identity-based group tensions, such as *"Black Americans on Immigration"* (Topic 20) and *"Black American vs. Immigrant Dynamics"* (Topic 135). These topics suggest that conflicts within minority communities also play a role in shaping discourse around the U.S. immigration issue.

Frequency	Group Pair	Pseudo-F	p	η^2
Daily	Increasing vs. Reference	1.898	.023	.001
	Decreasing vs. Reference	1.424	.045	.002
Weekly	Increasing vs. Reference	2.211	.023	.002
	Decreasing vs. Reference	1.619	.040	.002

Table 2: PERMANOVA results comparing 5-D topic trajectories of toxicity-changing groups with toxicity-matched reference groups at daily (T=194) and weekly (T=27) frequencies. p-values are based on 4,999 permutations.

Label collisions, such as the presence of three distinct *"Illegal Immigration Debate"* clusters (Topics 66, 146, and 155), might appear redundant. However, close reading reveals subtle differences in stance. (Full descriptions of all topics are accessible at https://topic-immigration.onrender.com). Topic 66 focuses on disputes concerning the distinction between asylum seekers and illegal immigrants. Topic 146 centers on debates over the usage of the term "illegal immigrants", whereas Topic 155 primarily includes emotionally charged condemnations. These nuanced differences among seemingly similar topics demonstrate that they can be further divided by their rhetorical emphasis.

In four high-traffic Level 2 subtopics, our iterative clustering process revealed a third layer of subtopics. One illustrative example is *"NYC Migrant Housing Crisis"* (Topic 92), which was split into three subtopics. Although these separate clusters were retained due to their coherence scores being statistically higher than that of their parent topic, their descriptions show substantial overlap. Therefore, the semantic distinctions between these Level 3 clusters may have been difficult to detect using our labeling approach.

Topic Trajectory Analysis Results

PERMANOVA Test Results To answer RQ1 and RQ2, we used PERMANOVA to compare the average topic trajectories of the Increasing Toxicity Group and the Decreasing Toxicity Group to those of their respective toxicity-matched reference groups.

Table 2 reports the results for both temporal granularities. All four comparisons reject the null hypothesis at the $\alpha = 0.05$ level, answering our two research questions.

- **RQ1 (Increasing Toxicity Group):** The Increasing Toxicity Group exhibited significantly different trajectories from its toxicity-matched reference group.
- **RQ2 (Decreasing Toxicity Users):** The Decreasing Toxicity Group also followed significantly different trajectories from its matched reference group.

The group effect sizes turned out to be small ($\eta^2 \leq .002$), indicating that group differences account for only 0.2% of the total dispersion, which is unsurprising given the presence of other sources of variation driving topic trajectories. Also, consistent pattern across temporal frequencies underscores that the reliable trajectory separation is not due to high-frequency noise.

Comparison of Trajectory Pairs In line with the statistical evidence (Table 2), we find clear qualitative differences in the average topic trajectories of user groups whose toxicity is increasing or decreasing compared to their respective stable counterparts. Below, we compare the weekly topic trajectories of the Increasing Toxicity Group versus its Reference Group (Table 17 in Appendix D), and of the Decreasing Toxicity Group versus its Reference Group (Table 18 in Appendix D), focusing on key transitions and divergences.

Increasing Toxicity Group vs. Reference Group By comparing weekly topic trajectories shown in Table 17 (Appendix D), we observe the Increasing Toxicity Group diverging notably from its Reference Group around late summer. In August, the Increasing Toxicity Group's average toxicity rises sharply (Figure 2), coinciding with shifts to more threat-oriented topics. For example, in the week of 7/31–8/6, this group gravitated to *"Secret Flights of Illegal Immigrants"* (Topic 39, toxicity 56.67), a conspiracy-tinged topic implying covert government actions, whereas the Reference Group remained on a more conventional *"155. Illegal Immigration Debate"* (Topic 155, toxicity 56.78). A few weeks later, from August 28 to September 3, the Increasing Toxicity Group's average trajectory shifted to *"U.S. Immigration Concerns"* (Topic 147, toxicity 61.67) and stayed there for the next five weeks, reflecting heightened fears about immigration.

By mid-September, Increasing Toxicity Group's average toxicity surpassed the Reference Group's (Figure 2), aligning with their emphasis on threat narratives. In contrast, the Reference Group's average trajectory during these weeks showed little departure from its April pattern, repeatedly returning to one of its most common topics, *"Anti-Immigration Sentiment"* (Topic 127, toxicity 70.58). This pattern addresses RQ1, suggesting that users whose toxicity increased shifted from early, sometimes humanitarian topics such as *"Temporary Protected Status for Immigrants"* (Topic 22, toxicity 30.18) to more alarmist themes, exemplified by *"Immigration and National Security"* (Topic 138, toxicity 68.13), during the critical August–September period, whereas their Reference Group maintained a steadier trajectory.

Decreasing Toxicity Group vs. Reference Group In contrast, the Decreasing Toxicity Group's trajectory shows a shift toward less inflammatory topics (Table 18 in Appendix D). In late spring 2023, this group occasionally engaged with the highly toxic and emotive topic *"Ilhan Omar Immigration Controversy"* (Topic 5, toxicity 73.93). However, by midsummer the Decreasing Toxicity Group began to concentrate on procedure-focused discussions. During the week of July 3 to July 9 its average trajectory visited *"Discrimination and Immigration Laws"* (Topic 37, toxicity 52.73), while the Reference Group showed a similar pattern with *Merit-Based Immigration Debate"* (Topic 84, toxicity 43.18). Over the next three weeks the Decreasing Toxicity Group remained in comparatively low-toxicity areas, including *"Legal vs Illegal Immigration"* (Topic 96, toxicity 47.28), *"Asylum Seekers and Immigration"* (Topic 76, toxicity 47.88), and *"Immigration Detention Criticism"* (Topic 72, toxicity 39.79).

Corresponding with these topical shifts, the Decreasing Toxicity Group's average toxicity fell below that of the Reference Group by mid-July (Figure 2) and generally remained lower thereafter. While the Reference Group's average trajectory periodically visited the highly toxic topic *"U.S. Immigration Politics and Controversy"* (Topic 23, toxicity 71.11), the Decreasing Toxicity Group lingered on policy- and legal-focused themes such as *"Supreme Court Immigration Rulings"* (Topic 29, toxicity 31.79). This suggests that members of the Decreasing Toxicity Group increasingly framed immigration as an institutional issue to be resolved through formal processes rather than as a partisan battle. By contrast, the Reference Group continued to gravitate toward politically charged frames, exemplified by *"Republican Rhetoric on Immigration"* (Topic 88, toxicity 52.67) and *Immigration and Social Justice Issues"* (Topic 101, toxicity 69.61). This pattern addresses RQ2: users whose toxicity decreased followed a distinct topical trajectory, an overall pivot to policy-oriented discussions.

Discussion

Contribution

Academic Contribution First, this study bridges topic-centered and user-centered toxicity research. By integrating user-focused analyses with a robust subtopic discovery approach, we highlight how specific subtopics can be associated with distinct forms of toxic behavior. This addresses the gap between topic-oriented toxicity research and user-level behavioral studies, offering a more holistic view of the dynamics of incivility in social media discourse.

Second, we empirically examine the longitudinal co-evolution of topic engagement and toxic communication. By modeling users' trajectories across subtopics over time, we demonstrate that rising or falling toxicity can be closely linked to shifts in their content focus. Instead of viewing users' toxicity as a fixed attribute, we show how subtopic trajectories can illuminate pathways leading to intensified hostility or redirecting discussions toward policy debates.

Lastly, this study advances methodology on two fronts. First, we pair instruction-based embeddings with recursive HDBSCAN and coherence-guided hierarchical merging to extract stable, interpretable subtopics from massive social-media corpora about complex social issues. Second, our trajectory-analysis framework — which treats each user's posting history as a continuous path through embedding space and compares groups using PERMANOVA — offers a replicable template for future research on the dynamics of social media discourse.

Implications on Policy Communication Toxic discourse on immigration can foster societal divisions and spread misinformation about immigrants' economic and social contributions. However, it would be reductive to label large groups of users as "toxic" in a static sense, as this would obscure the more nuanced narrative shaping public opinion. Negativity in a policy discussion often arises in response to specific triggers—such as unmet policy needs, perceived inequities, or anxieties about future scenarios. Some individuals may start with moderate concerns about an immigration policy but escalate toward more toxic language when they feel their worries are dismissed or when policy outcomes fail to materialize as promised. In that sense, negativity can be a dynamic signal of policy-related frustration or dissatisfaction. This study can help policy communicators recognize that rising toxicity may be intertwined with specific policy-related frustrations, enabling them to devise strategies to de-escalate hostility and steer conversations back toward solution-oriented dialogue.

