




ALETHEA

PromptPasta: Stirring the Pot of Online Discourse

How Generative AI Is Being Used and Confused
by Inauthentic Influence Campaigns

By The Alethea Team

08.2025



PromptPasta: Stirring the Pot of Political Discourse

How Generative AI Is Being Used and Confused by Inauthentic Influence Campaigns

EXECUTIVE SUMMARY

Alethea identified a network of **at least 400 X accounts** exhibiting inauthentic behavior and coordinated narrative promotion. The network used large language models (LLMs) to generate subtly varied cospasta-style content—a technique **Alethea terms *PromptPasta***—to amplify messages aligned with key figures and policy positions associated with the second Trump Administration.¹ The outputs of PromptPasta are **nuanced**, with **variations in phrasing** among the posts that correspond to each LLM prompt, making them **harder to detect** than identical posts shared via cospasta tactics.

The network's scale, speed, and narrative discipline illustrate how generative AI is **lowering the barrier** to sophisticated influence operations. While attribution remains

¹ Cospasta is a block of text that is repeatedly copied and pasted across the internet, often to spread memes, jokes, or spam. It typically becomes viral due to its humorous, absurd, or provocative content.

uncertain, behavioral indicators—batch account creation, reused profile images, identical replies, and lack of original posts—strongly suggest **inauthentic coordination**.

During the 2024 election, the network shared content promoting **distrust of Kamala Harris**, advocating for **Trump’s re-election**, and later praising figures in his second administration, including **Press Secretary Karoline Leavitt and HHS Secretary Robert F. Kennedy Jr.** Recently, some accounts shifted to sharing **criticism of the Administration’s handling of the Jeffrey Epstein files** and posted erratic commentary on Trump officials and Rachel Maddow. X suspended at least 70 of the accounts during Alethea’s investigation.

Although not linked to a specific actor or foreign influence, the network’s use of generative AI **suggests that these technologies pose a growing risk** to democratic discourse. Its evolving tactics offer early insights into how AI may be weaponized at scale in future influence campaigns.

NARRATIVES

Sowing Distrust, Growing Doubt

The network posted its earliest content shortly before the 2024 Presidential Election. These posts conveyed a growing sense of distrust and dissatisfaction among voters regarding former Vice President Kamala Harris’s candidacy. Many of the accounts expressed their reluctance to support her, citing concerns about her integrity, including her alleged alignment with “Big Pharma” and perceived corruption within the Democratic Party.



Election day post expressing displeasure with both candidates

Alethea observed content that did not express support for either candidate for president. For example, there were posts that expressed a lack of support for President Trump while also lamenting that Vice President Harris was no longer an acceptable candidate. In some cases, the posts attributed this shift to the perceptions that the Democratic Party held an increasingly pro-war stance about Gaza and Ukraine and used hashtags such as #NoVote, #NotVoting, and #NoMoreWar.

Other accounts shared content emphasizing a perceived threat from the political left, with posts containing claims that the left will “never stop trying to destroy President Trump” and similar messages. The posts also contained calls for unity among Republican voters to ensure Trump’s success in the upcoming election, and suggested a strong mobilization of Trump supporters would ensure election integrity and prevent another “Russian collusion hoax” meant to delegitimize Trump’s victory.



Post mentioning a vote for Trump and warning of another Russian collusion hoax

Many of the accounts appeared to be attempting to develop personas intended to seem like real Trump supporters. Alethea observed posts that began with phrases like “As a conservative,” “As a patriotic American,” or “As an America First Republican.” Others further specified a state of residency, starting a post, for example, with the phrase “I’m a proud Arizona Republican.”




Post expressing a vote for Trump in Arizona

Trump's First 100 Days and the Press Secretary

Since January 2025, this coordinated inauthentic network has generally shared content supporting the second Trump Administration's policy positions and expressing approval of White House officials like Karoline Leavitt, Robert F Kennedy Jr., Pam Bondi, and Kash Patel. Since its inception, the network has most frequently used the hashtag #MAGA, followed by the hashtag #AmericaFirst.

Starting in early February 2025, the accounts consistently highlighted White House Press Secretary Karoline Leavitt's ability to communicate the Trump Administration's actions and policies and praised her for her "honesty," "transparency," and "straightforward approach," using PromptPasta style replies. The network exclusively posted this type of content as replies to larger media and influencer posts, a tactic Alethea has observed in other influence operations that may be intended to increase the visibility of posts.