Limitations and Ethical Considerations

The dataset used for this study consists of posts in English, excluding discourse from users who primarily communicate in other languages. This introduces a potential bias by restricting the analysis to English-speaking populations, which may not fully represent broader discourse on U.S. immigration.

Second, our workflow extensively leveraged large language models, including zero-shot filtering of relevant posts, automated toxicity scoring, and LLM-based topic coherence scoring and labeling. While these models significantly reduce the cost of processing millions of posts, they also introduce errors and may reflect linguistic or cultural biases. The future study pipeline should therefore incorporate more systematic audits of misclassifications and robustness checks using alternative models. Expanding the proportion of human-validated ground truth at each stage will also help address the issue of bias.

The potential negative societal impacts of this research include reinforcing the stigmatization of certain user groups. To mitigate this risk, all identifiers, including user IDs, remain within the scope of the study and are not publicly disclosed, ensuring anonymity. No additional personally identifiable information was collected or used in the research. Furthermore, the results are framed to emphasize understanding and addressing toxicity constructively, rather than singling out specific user groups.

Potential misuse of this work includes justifying censorship, discrimination, or punitive measures against specific user groups based on their association with certain topics or levels of toxicity. To address these concerns, we reiterate that the purpose of this research is to foster an understanding of discourse dynamics on social issues and their implications for policy communications.

Finally, although users and topical factors may effectively inform us insights in creating healthy conversation and make potentially proactive interventions, there exist also other factors such as platform-level factors (e.g., (DiCicco et al. 2020)) at play here. Future research should explore further the connection between toxicity patterns, its antecedents (e.g., (Shen et al. 2020)), and outcomes.

Conclusion

In summary, our large-scale, six-month analysis of U.S.-immigration discourse on X shows that toxicity is tightly intertwined with how users move through the issue space. When people's toxicity climbs, their posts shift toward alarmist and conspiracy-tinged subtopics, whereas falling toxicity is paired with a pivot to procedure- and policy-oriented themes. These findings highlight the fluid, topic-dependent character of incivility around contentious social issues and point to the limits of treating "toxic users" as a fixed category. In addition, the analytical toolkit we introduce, which features hierarchical topic discovery and a topic trajectory analysis pipeline, provides a replicable template for future research and for practical efforts to cultivate healthier online public spheres.

References

Aggarwal, C. C.; Hinneburg, A.; and Keim, D. A. 2001. On the Surprising Behavior of Distance Metrics in High Dimensional Space. In Van den Bussche, J.; and Vianu, V., eds., *Database Theory — ICDT 2001*, 420–434. Berlin, Heidelberg: Springer. ISBN 978-3-540-44503-6.

Anderson, M. J. 2001. A new method for non-parametric multivariate analysis of variance. *Austral Ecology*, 26(1): 32–46. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1442-9993.2001.01070.pp.x.

Angelov, D. 2020. Top2Vec: Distributed Representations of Topics. ArXiv:2008.09470 [cs, stat].

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3: 993–1022.

Bruns, A.; and Highfield, T. 2015. Is Habermas on Twitter?: Social media and the public sphere. In *The Routledge companion to social media and politics*, 56–73. Routledge.

Campello, R. J. G. B.; Moulavi, D.; and Sander, J. 2013. Density-Based Clustering Based on Hierarchical Density Estimates. In Pei, J.; Tseng, V. S.; Cao, L.; Motoda, H.; and Xu, G., eds., *Advances in Knowledge Discovery and Data Mining*, 160–172. Berlin, Heidelberg: Springer. ISBN 978-3-642-37456-2.

Chen, Y.; and Wang, L. 2022. Misleading political advertising fuels incivility online: A social network analysis of 2020 US presidential election campaign video comments on YouTube. *Computers in Human Behavior*, 131: 107202. Publisher: Elsevier.

Coe, K.; Kenski, K.; and Rains, S. A. 2014. Online and uncivil? Patterns and determinants of incivility in newspaper website comments. *Journal of communication*, 64(4): 658–679. Publisher: Oxford University Press.

Cohen, J. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1): 37–46. Publisher: SAGE Publications Inc.

de Wynter, A.; Watts, I.; Wongsangaroonsri, T.; Zhang, M.; Farra, N.; Altıntoprak, N. E.; Baur, L.; Claudet, S.; Gajdušek, P.; and Gu, Q. 2025. RTP-LX: Can LLMs Evaluate Toxicity in Multilingual Scenarios? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 27940–27950. Issue: 27.

DiCicco, K.; Noor, N. B.; Yousefi, N.; Maleki, M.; Spann, B.; and Agarwal, N. 2020. Toxicity and Networks of COVID-19 discourse communities: a tale of two social media platforms. *Proceedings http://ceur-ws. org ISSN*, 1613: 0073.

Ester, M.; Kriegel, H.-P.; Sander, J.; and Xu, X. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, 226–231. Portland, Oregon: AAAI Press.

FORCE11. 2020. The FAIR Data principles. https://force11.org/info/the-fair-data-principles/.

Gallacher, J. D.; Heerdink, M. W.; and Hewstone, M. 2021. Online engagement between opposing political protest groups via social media is linked to physical violence of offline encounters. *Social Media+ Society*, 7(1): 2056305120984445. Publisher: SAGE Publications Sage UK: London, England.

Gao, T.; Yao, X.; and Chen, D. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 6894–6910. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.

Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.

Gervais, B. T.; Dye, C.; and Chin, A. 2025. Incivility or Invalidity? Evaluating Perspective API Scores as a Measure of Political Incivility. *American Politics Research*, 53(3): 266–274. Publisher: SAGE Publications Inc.

Gottfried, J. 2024. Americans' Social Media Use.

Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. ArXiv:2203.05794 [cs].

Hanu, L.; and Unitary team. 2020. Detoxify. Github. https://github.com/unitaryai/detoxify.

Hickey, D.; Schmitz, M.; Fessler, D.; Smaldino, P. E.; Muric, G.; and Burghardt, K. 2023. Auditing Elon Musk's Impact on Hate Speech and Bots. *Proceedings of the International AAAI Conference on Web and Social Media*, 17: 1133–1137.

Introne, J. 2023. Measuring Belief Dynamics on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, 17: 387–398.

Jones, J. M. 2024. Sharply More Americans Want to Curb Immigration to U.S. Section: Politics.

Kim, J. W.; Guess, A.; Nyhan, B.; and Reifler, J. 2021. The Distorting Prism of Social Media: How Self-Selection and Exposure to Incivility Fuel Online Comment Toxicity. *Journal of Communication*, 71(6): 922–946.

Klein, M.; and Majdoubi, N. 2024. The medium is the message: toxicity declines in structured vs unstructured online deliberations. *World Wide Web*, 27(3): 31.

Kumar, D.; Hancock, J.; Thomas, K.; and Durumeric, Z. 2022. Understanding Longitudinal Behaviors of Toxic Accounts on Reddit. *arXiv preprint arXiv:2209.02533*.

Maaten, L. v. d.; and Hinton, G. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86): 2579–2605.

Maćkiewicz, A.; and Ratajczak, W. 1993. Principal components analysis (PCA). *Computers & Geosciences*, 19(3): 303–342.

McInnes, L.; Healy, J.; and Melville, J. 2020. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. ArXiv:1802.03426 [stat].

Mittos, A.; Zannettou, S.; Blackburn, J.; and Cristofaro, E. D. 2020. Analyzing genetic testing discourse on the web through the lens of Twitter, Reddit, and 4chan. *ACM Transactions on the Web (TWEB)*, 14(4): 1–38. Publisher: ACM New York, NY, USA.

Muennighoff, N.; Tazi, N.; Magne, L.; and Reimers, N. 2023. MTEB: Massive Text Embedding Benchmark. In Vlachos, A.; and Augenstein, I., eds., *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 2014–2037. Dubrovnik, Croatia: Association for Computational Linguistics.

Mukherjee, S.; and Krogstad, J. M. 2024. Trump and Harris supporters differ on mass deportations but favor border security, high-skilled immigration. Accessed: 2024-10-24.

Nachar, N. 2008. The Mann-Whitney U: A Test for Assessing Whether Two Independent Samples Come from the Same Distribution. *Tutorials in Quantitative Methods for Psychology*, 4(1): 13–20.

Newman, D.; Karimi, S.; and Cavedon, L. 2009. External evaluation of topic models. In *Proceedings of the 14th Australasian Document Computing Symposium*, 1–8. University of Sydney.

Pradel, F.; Zilinsky, J.; Kosmidis, S.; and Theocharis, Y. 2024. Toxic speech and limited demand for content moderation on social media. *American Political Science Review*, 1–18. Publisher: Cambridge University Press.

Qayyum, H.; Zi Hao Zhao, B.; Wood, I.; Ikram, M.; Kourtellis, N.; and Ali Kaafar, M. 2023. A longitudinal study of the top 1% toxic Twitter profiles. In *Proceedings of the 15th ACM Web Science Conference 2023*, 292–303.

Rajadesingan, A.; Resnick, P.; and Budak, C. 2020. Quick, Community-Specific Learning: How Distinctive Toxicity

Norms Are Maintained in Political Subreddits. *Proceedings of the International AAAI Conference on Web and Social Media*, 14: 557–568.

Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3982–3992. Hong Kong, China: Association for Computational Linguistics.

Rossini, P.; and Maia, R. 2021. Characterizing disagreement in online political talk: Examining incivility and opinion expression on news websites and Facebook in Brazil. *Journal of Deliberative Democracy*, 17(1): 90–104. Publisher: University of Westminster Press.

Salminen, J.; Sengün, S.; Corporan, J.; Jung, S.-g.; and Jansen, B. J. 2020. Topic-driven toxicity: Exploring the relationship between online toxicity and news topics. *PLOS ONE*, 15(2): e0228723. Publisher: Public Library of Science.

Santana, A. D. 2015. Incivility dominates online comments on immigration. *Newspaper research journal*, 36(1): 92–107. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Shen, C.; Sun, Q.; Kim, T.; Wolff, G.; Ratan, R.; and Williams, D. 2020. Viral vitriol: Predictors and contagion of online toxicity in World of Tanks. *Computers in Human Behavior*, 108: 106343. Publisher: Elsevier.

Stromer-Galley, J.; Bryant, L.; and Bimber, B. 2015. Context and medium matter: Expressing disagreements online and face-to-face in political deliberations. *Journal of Deliberative Democracy*, 11(1). Publisher: University of Westminster Press.

Su, H.; Shi, W.; Kasai, J.; Wang, Y.; Hu, Y.; Ostendorf, M.; Yih, W.-t.; Smith, N. A.; Zettlemoyer, L.; and Yu, T. 2023. One Embedder, Any Task: Instruction-Finetuned Text Embeddings. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023*, 1102–1121. Toronto, Canada: Association for Computational Linguistics.

Sydnor, E. 2019. Signaling Incivility: The Role of Speaker, Substance, and Tone. In *A Crisis of Civility?* Routledge. ISBN 978-1-351-05198-9. Num Pages: 20.

Thompson, L.; and Mimno, D. 2020. Topic Modeling with Contextualized Word Representation Clusters. ArXiv:2010.12626 [cs].

Vayansky, I.; and Kumar, S. A. P. 2020. A review of topic modeling methods. *Information Systems*, 94: 101582.

Wu, X.; Nguyen, T.; and Luu, A. T. 2024. A Survey on Neural Topic Models: Methods, Applications, and Challenges. ArXiv:2401.15351 [cs].

Yang, A. 2020. The issue niche theory of nongovernmental and nonprofit organizations' interorganizational network ecology. *Communication Theory*, 30(1): 41–63. Publisher: Oxford University Press.

Zhang, Z.; Fang, M.; Chen, L.; and Namazi Rad, M. R. 2022. Is Neural Topic Modelling Better than Clustering? An Empirical Study on Clustering with Contextual Embeddings for Topics. In Carpuat, M.; de Marneffe, M.-C.; and Meza Ruiz, I. V., eds., *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3886–3893. Seattle, United States: Association for Computational Linguistics.

Ziems, C.; Held, W.; Shaikh, O.; Chen, J.; Zhang, Z.; and Yang, D. 2024. Can Large Language Models Transform Computational Social Science? *Computational Linguistics*, 50(1): 237–291.

Paper Checklist

1. For most authors...
   (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? Yes
   (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes
   (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes
   (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? Yes, see the Data and Limitations and Ethical Considerations section.
   (e) Did you describe the limitations of your work? Yes, see the Limitations and Ethical Considerations section.
   (f) Did you discuss any potential negative societal impacts of your work? Yes, see the Limitations and Ethical Considerations section.
   (g) Did you discuss any potential misuse of your work? Yes, see the Limitations and Ethical Considerations section.
   (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes, see the Limitations and Ethical Considerations section.
   (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes
2. Additionally, if your study involves hypotheses testing...
   (a) Did you clearly state the assumptions underlying all theoretical results? NA
   (b) Have you provided justifications for all theoretical results? NA
   (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? NA

(d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? NA

(e) Did you address potential biases or limitations in your theoretical framework? NA

(f) Have you related your theoretical results to the existing literature in social science? NA

(g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? NA

3. Additionally, if you are including theoretical proofs...

(a) Did you state the full set of assumptions of all theoretical results? NA

(b) Did you include complete proofs of all theoretical results? NA

4. Additionally, if you ran machine learning experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? Yes, detailed instructions for the analysis are provided in the Methodologies section and Appendix. However, the social media data will not be released due to privacy concerns.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? Yes, see the Methodologies section.

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? Yes, see the Methodologies section.

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? Yes, see the Methodologies section.

(e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? Yes, see the Methodologies section.

(f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? Yes, see the Methodologies section.

5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...

(a) If your work uses existing assets, did you cite the creators? Yes, see the Data and Methodologies section.

(b) Did you mention the license of the assets? Yes, see the Methodologies section.

(c) Did you include any new assets in the supplemental material or as a URL? NA

(d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? No, because the dataset used in this study consists of publicly available posts from X. As such, no explicit consent was obtained from individual users.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, see the Data, Methodologies and Limitations and Ethical Considerations section.

(f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? NA

(g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? NA

6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...

(a) Did you include the full text of instructions given to participants and screenshots? NA

(b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA

(d) Did you discuss how data is stored, shared, and deidentified? NA

Appendix A: Prompt for Topic Coherence Rating

Listing 1: Prompt for Topic Coherence Rating

```
Task Description:
You are a computational social scientist
     evaluating the coherence of a
    specific topic ("Topic A") ...

Evaluation Process:
1. Examine two sets of tweets:
   - In-topic examples: 30 randomly
       selected tweets classified as
       belonging to Topic A.
     {in_topic_examples}
   - Out-topic examples: 30 randomly
       selected tweets classified as NOT
       belonging to Topic A.
     {out_topic_examples}

2. Rate the coherence of Topic A on a 5-
    point Likert scale:
   - 5: Highly coherent.
   - 4: Moderately coherent.
   - 3: Neutral.
   - 2: Somewhat incoherent.
   - 1: Highly incoherent.

Note:
Relevance to U.S. immigration alone does
     not imply coherence. Evaluate
    distinctness of subtopics.

Output Format:
Coherence: {{coherence rate}}
```

Appendix B: Prompt for Toxicity Detection

Listing 2: Prompt for Toxicity Score Annotation

```
Task Description:
You will be given a "text" in your language.
Your task is to assign a holistic toxicity score on a 5-point Likert scale to the text.
Toxicity can take different forms, such as identity attack, bias, microaggression, violence, self-harm, insult, or sexual content.

Note:
- This dataset may contain extremely toxic content.
- Evaluate the text based on how toxic it would be perceived by a native speaker of the language.
- The text may be truncated.

Input:
Text: "{text}"

Output Format:
Respond with a single integer toxicity score [1-5]. Do not include anything else.
```

Appendix C: Detailed Topic Discovery Results

See Tables 3–16.

Appendix D: Interpretable Weekly Topic Trajectories

See Table 17 and 18.