Post applauding Press Secretary Karoline Leavitt for her "transparency," The network frequently referred to Leavitt as "transparent" and "honest."

At the time of writing, one of the network clusters continued to share content praising the Press Secretary, even as another cluster began to share posts attacking the Administration for a perceived lack of transparency around the Epstein files. The


discord observed within this network may be indicative of the challenges associated with using LLMs at scale for this type of campaign, as opposed to human operators crafting the messaging. For example, instead of consistently publishing content aligned with the Trump Administration’s policies, as might be expected from a human-backed network, the LLMs providing content for this network appear to have jumped on the bandwagon with online users who were critical of the Administration’s recent announcements related to the Epstein files.

RFK Jr.’s Appointment to HHS and the Make America Healthy Again Agenda

A sizable portion of the content posted by this coordinated inauthentic network focused on the promotion of RFJ Jr.’s appointment to HHS Secretary and his “Make America Healthy Again” (MAHA) agenda. Alethea identified posts replying positively to posts that described alleged vaccine injuries, supporting the fight against perceived corruption among health agencies and Big Pharma, and portraying RFK Jr. as a champion of the anti-vaccine movement.




Post applauding RFK, Jr.’s confirmation.



Post applauding RFJ Jr.'s "Make America Healthy Again" initiatives.


Epstein Betrayal: Turning on the Administration

In May and June 2025, a cluster of approximately 140 accounts within the network shifted their focus of their content to the growing controversy surrounding the Administration's handling of the Epstein files. Both within this cluster of accounts and on social media at large, many users questioned when or whether AG Pam Bondi was going to release the Epstein files. Initially, these accounts frequently shared posts calling on the Trump Administration to be more transparent about their plans to release the files, emphasizing the public's desire for accountability. Alethea observed posts shared by many of the 140 accounts praising AG Pam Bondi, Kash Patel, and President Trump for signaling that they would uphold Trump's campaign promise of releasing the Epstein files.




Post supporting the Administration's promises to release the Epstein Files.

After the U.S. Department of Justice announced on July 7, 2025, that the Epstein “client list” would not be released, contradicting earlier comments by AG Pam Bondi, the network cluster split into two camps. After a pause in activity from late June through early July, some accounts began to share content defending the decision, while others shared harsh criticism of AG Bondi and President Trump. This inconsistency highlights the reflective nature of LLM-based automation, which draws on broader trends in news and social media content rather than strictly adhering to a preset script written by an actor.



Post suggesting that Trump is lying about the Epstein files

During the lead-up to the July 7 announcement, accounts in the network praised Attorney General Pam Bondi and FBI Director Kash Patel; however, when the memo that was ultimately released did not live up to expectations, the same account pivoted to calling Epstein a “Mossad agent” and alleging there was a cover-up.



Post praising Pam Bondi for the expected release of the Epstein files.



Post criticizing Pam Bondi on her handling of the Epstein files and calling Epstein a Mossad operative.

Both clusters of accounts exhibited nearly identical structural and behavioral characteristics, such as specific batch creation dates, lack of original media, and occasional reposts of generally pro-Trump influencers, further evidence of this network's inauthenticity.

Local and National Policy Discussion

In the weeks following the 2024 election, the network continued to reinforce Trump Administration-aligned policy positions, particularly on election integrity, immigration, and transgender athletes in women's sports. Alethea observed accounts posting content that amplified division among a select group of Democratic House members who took unpopular votes in support of Trump's SAVE Act or supported Trump's proposed legislation about banning transgender athletes from competing in women's sports. At times, the accounts conflated discussion of transgender rights with commentary about the SAVE Act, suggesting that the LLMs did not correctly interpret them as disparate topics.

Alethea repeatedly observed that the accounts in this network replied positively to accounts that showed support for the Trump Administration's SAVE Act, which mandates documentary proof of U.S. citizenship for federal voting registration and is awaiting action in the Senate. While many of the replies criticized Democrats who opposed the bill, some expressed appreciation for a few Democrats who voted in favor.



Post praising a handful of House Democrats for supporting President Trump's SAVE Act.

The network amplified the belief that the participation of transgender women in women's sports leagues was detrimental to women's competitiveness, privacy, and safety. Alethea observed content that supported "banning of biological men from competing in female-designated sports," and used this narrative to attack vulnerable Democratic House members, including Representatives Susie Lee (D-NV), Marcy Kaptur (D-OH), Eugene Vindman (D-VA), Marie Gluesenkamp Perez (D-WA), Laura Gillen (D-NY), Ron Davis (D-NC), Henry Cuellar (D-TX), Ed Case (D-HI), and Jared Golden (D-ME).