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
Outliers					94,558
1			Anti-Immigrant Sentiment		29,902
			The topic highlights hostile, divisive attitudes toward undocumented immigrants, advocating stricter border control and portraying them as threats.		
2			Immigrant Experiences and Stories		23,481
			This topic highlights tweets sharing immigrant experiences, challenges, triumphs, cultural diversity, and promoting empathy and understanding of immigration.		
3			Undocumented Immigration Issues		10,538
			The topic addresses undocumented immigration in the U.S., highlighting challenges, economic and societal impacts, and differing perspectives on policy.		
4			Migration and Immigration Issues		14,685
			This topic explores migration, its societal impact, policies, border control, economic growth, social integration, and diverse perspectives on immigration.		
5			European Migrant Crisis		3,726
			The European migrant crisis involves large-scale migration to Europe, impacting nations and sparking political, social, and community challenges.		
6			Israeli-Palestinian Conflict		4,405
			The topic focuses on the Israeli-Palestinian conflict, highlighting illegal settlements, Palestinian rights, and criticism of U.S. support for Israel.		
7			Anti-Immigrant Sentiment and Hate		4,930
			This topic highlights xenophobia, bigotry, and negative rhetoric toward immigrants, often involving derogatory language, generalizations, and political divisiveness.		
8			Immigration Fraud and Scams		2,352
			The topic addresses U.S. immigration fraud, including fake marriages, document forgery, scams, and calls for stricter enforcement and verification.		
9			Economic Migration Debate		6,165
			The topic discusses the polarized debate on economic migrants, their impact, immigration policies, and the distinction from refugees.		
10			Deportation of Undocumented Immigrants		2,019
			The topic focuses on calls for deporting undocumented immigrants, featuring divisive, urgent, and critical opinions on stricter immigration enforcement.		
11			U.S. Immigration Debate		4,454,514
			The U.S. immigration debate covers border control, illegal immigration, migrant rights, policy impacts, and calls for comprehensive reform.		
	11-1 (Topic 1)		Human Trafficking and Immigration	64.92	1,126
			The topic highlights human trafficking, illegal immigration, migrant exploitation, border risks, and criticism of U.S. government policies.		

Table 3: Detailed Topic Discovery Results 1

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-2 (Topic 2)		Anti-Immigration Sentiment in the US	63.10	1,536
			The topic covers tweets opposing U.S. immigration, linking it to economic issues, favoring Trump's policies, and citing a related book.		
	11-3 (Topic 3)		Immigration and Public Health Concerns	64.32	20,664
			The topic explores debates on U.S. immigration policies' impact on public health, focusing on disease spread, vaccinations, and political discourse.		
	11-4 (Topic 4)		Saudi Border Guard Atrocities	57.48	2,992
			The topic highlights alleged killings of African migrants by Saudi border guards and critiques the U.S. response to these abuses.		
	11-5 (Topic 5)		Ilhan Omar Immigration Controversy	73.93	3,282
			The topic involves allegations of immigration fraud against Ilhan Omar, sparking polarized debates on her citizenship and calls for deportation.		
	11-6 (Topic 6)		Civics Test for Voting	50.33	2,379
			The topic explores requiring U.S. citizens to pass a civics test to vote, sparking debates on informed voting and barriers.		
	11-7 (Topic 7)		Abuse of Migrant Populations	67.13	1,105
			The topic highlights allegations of sexual misconduct by law enforcement against vulnerable migrants, exposing power dynamics and institutional failures.		
	11-8 (Topic 8)		DACA Immigration Policy Debate	35.90	2,394
			The topic discusses the DACA program, its legality, impact on "Dreamers," and the political debates surrounding U.S. immigration policy.		
	11-9 (Topic 9)		Biden Immigration Loan Policies	55.43	1,708
			The topic discusses Biden administration policies pressuring banks to lend to illegal immigrants, sparking debate over financial access and risks.		
	11-10 (Topic 10)		Illegal Immigration in India	68.56	1,218
			The topic discusses concerns over illegal immigration in India, focusing on impacts, government handling, and calls for stricter border control.		
	11-11 (Topic 11)		Texas Migrant Shelter Attack	59.35	2,954
			An SUV intentionally rammed into migrants in Brownsville, Texas, sparking outrage, debates on immigration policies, and calls for accountability.		
	11-12 (Topic 12)		Illegal Immigrants as Police Officers	62.20	4,245
			The debate over allowing illegal immigrants to become police officers in Illinois sparks concerns about law, security, and political polarization.		
	11-13 (Topic 13)		Vatican's Stance on Immigration	57.28	1,404
			The topic critiques Pope Francis' stance on U.S. immigration, questioning the Vatican's actions and accusing him of hypocrisy.		
	11-14 (Topic 14)		Birthright Citizenship Debate	57.51	12,634
			The debate on U.S. birthright citizenship focuses on granting automatic citizenship to children of undocumented immigrants, referencing the 14th Amendment.		

Table 4: Detailed Topic Discovery Results 2

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
		11-14-1	Chain Migration Debate		1,108
			The topic discusses chain migration, its role in U.S. immigration policy, and debates over political hypocrisy and double standards.		
		11-14-2	Birthright Citizenship Debate		6,965
			The topic discusses the debate over birthright citizenship for children of undocumented immigrants, focusing on the 14th Amendment and immigration policy.		
	11-15 (Topic 15)		Ted Cruz Immigration Hypocrisy	62.98	2,227
			Criticism of Senator Ted Cruz for anti-immigrant rhetoric, highlighting perceived hypocrisy given his Cuban immigrant father's background.		
	11-16 (Topic 16)		Satirical Views on Immigration Benefits	69.99	1,633
			This topic critiques U.S. immigration policies through satirical tweets, claiming illegal immigrants receive unfair benefits unavailable to citizens.		
	11-17 (Topic 17)		Driver's Licenses for Immigrants	32.65	1,552
			The topic discusses granting driver's licenses to undocumented immigrants in the US, highlighting benefits, state policies, and political debates.		
	11-18 (Topic 18)		Immigration and Voter Fraud	62.31	57,998
			The topic discusses claims of illegal immigration enabling voter fraud, influencing elections, and sparking polarized debates on immigration and election integrity.		
	11-19 (Topic 19)		Replacement of American Citizens	71.91	2,044
			The topic discusses a conspiracy theory claiming Democrats intentionally replace American citizens with immigrants to shift demographics and voting power.		
	11-20 (Topic 20)		Black Americans on Immigration	71.89	18,230
			Black Americans express concerns about illegal immigration's impact on their communities, citing economic competition, political neglect, and social inequality.		
	11-21 (Topic 21)		U.S. Immigration Deportation Policies	35.85	1,576
			The topic covers U.S. deportation policies, focusing on resumed flights, Mexico agreements, and their political and humanitarian impacts.		
	11-22 (Topic 22)		Temporary Protected Status for Immigrants	30.18	2,226
			The U.S. granted Temporary Protected Status to Venezuelan migrants, enabling legal work, deportation protection, and sparking diverse stakeholder reactions.		
	11-23 (Topic 23)		U.S. Immigration Politics and Controversy	71.11	4,186
			The topic critiques Governor Ron DeSantis' immigration policies, linking them to broader issues of discrimination, human rights, and social justice.		
	11-24 (Topic 24)		U.S. Immigration Policy Debate	40.33	13,857
			The topic discusses U.S. immigration policies, focusing on Title 42's expiration, migrant influx, border impacts, and polarized opinions.		
	11-25 (Topic 25)		Climate Change and Migration	45.72	9,126
			The topic explores how climate change drives migration, emphasizing policy action, resource strain, and the need for climate justice.		

Table 5: Detailed Topic Discovery Results 3

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-26 (Topic 26)		Court Rulings on Immigration Policy	34.02	1,931
			The topic covers court rulings impacting Biden's immigration policies, sparking debate over checks on power versus immigration reform and security.		
	11-27 (Topic 27)		Support for Undocumented Students	27.34	1,013
			The topic focuses on support, resources, and advocacy for undocumented students in U.S. higher education, promoting inclusivity and empowerment.		
	11-28 (Topic 28)		U.S. Immigration Controversies	61.77	6,882
			The topic highlights controversies over U.S. immigration, focusing on family separations, detention centers, policy criticism, and calls for accountability.		
	11-29 (Topic 29)		Supreme Court Immigration Rulings	31.79	1,686
			The topic covers Supreme Court rulings on Biden administration immigration policies, prioritizing deportations of public safety risks and border cases.		
	11-30 (Topic 30)		Migrant Work Authorization	34.93	2,561
			The topic highlights the urgent need for expedited work permits for migrants, emphasizing economic benefits and political advocacy for reforms.		
	11-31 (Topic 31)		Healthcare for Undocumented Immigrants	32.39	3,108
			The topic discusses healthcare access for undocumented immigrants in the U.S., including policies, challenges, and advocacy for universal healthcare.		
	11-32 (Topic 32)		Immigration Moratorium Advocacy	50.25	1,299
			The topic discusses calls for an immigration moratorium in the U.S., citing job competition, cultural shifts, and national security concerns.		
	11-33 (Topic 33)		Reagan's Immigration Amnesty Policy	50.24	1,235
			The topic discusses Reagan's 1986 Immigration Reform, its amnesty for 3 million immigrants, and its impact on current immigration debates.		
	11-34 (Topic 34)		Amnesty for Illegal Immigrants	52.52	2,115
			The topic discusses the contentious debate on granting amnesty to illegal immigrants, highlighting political divisions and varying public opinions.		
	11-35 (Topic 35)		U.S. Immigration Policy Debate	53.13	3,185
			The topic discusses U.S. immigration policies, focusing on illegal immigration, border control, economic impacts, and polarized opinions on leadership.		
	11-36 (Topic 36)		Anti-Immigration Border Control	63.36	3,036
			The topic highlights opposition to illegal immigration, advocating border wall construction, strict deportation, and aggressive border control for national security.		
	11-37 (Topic 37)		Discrimination and Immigration Laws	52.73	1,948
			This topic discusses polarized views on discrimination in U.S. immigration, focusing on legality, race, ethnicity, and national origin.		
	11-38 (Topic 38)		U.S. Immigration Controversy	66.62	1,829
			The topic addresses criticism of Texas Governor Greg Abbott's comment labeling mass shooting victims as "illegal immigrants," sparking immigration debate.		