Post criticizing House Democrats who don't support President Trump's transgender ban in women's sports.




Post criticizing Rep. Susie Lee's vote on President Trump's policy initiatives.

Rachel Maddow, a Tale of Two Sentiments

The accounts within the network posted less frequently between December 19 and January 14, 2025, and ramped up activity in late January after President Trump took office. On January 13, 2025, [Rachel Maddow announced](#) that she was returning to the anchor chair at MSNBC five nights a week and would expand her on-air presence for the first 100 days of President Trump's Administration.

For the remainder of January 2025, Rachel Maddow was a top subject of this inauthentic LLM-based PromptPasta network. At this stage, mentions of Maddow accounted for approximately 20% of the network's total post volume and were concurrent with other posts that shared positive mentions of President Trump and a brief spike in activity about FBI Director Kash Patel.


The network continued to share content mentioning Maddow, albeit at a declining rate, through mid-May. The accounts in this network both criticized and praised Rachel Maddow's reporting on President Trump's inauguration and 100-day agenda. At times, some of the accounts within the network expressed diametrically opposing views.



Post reflecting a negative view of Rachel Maddow



Post reflecting a neutral view of Rachel Maddow



Post reflecting a positive view of Rachel Maddow



Post from the same account as above, but this time with a negative view, and expressing the same sentiments about left-wing propaganda as other negative depictions from the network

Other Narratives

This network has responded to, or in much smaller volumes, commented on, other prominent topics of discussion related to domestic and foreign policy. Accounts mentioned various issues, including US policy towards China and tariffs, President Trump's strong relationship with Israel, President Trump's proposals for postwar Gaza, the potential for peace in Ukraine, DEI policies, and immigration. The overall volume of content expressing support for specific policy positions or the Administration's initiatives was relatively minimal as compared to the discussion of specific members of the Administration or other public figures.

EVIDENCE OF LLM USAGE

The network exclusively published content in the form of replies to posts outside of the network, which were often posted between 5 and 15 minutes after the publication of the external posts, though at times as quickly as a minute. Among the posts mentioning a specific person or issue, the language varied subtly while repeatedly using specific adjectives and maintaining the same overall sentiment. For example, in posts describing Karoline Leavitt, the terms "honest" and "transparent" were frequently used together to express praise for her performance as Press Secretary.

The posts shared by this network did not contain any links to external websites or articles and were entirely text-based, with a small minority of the posts using hashtags and emojis. Alethea did not observe the use of images or other multimedia content among this network's posts; notably, we saw no AI-generated images associated with these posts. It is possible that the actor behind this network chose to avoid using AI-generated media to help evade detection and increase believability, as AI image quality can vary greatly from model to model.

Alethea identified at least one instance of what appeared to be an accidental posting of an error message from an LLM by a network asset. After experimenting with ChatGPT, Claude, Gemini, Grok, and Llama, among others, we found that Gemini 2.5 Flash produces the exact same message after receiving a prompt that could violate its policies, though at other times, Gemini responded to the same prompt without returning an error. While not conclusive, this suggests the actor may have used Gemini to create at least some of this content.


When asked about the error message, Gemini returned the following:

This response sounds like it came from a Google model, such as Gemini or a similar large language model developed by Google.

The phrase “Apologies for any confusion, but I’m simply here to assist you with your request and answer your questions to the best of my ability” is a common boilerplate or canned response used by Google’s GenAI to clarify its purpose and limitations, especially if a user’s prompt might imply it has personal feelings, opinions, or a memory beyond the current conversation.



Post containing what appears to be an LLM-generated reply informing the user that the AI was unable to complete their request; Alethea was able to replicate this response word for word with Google’s Gemini AI



Reply using Gemini 2.5 Flash, utilizing the same language as seen in the post from the network asset

NETWORK CHARACTERISTICS

Over the last 90 days, this 400-account network has shared an average of 390 posts per day. The narratives in the network's content promote White House officials like Karoline Leavitt, Robert F Kennedy Jr., and Kash Patel and generally align with the Trump Administration's agendas and policies.

Some elements of the network, particularly the focus on local politics, may indicate a desire to build believable personas that suggest to the network's audience that their views are shared by others. Additionally, many of the posts attempted to specify their location in addition to their political affiliation. The network may have targeted local voting behaviors to build a sense of relatability and thereby have greater influence over the social media users viewing its posts.