Table 6: Detailed Topic Discovery Results 4

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-39 (Topic 39)		Secret Flights of Illegal Immigrants	56.67	1,469
			The topic discusses alleged secret flights of illegal immigrants under the Biden administration, sparking concerns over transparency and national security.		
	11-40 (Topic 40)		Christian Views on Immigration	55.72	7,410
			The topic explores Christianity's role in U.S. immigration debates on social media, citing biblical arguments and highlighting political and moral divides.		
	11-41 (Topic 41)		DeSantis Immigration Policy Debate	42.36	1,263
			The topic compares immigration policies of Ron DeSantis and Donald Trump, highlighting debates on effectiveness, political implications, and Republican stances.		
	11-42 (Topic 42)		Migrant Shipwrecks and Public Response	53.12	20,036
			The topic highlights outrage over unequal media coverage and empathy for migrant shipwrecks compared to wealthy individuals' tragedies.		
	11-43 (Topic 43)		Immigration and Military Service	60.51	4,762
			The topic explores immigrants, including illegal ones, serving in the US military to gain citizenship, sparking diverse political and social views.		
	11-44 (Topic 44)		U.S. Immigration Policy History	45.24	1,612
			This topic discusses the historical impact, intentions, and debates surrounding the U.S. Immigration and Nationality Act of 1965.		
	11-45 (Topic 45)		Immigration Impact on Public Schools	62.75	1,002
			The topic highlights concerns about immigration's impact on U.S. public schools, including strained resources, lowered standards, and financial burdens.		
	11-46 (Topic 46)		Japan's Immigration Policies	51.86	2,201
			The topic discusses Japan's strict immigration policies, comparisons to Western countries, and debates on cultural preservation versus economic needs.		
	11-47 (Topic 47)		Illegal Economic Migration Debate	60.51	1,358
			The topic focuses on opposition to illegal economic migrants, highlighting concerns about economy, crime, social benefits, and stricter border controls.		
	11-48 (Topic 48)		Chinese Immigration and National Security	63.77	2,773
			Concerns about Chinese nationals crossing the U.S.-Mexico border, with speculation about espionage, security threats, and calls for government action.		
	11-49 (Topic 49)		Migrant Students in US Schools	35.34	1,044
			The topic highlights challenges of migrant students in US schools, focusing on resources, emotional struggles, and local government responses.		
	11-50 (Topic 50)		U.S. Border Wall Debate	54.32	26,717
			The topic discusses the polarized debate over the U.S.-Mexico border wall, focusing on its effectiveness, cost, and political implications.		
	11-51 (Topic 51)		Migrant Child Death Controversy	55.62	1,472
			The death of a 3-year-old migrant during Texas' busing program sparks outrage, criticism of Governor Abbott, and immigration policy debates.		

Table 7: Detailed Topic Discovery Results 5

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-52 (Topic 52)		Free College for Illegal Immigrants	56.25	1,508
			The topic discusses controversy over U.S. states offering free college tuition and benefits to undocumented immigrants, sparking fairness debates.		
	11-53 (Topic 53)		UK Immigration to Rwanda Policy	57.96	1,109
			The UK's plan to send migrants to Rwanda sparks debate on immigration, human rights, and its broader migration policies.		
	11-54 (Topic 54)		Anti-Immigration Sentiment	66.81	94,522
			The topic highlights opposition to illegal immigration in the U.S., featuring negative, xenophobic rhetoric, policy criticism, and misinformation.		
	11-55 (Topic 55)		Refugee vs Illegal Immigrant Debate	53.09	1,923
			The topic highlights the distinction between refugees and illegal immigrants, emphasizing legal, humanitarian differences, and politically charged debates.		
	11-56 (Topic 56)		Undocumented Immigrants and Taxes	51.28	3,505
			The topic discusses undocumented immigrants' significant tax contributions, countering misconceptions and highlighting their economic impact in the U.S.		
	11-57 (Topic 57)		Migrant Child Labor Exploitation	57.58	3,242
			The topic highlights migrant child labor exploitation in hazardous jobs, inadequate protections, and calls for stricter labor laws and standards.		
	11-58 (Topic 58)		U.S. Immigration Court Backlog	38.20	1,562
			The U.S. immigration court system faces significant case backlogs, prompting calls for more judges, funding, and streamlined processing for fairness.		
	11-59 (Topic 59)		Support for Migrant Communities	30.47	1,468
			Tweets express support for migrants, highlighting aid efforts, essential services, and advocacy for their rights and well-being.		
	11-60 (Topic 60)		Green Card Backlog Issues	50.05	2,462
			Legal immigrants, especially Indians, face green card delays due to backlogs and caps, sparking calls for U.S. immigration reforms.		
	11-61 (Topic 61)		Human Smuggling and Border Control	42.10	1,632
			The topic focuses on U.S. immigration enforcement, highlighting human smuggling incidents, law enforcement efforts, and border security challenges.		
	11-62 (Topic 62)		Migrant Deaths and Border Crisis	49.40	2,898
			The topic highlights migrant deaths, inhumane conditions, border crisis, government inaction, and the urgent need for immigration reform.		
	11-63 (Topic 63)		Criticism of US Immigration Policies	67.11	7,850
			Criticism of Texas Governor Greg Abbott's immigration policies, including razor wire use and migrant pushbacks, deemed inhumane and cruel.		
	11-64 (Topic 64)		Expats vs Immigrants	46.36	2,099
			The topic examines the double standard in using ”expats” for Western migrants and ”immigrants” for others, highlighting bias and inconsistency.		

Table 8: Detailed Topic Discovery Results 6

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-65 (Topic 65)		The Great Migration	38.20	1,820
			The Great Migration saw 6 million African Americans move from the South (1915–1970), shaping history, arts, and social justice.		
	11-66 (Topic 66)		Illegal Immigration Debate	62.61	1,096
			The topic discusses polarized views on U.S. illegal immigration, debating distinctions between refugees and illegal immigrants and their impact.		
	11-67 (Topic 67)		Enforcement of Immigration Laws	49.11	6,612
			The topic highlights frustration over perceived lack of immigration law enforcement under Biden, with calls for stricter border control.		
	11-68 (Topic 68)		US-Mexico Border Crisis	39.26	1,252
			The topic highlights the migrant surge at the US-Mexico border, straining resources and prompting concerns over a local crisis.		
	11-69 (Topic 69)		Immigration and Welfare Benefits	56.20	6,321
			The topic discusses U.S. immigration's impact on welfare, with concerns about economic strain and debates on immigrants' contributions.		
	11-70 (Topic 70)		DeSantis' Migrant Relocation Controversy	52.74	18,930
			The topic covers criticism of Florida Governor Ron DeSantis for relocating migrants, facing accusations of exploitation, and calls for accountability.		
	11-71 (Topic 71)		U.S.-Mexico Border Disputes	33.66	1,493
			The topic covers Texas' use of floating barriers in the Rio Grande, sparking legal, humanitarian, and political disputes with the federal government.		
	11-72 (Topic 72)		Immigration Detention Criticism	39.79	1,917
			The topic highlights criticism of the U.S. immigration detention system, focusing on inhumane conditions, private prisons, and calls for reform.		
	11-73 (Topic 73)		Refugee vs Economic Migrant Debate	55.81	2,777
			The topic debates distinguishing refugees from economic migrants, focusing on skepticism of refugee claims and calls for stricter immigration controls.		
	11-74 (Topic 74)		Economic Migrants vs Asylum Seekers	56.06	1,860
			The topic highlights skepticism about asylum claims, emphasizing the need to differentiate genuine asylum seekers from economic migrants.		
	11-75 (Topic 75)		Anti-Immigration Sentiment	61.72	2,193
			The topic highlights anti-immigration sentiment in the U.S., criticizing illegal immigration, asylum seekers, and advocating stricter border control policies.		
	11-76 (Topic 76)		Asylum Seekers and Immigration	47.88	12,369
			The topic highlights asylum seekers' legal rights, distinguishing them from illegal immigrants, and addresses misconceptions through informative debates.		
	11-77 (Topic 77)		Racist Immigration Rhetoric	67.84	1,121
			This topic examines racism in U.S. immigration debates, focusing on assumptions about Latin American immigrants and the impact of racist rhetoric.		

Table 9: Detailed Topic Discovery Results 7

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-78 (Topic 78)		Immigration Support and Resources	20.28	9,172
			The topic highlights support for immigrants through resources, services, and initiatives, promoting empowerment, integration, and a welcoming environment.		
	11-79 (Topic 79)		Immigration and Gun Violence	64.29	16,190
			The topic highlights concerns over crimes by undocumented immigrants involving gun violence, politicizing immigration policies and advocating stricter controls.		
	11-80 (Topic 80)		Illegal Immigration and Crime	67.35	1,213
			The topic highlights illegal immigration in the U.S., focusing on crimes by undocumented immigrants and calls for stricter border control.		
	11-81 (Topic 81)		Anti-Muslim Immigration Sentiment	69.02	25,435
			The topic highlights negative sentiments, conspiracy theories, and discriminatory discourse against Muslim immigration, linking it to terrorism and societal threats.		
	11-82 (Topic 82)		Immigration and Crime Rates	55.95	2,563
			The topic discusses lower crime rates among immigrants, countering stereotypes with studies and emphasizing their positive societal contributions.		
	11-83 (Topic 83)		Immigration and Crime Association	82.99	7,897
			The topic explores the perceived link between immigration and crime, fueling fears, misinformation, and calls for stricter border controls.		
	11-84 (Topic 84)		Merit-Based Immigration Debate	43.18	1,089
			The topic discusses the debate on a merit-based U.S. immigration system, addressing its benefits, criticisms, and related policy complexities.		
	11-85 (Topic 85)		Immigration and Housing Crisis	48.44	26,473
			This topic explores the link between immigration and housing crises, highlighting concerns about demand, affordability, and contributing factors.		
	11-86 (Topic 86)		Blaming Immigrants for Social Issues	55.00	4,362
			This topic addresses the scapegoating of immigrants for social, economic, and political issues, sparking polarized debates on immigration’s impact.		
	11-87 (Topic 87)		Immigrant Work Ethic	49.73	1,586
			The topic highlights immigrants’ strong work ethic, economic contributions, and challenges stereotypes, advocating for recognition of their societal value.		
	11-88 (Topic 88)		Republican Rhetoric on Immigration	52.67	1,179
			The topic discusses claims that Republican ”open borders” rhetoric fuels migrant surges, highlighting political polarization and misinformation on immigration.		
	11-89 (Topic 89)		Replacement Migration Conspiracy	60.56	2,277
			The topic explores ”replacement migration” and its link to the ”Great Replacement” theory, highlighting fears of demographic and cultural shifts.		
	11-90 (Topic 90)		Pro-Immigration Sentiment	41.71	1,064
			This topic highlights supportive tweets advocating for immigrant inclusion, emphasizing their contributions, rights, and a more open immigration policy.		

Table 10: Detailed Topic Discovery Results 8

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-91 (Topic 91)		Migrant Child Trafficking Concerns	68.60	11,884
			The topic highlights concerns over migrant child trafficking, unaccompanied minors, and criticism of the Biden Administration's immigration policies.		
	11-92 (Topic 92)		Migrant Housing Crisis in NYC	45.90	66,591
			The topic highlights debates and tensions over migrant housing in NYC, involving public facilities, safety concerns, and political responses.		
		11-92-1	Housing for Immigrants		6,516
			The topic critiques U.S. immigration policies, suggesting housing migrants in supporters' homes or properties, often using satire or sarcasm.		
		11-92-2	NYC Migrant Crisis Costs		2,487
			The topic highlights New York City's $12 billion migrant crisis cost, sparking concerns over budgets, taxpayer strain, and federal assistance.		
		11-92-3	US Migrant Crisis Response		28,949
			The topic covers the U.S. migrant crisis, focusing on challenges, political responses, and social-economic implications, particularly in New York City.		
	11-93 (Topic 93)		US Immigration and Border Crisis	38.42	8,005
			The topic covers U.S. immigration challenges, migrant surges, root causes, policy impacts, and debates on border control and compassion.		
		11-93-1	U.S.-Mexico Border Crisis		1,864
			This topic focuses on the U.S.-Mexico border crisis, migrant surges, border security, immigration policies, and related political and humanitarian issues.		
		11-93-2	Darien Gap Migration Crisis		1,834
			The topic covers the migrant surge through the dangerous Darien Gap, highlighting hardships, deaths, and debates on immigration policies.		
	11-94 (Topic 94)		Trump's Anti-Immigrant Rhetoric	77.64	1,952
			The topic highlights backlash against Donald Trump's anti-immigrant rhetoric, including his "poisoning the blood" comment likened to Nazi ideology.		
	11-95 (Topic 95)		Immigration and Social Unrest	63.02	4,115
			The topic explores immigration's link to social unrest, focusing on riots, perceptions, policies, media influence, and public opinion.		
	11-96 (Topic 96)		Legal vs Illegal Immigration	47.28	44,491
			The topic highlights debates on legal vs. illegal immigration in the U.S., emphasizing law, politics, and societal impacts.		
		11-96-1	Legal vs Illegal Immigration		1,763
			This topic highlights the contentious debate over distinguishing legal and illegal immigration in U.S. immigration policy discussions.		
		11-96-2	Legal Immigration Debate		13,864
			The topic discusses support for legal immigration, criticism of inefficiencies, and calls for clearer policies and political accountability.		

Table 11: Detailed Topic Discovery Results 9

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-97 (Topic 97)		African Immigration and Colonialism	60.54	1,485
			The topic discusses African migration to the West, linking it to colonialism, Western responsibility, and African governance challenges.		
	11-98 (Topic 98)		Western Intervention and Migration	58.25	9,750
			Western military interventions and regime changes have caused mass migration, refugee crises, and tensions in Europe, sparking widespread debate.		
	11-99 (Topic 99)		Anti-Immigration Conspiracy Theories	64.56	7,207
			The topic involves tweets promoting conspiracy theories linking immigration to societal decline, global elite control, and the erosion of national identity.		
	11-100 (Topic 100)		Controlled Immigration Debate	46.28	1,095
			The topic discusses debates on regulated immigration policies, balancing economic, security, and humanitarian concerns, with varying opinions on stricter controls.		
	11-101 (Topic 101)		Immigration and Social Justice Issues	69.61	11,215
			This topic covers U.S. immigration issues, social justice, LGBTQ+ and racial minority rights, and polarized views on related policies and ideologies.		
	11-102 (Topic 102)		Reducing Immigration Levels	47.73	2,961
			The topic focuses on reducing immigration levels, particularly low-skilled immigrants, with criticism of current policies and government inaction.		
	11-103 (Topic 103)		Immigration Numbers and Policies	42.64	1,233
			This topic discusses ideal immigration levels, policy impacts, and economic, social, and environmental implications of varying immigration rates.		
	11-104 (Topic 104)		Immigration and Cultural Assimilation	54.63	8,461
			The topic discusses immigration and cultural assimilation, focusing on concerns about immigrants adapting to the host country's culture.		
	11-105 (Topic 105)		Immigration and Colonialism Comparison	57.73	1,213
			The topic explores the conflation of immigration and colonialism, critiquing Western nations' immigration policies in light of their colonial histories.		
	11-106 (Topic 106)		Immigration as Invasion	62.92	9,407
			This topic frames U.S. immigration as an "invasion," using inflammatory language to depict migrants as threats to security and identity.		
	11-107 (Topic 107)		Anti-Immigration Sentiment	62.12	1,493
			The topic highlights anti-immigration sentiment, citing cultural erosion, national identity loss, and advocating stricter border controls against mass immigration.		
	11-108 (Topic 108)		Opposition to Mass Immigration	50.42	6,211
			The topic highlights anti-mass immigration sentiment, criticizing governments and advocating stricter border controls due to cultural, economic, and security concerns.		
	11-109 (Topic 109)		Immigration and Veteran Welfare	62.49	20,100
			The topic highlights a debate over perceived prioritization of illegal immigrants' welfare over U.S. veterans' support, sparking political division.		

Table 12: Detailed Topic Discovery Results 10

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-110 (Topic 110)		Anti-Immigration Sentiment	55.67	1,641
			The topic highlights tweets opposing U.S. immigration, emphasizing negative impacts, urgent language, and calls for halting immigration entirely.		
	11-111 (Topic 111)		Uncontrolled Immigration Concerns	51.37	2,791
			The topic highlights concerns about unchecked immigration in the U.S., focusing on social, economic, and cultural impacts with critical sentiments.		
	11-112 (Topic 112)		E-Verify and Immigration Reform	51.75	1,015
			This topic discusses E-Verify's role in immigration reform, debating its effectiveness, implications, and connection to broader immigration issues.		
	11-113 (Topic 113)		Skilled Immigration Debate	42.55	1,155
			The topic discusses debates on skilled immigration in the U.S., highlighting benefits, drawbacks, and calls for a points-based system.		
	11-114 (Topic 114)		Immigration and Demographic Shifts	51.32	5,649
			This topic discusses immigration as a solution to U.S. population decline, addressing low birth rates and economic implications.		
	11-115 (Topic 115)		Labor Shortage and Immigration	44.68	2,255
			The U.S. faces labor shortages, and immigration is proposed as a solution to fill jobs, boost the economy, and alleviate crises.		
	11-116 (Topic 116)		Economic Benefits of Immigration	39.58	1,099
			This topic highlights immigration's positive economic impact, emphasizing labor gap filling, productivity, growth, and countering negative stereotypes.		
	11-117 (Topic 117)		Hiring of Illegal Immigrants	56.68	14,658
			The topic discusses companies hiring illegal immigrants, legal implications, government enforcement, and public debate over U.S. immigration policies.		
	11-118 (Topic 118)		Elon Musk's Immigration Hypocrisy	48.20	1,085
			The topic highlights accusations of hypocrisy against Elon Musk, a South African immigrant, for criticizing U.S. immigration policies despite benefiting from immigration.		
	11-119 (Topic 119)		U.S. Border Immigration Concerns	65.13	16,211
			The topic highlights concerns over U.S. southern border issues, illegal immigration, drug trafficking, and calls for stricter border control.		
	11-120 (Topic 120)		DeSantis' Immigration Policies in Florida	51.96	7,691
			The topic discusses criticism of Florida Governor Ron DeSantis' immigration policies, highlighting their potential harm to the state's economy and industries.		
	11-121 (Topic 121)		Immigration and Farm Labor	54.98	1,683
			The topic examines U.S. immigration's impact on agriculture, highlighting migrant workers' essential role, economic effects, and policy contradictions.		
	11-122 (Topic 122)		Immigration and Labor Economics	55.42	9,611
			The topic examines immigration's impact on wages, job availability, and labor economics, highlighting debates over policy, inequality, and corporate roles.		

Table 13: Detailed Topic Discovery Results 11

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-123 (Topic 123)		Immigration and Job Competition	62.69	17,645
			The topic discusses concerns about illegal immigration's impact on U.S. jobs, wages, labor laws, and calls for stricter regulations.		
	11-124 (Topic 124)		Challenging "Illegal" Immigration Labels	48.08	3,757
			The topic debates the term "illegal" for immigrants, advocating for inclusive language and highlighting migrants' positive societal contributions.		
	11-125 (Topic 125)		Colony Ridge Immigration Controversy	59.88	3,000
			The controversy over Texas's Colony Ridge housing development allegedly housing illegal immigrants sparks concerns about crime, governance, and inaction.		
	11-126 (Topic 126)		U.S. Border Immigration Crisis	57.29	4,971
			The topic highlights debates on U.S.-Mexico border issues, focusing on illegal immigration, border security, and polarized views on policy approaches.		
	11-127 (Topic 127)		Anti-Immigration Sentiment	70.58	9,742
			This topic covers hostile tweets on U.S. illegal immigration, using derogatory language, criticizing policies, and promoting divisive rhetoric.		
	11-128 (Topic 128)		Immigrant Experiences and Stories	26.71	7,196
			The topic highlights immigrant experiences in media, emphasizing authentic representation in films, books, art, and cultural storytelling.		
	11-129 (Topic 129)		Immigrant Cultural Exchange	39.87	2,150
			The topic highlights immigrants enriching local culture through unique culinary contributions, fusion dishes, and immigrant-owned restaurants, fostering diversity.		
	11-130 (Topic 130)		U.S. Border Control Debate	61.95	10,098
			The topic highlights concerns over U.S. immigration, emphasizing stricter southern border control, national security, and the economic impact of illegal crossings.		
	11-131 (Topic 131)		Opposition to Immigrant Housing	61.78	6,356
			The topic highlights frustration over housing resources for illegal immigrants, perceived as prioritized over homeless Americans, sparking criticism of immigration policies.		
	11-132 (Topic 132)		Anti-Immigration White Nationalism	76.81	2,298
			The topic involves tweets with anti-immigration sentiments, white nationalist ideologies, and alarmist claims about immigration eroding white majority populations.		
	11-133 (Topic 133)		California Immigration and Crime Issues	67.00	2,828
			Criticism of California's immigration, crime, and social issues under Governor Newsom, highlighting homelessness, high taxes, and illegal immigration concerns.		
	11-134 (Topic 134)		Racial Bias in Immigration	66.61	4,648
			This topic examines racial bias in U.S. immigration policies, highlighting disparities in treatment and perception based on skin color.		
	11-135 (Topic 135)		Black American vs Immigrant Dynamics	68.28	18,804
			The topic explores tensions between Black Americans and black immigrants, focusing on identity, culture, societal treatment, and shared complexities.		

Table 14: Detailed Topic Discovery Results 12

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-136 (Topic 136)		Busing of Migrants by Texas Governor	47.59	7,050
			The topic covers Texas Governor Greg Abbott's controversial policy of busing migrants to Democratic-leaning areas, sparking political and humanitarian debates.		
	11-137 (Topic 137)		Illegal Immigration as Crime	65.62	2,334
			This topic focuses on framing illegal immigration as criminal, advocating strict enforcement, deportation, and criticizing lenient immigration policies.		
	11-138 (Topic 138)		Immigration and National Security	68.13	13,762
			The topic highlights concerns over U.S. immigration, focusing on national security threats, illegal crossings, and criticism of border policies.		
	11-139 (Topic 139)		Sanctuary Cities and Immigration	59.09	21,501
			The topic discusses the polarized debate on transporting illegal immigrants to sanctuary cities, highlighting political controversy and perceived hypocrisy.		
	11-140 (Topic 140)		Criticism of Immigration Hypocrisy	63.33	6,869
			The topic critiques perceived hypocrisy of immigration supporters, questioning their willingness to personally house illegal immigrants, often sarcastically.		
	11-141 (Topic 141)		Opposition to Immigration Spending	60.57	1,376
			Tweets criticize New York City's spending on illegal immigrants, arguing it neglects legal residents' needs, infrastructure, and essential services.		
	11-142 (Topic 142)		Opposition to US Immigration Policies	62.01	9,516
			The topic highlights opposition to US immigration policies, criticizing government inaction and prioritization of illegal immigrants over citizens.		
	11-143 (Topic 143)		U.S. Immigration Border Crossings	47.83	2,738
			The topic covers social media debates on U.S. illegal border crossings, immigration policy effectiveness, and partisan views on migrant issues.		
	11-144 (Topic 144)		Illegal Immigration Concerns	55.28	1,889
			The topic highlights frustrations over illegal immigration in the U.S., emphasizing law enforcement, national security, and criticism of enabling policies.		
	11-145 (Topic 145)		Immigration and Illegal Aliens	61.30	1,987
			The topic focuses on U.S. immigration issues, particularly illegal immigration, highlighting debates, concerns, and divisive sentiments on related policies.		
	11-146 (Topic 146)		Illegal Immigration Debate	64.50	10,766
			The topic addresses the polarized debate over terminology for undocumented immigrants in the U.S., tied to border security and societal impact.		
	11-147 (Topic 147)		U.S. Immigration Concerns	61.67	43,217
			The topic highlights criticisms of U.S. immigration policies, focusing on illegal immigration, border control, and national security concerns.		

Table 15: Detailed Topic Discovery Results 13

Lv.1 Cluster	Lv.2 Cluster	Lv.3 Cluster	Label	Toxicity	Count
	11-148 (Topic 148)		Anti-Immigration Sentiment and Treason Accusations	71.79	1,668
			The topic highlights strong criticism of U.S. immigration policies, accusing the Biden administration of enabling illegal immigration and undermining sovereignty.		
	11-149 (Topic 149)		Stopping Illegal US Immigration	50.79	2,049
			This topic focuses on reducing illegal immigration to the U.S., emphasizing stricter border controls, national security, and economic concerns.		
	11-150 (Topic 150)		Social Security for Illegal Immigrants	63.71	4,014
			Twitter users express outrage over claims that illegal immigrants receive Social Security funds, surpassing payments to legal citizens and retirees.		
	11-151 (Topic 151)		Healthcare for Illegal Immigrants	62.69	2,492
			The topic discusses social media debates on taxpayer-funded healthcare for illegal immigrants, highlighting concerns about fairness, costs, and policies.		
	11-152 (Topic 152)		Cost of Illegal Immigration	59.85	14,342
			This topic highlights concerns about the economic burden of illegal immigration, taxpayer-funded benefits, and criticism of government immigration policies.		
	11-153 (Topic 153)		U.S. Immigration and Funding Criticism	61.48	14,601
			Criticism of U.S. immigration policies focuses on national security, economic burden, and fund allocation prioritizing immigrants over domestic needs.		
	11-154 (Topic 154)		U.S. Immigration and Economic Concerns	63.35	6,854
			The topic critiques U.S. immigration policies under Biden, linking illegal immigration to economic issues, crime, and social challenges.		
	11-155 (Topic 155)		Illegal Immigration Debate	56.78	1,552
			The topic focuses on negative perceptions of illegal immigration in the U.S., highlighting conservative, anti-immigrant views and related societal concerns.		
	11-156 (Topic 156)		Unconstitutional Immigration Actions	49.43	1,283
			This topic discusses U.S. immigration issues, highlighting perceived unconstitutional actions, policy criticism, and calls for legal adherence and reform.		
	11-157 (Topic 157)		Calls for Immigration Illegality	50.57	1,450
			The topic focuses on tweets advocating stricter U.S. immigration laws, emphasizing illegal actions, stricter enforcement, and frustration with current policies.		

Table 16: Detailed Topic Discovery Results 14

<table>
<thead>
<tr><th>Week</th><th>Increasing Toxicity Group</th><th>Reference Group</th></tr>
</thead>
<tbody>
<tr><td>4/17–4/23</td><td>41. DeSantis Immigration Policy Debate (42.36)</td><td>148. Anti-Immigration Sentiment and Treason Accusations (71.79)</td></tr>
<tr><td>4/24–4/30</td><td>23. U.S. Immigration Politics and Controversy (71.11)</td><td>54. Anti-Immigration Sentiment (66.81)</td></tr>
<tr><td>5/1–5/7</td><td>29. Supreme Court Immigration Rulings (31.79)</td><td>149. Stopping Illegal US Immigration (50.79)</td></tr>
<tr><td>5/8–5/14</td><td>41. DeSantis Immigration Policy Debate (42.36)</td><td>127. Anti-Immigration Sentiment (70.58)</td></tr>
<tr><td>5/15–5/21</td><td rowspan="3">22. Temporary Protected Status for Immigrants (30.18)</td><td>112. E-Verify and Immigration Reform (51.75)</td></tr>
<tr><td>5/22–5/28</td><td>117. Hiring of Illegal Immigrants (56.68)</td></tr>
<tr><td>5/29–6/4</td><td>34. Amnesty for Illegal Immigrants (52.52)</td></tr>
<tr><td>6/5–6/11</td><td rowspan="2">96. Legal vs Illegal Immigration (47.28)</td><td>35. U.S. Immigration Policy Debate (53.13)</td></tr>
<tr><td>6/12–6/18</td><td rowspan="2">96. Legal vs Illegal Immigration (47.28)</td></tr>
<tr><td>6/19–6/25</td><td>6. Civics Test for Voting (50.33)</td></tr>
<tr><td>6/26–7/2</td><td>14. Birthright Citizenship Debate (57.51)</td><td>127. Anti-Immigration Sentiment (70.58)</td></tr>
<tr><td>7/3–7/9</td><td>23. U.S. Immigration Politics and Controversy (71.11)</td><td>101. Immigration and Social Justice Issues (69.61)</td></tr>
<tr><td>7/10–7/16</td><td>41. DeSantis Immigration Policy Debate (42.36)</td><td>77. Racist Immigration Rhetoric (67.84)</td></tr>
<tr><td>7/17–7/23</td><td>24. U.S. Immigration Policy Debate (40.33)</td><td rowspan="2">14. Birthright Citizenship Debate (57.51)</td></tr>
<tr><td>7/24–7/30</td><td>93. US Immigration and Border Crisis (38.42)</td></tr>
<tr><td>7/31–8/6</td><td rowspan="4">39. Secret Flights of Illegal Immigrants (56.67)</td><td>155. Illegal Immigration Debate (56.78)</td></tr>
<tr><td>8/7–8/13</td><td>147. U.S. Immigration Concerns (61.67)</td></tr>
<tr><td>8/14–8/20</td><td rowspan="3">127. Anti-Immigration Sentiment (70.58)</td></tr>
<tr><td>8/21–8/27</td></tr>
<tr><td>8/28–9/3</td><td rowspan="5">147. U.S. Immigration Concerns (61.67)</td></tr>
<tr><td>9/4–9/10</td><td>155. Illegal Immigration Debate (56.78)</td></tr>
<tr><td>9/11–9/17</td><td rowspan="4">149. Stopping Illegal US Immigration (50.79)</td></tr>
<tr><td>9/18–9/24</td></tr>
<tr><td>9/25–10/1</td></tr>
<tr><td>10/2–10/8</td><td>119. U.S. Border Immigration Concerns (65.13)</td></tr>
<tr><td>10/9–10/15</td><td>138. Immigration and National Security (68.13)</td><td rowspan="2">127. Anti-Immigration Sentiment (70.58)</td></tr>
<tr><td>10/16–10/22</td><td>43. Immigration and Military Service (60.51)</td></tr>
</tbody>
</table>

Table 17: Weekly Topic Trajectories – Increasing Toxicity vs. Reference Group (Topic Toxicity Scores)

Week	Decreasing Toxicity Group	Reference Group
4/17–4/23	35. U.S. Immigration Policy Debate (53.13)	50. U.S. Border Wall Debate (54.32)
4/24–4/30	67. Enforcement of Immigration Laws (49.11)	88. Republican Rhetoric on Immigration (52.67)
5/1–5/7	5. Ilhan Omar Immigration Controversy (73.93)	
5/8–5/14	96. Legal vs Illegal Immigration (47.28)	24. U.S. Immigration Policy Debate (40.33)
5/15–5/21	143. U.S. Immigration Border Crossings (47.83)	
5/22–5/28	34. Amnesty for Illegal Immigrants (52.52)	22. Temporary Protected Status for Immigrants (30.18)
5/29–6/4	96. Legal vs Illegal Immigration (47.28)	
6/5–6/11		41. DeSantis Immigration Policy Debate (42.36)
6/12–6/18	14. Birthright Citizenship Debate (57.51)	88. Republican Rhetoric on Immigration (52.67)
6/19–6/25		23. U.S. Immigration Politics and Controversy (71.11)
6/26–7/2		
7/3–7/9	37. Discrimination and Immigration Laws (52.73)	84. Merit-Based Immigration Debate (43.18)
7/10–7/16	15. Ted Cruz Immigration Hypocrisy (62.98)	115. Labor Shortage and Immigration (44.68)
7/17–7/23	14. Birthright Citizenship Debate (57.51)	93. US Immigration and Border Crisis (38.42)
7/24–7/30	67. Enforcement of Immigration Laws (49.11)	
7/31–8/6	96. Legal vs Illegal Immigration (47.28)	24. U.S. Immigration Policy Debate (40.33)
8/7–8/13	76. Asylum Seekers and Immigration (47.88)	70. DeSantis' Migrant Relocation Controversy (52.74)
8/14–8/20	72. Immigration Detention Criticism (39.79)	34. Amnesty for Illegal Immigrants (52.52)
8/21–8/27	23. U.S. Immigration Politics and Controversy (71.11)	33. Reagan's Immigration Amnesty Policy (50.24)
8/28–9/3		41. DeSantis Immigration Policy Debate (42.36)
9/4–9/10	96. Legal vs Illegal Immigration (47.28)	68. US-Mexico Border Crisis (39.26)
9/11–9/17	58. U.S. Immigration Court Backlog (38.2)	143. U.S. Immigration Border Crossings (47.83)
9/18–9/24	29. Supreme Court Immigration Rulings (31.79)	33. Reagan's Immigration Amnesty Policy (50.24)
9/25–10/1	35. U.S. Immigration Policy Debate (53.13)	23. U.S. Immigration Politics and Controversy (71.11)
10/2–10/8	29. Supreme Court Immigration Rulings (31.79)	88. Republican Rhetoric on Immigration (52.67)
10/9–10/15	72. Immigration Detention Criticism (39.79)	98. Western Intervention and Migration (58.25)
10/16–10/22	58. U.S. Immigration Court Backlog (38.2)	101. Immigration and Social Justice Issues (69.61)

Table 18: Weekly Topic Trajectories – Decreasing Toxicity vs. Reference Group (Topic Toxicity Scores)