Indicators of Inauthenticity

A majority of the accounts in the network were created in batches on the same date. Occasionally, batch creation can happen organically. For instance, in the wake of a major world event, more people will create accounts than expected on an average day to discuss the event. Because the creation dates of the accounts in this network do not correspond to any events of great enough magnitude to organically increase account creation rates, and the accounts exhibit other overlapping behaviors, we assess that their shared creation dates are indicative of inauthentic activity.

2024-10-24	3:23:58	1849290694596	Chass Pauline	chass_paul1861
2024-10-24	3:26:23	1849291301713	Dean Schauf	dean_schau500
2024-10-24	3:38:16	1849294294148	Homme Lucille	HommeLuci778
2024-10-24	3:39:55	1849294703554	Roscoe Snover	SnoverR7687
2024-10-24	3:51:16	1849297562987	Dylla Emory	DyllaEmory1897
2024-10-24	3:53:07	1849298031625	Magda Zaibel	zaibel3961
2024-10-24	4:13:23	1849303132817	Bart Mclucas	BartMcluca7030
2024-10-24	4:19:45	1849304730822	Olive Basua	OBasua50469
2024-10-24	4:35:09	1849308607206	Andria Bowdle	AndriaBowd947
2024-10-24	4:39:00	1849309574832	Halgas Wilbur	HWilbur40845
2024-10-24	5:00:03	1849314874511	Molloy Kiley	MKiley95311
2024-10-24	5:18:43	1849319573897	Madison Akles	MAkles59558
2024-10-24	6:23:53	1849335970778	Arrocho Carol	ArrochoC19594
2024-10-24	6:38:34	1849339669747	Mulkin Keren	MulkinK89764


Example of batch creation behavior exhibited by the network

The network's accounts predominantly used stolen profile pictures from other authentic social media accounts, stock photography sites, and obscure memes found on 4chan mirrored forums. In a few cases, different accounts within the network used the same profile image. The lack of original profile images is often an indicator of inauthenticity.



Examples of batch-created profiles reusing the same profile image

Other accounts seemed to display account name and profile picture mismatches, such as using a traditionally masculine name, but having what appear to be feminine profile pictures. While there are certainly users for whom this perceived mismatch is true to their identity, in light of the other indicators of inauthenticity, this appears most likely to be the result of the automated batch creation of these accounts. Alethea has observed this type of persona mismatch in other investigations of influence operations.




Examples of possible display name mismatches

In addition, a number of content errors were also observed that indicate the use of automation tools to create content without human review. For example, the network shared the hashtags #Trump2021 and #Trump2025, or misspelled the Press Secretary’s name as Caroline or Carolyn instead of Karoline Leavitt, while others had errors like “#MAGApic.twitter.com/d0Ym9msPBI.” Interestingly, this example appears to imply that an image was supposed to be attached, which would have been an outlier among the rest of the posts in the network, which did not contain any media.



Post demonstrating an error in the LLM output



Post using an alternative spelling of Karoline Leavitt and using an incongruent hashtag, #Trump2021


Indicators of Coordination

Alethea's Artemis platform detected instances of coordination among these accounts, primarily based on account attributes and posting behaviors. However, the use of LLM-generated text made detection across the broader network more challenging due to the slight variations in phrasing across posts, which we have dubbed PromptPasta. This differs from the less advanced tactic known as copypasta, which involves the sharing of identical posts by a network of accounts that can be easily matched and flagged as coordination by tools like Artemis. As more adversaries turn to generative AI for information operations, detecting automated campaigns will become more challenging.




Alethea's Artemis platform detects coordinated posting between accounts within the network

The network's posting pattern over time was fairly consistent, with activity occurring between approximately 8 AM and 8 PM Eastern time, corresponding with Western Hemisphere users' expected waking hours. The majority of content was concentrated between 9 AM and 6 PM. The network largely only posted on weekdays; it shared a relatively negligible volume of posts on weekends. This further supports the assessment that the accounts are inauthentic and demonstrates their coordinated nature because the content collected in this investigation was almost always operated during business hours.



Polar plot showing the posting behavior of the network by hour



Posting behavior for weekdays (left) and weekends (right)

ABOUT ALETHEA:

Powered by industry-first proprietary data, Alethea's machine learning platform Artemis detects risky, harmful, and fabricated narratives across the information ecosystem, enabling customers to proactively defend against perception manipulation. With deeper, first-party access to niche pockets of the internet, Artemis prevents narrative manipulation before it impacts brand reputation, key assets, physical safety, and market value. Alethea was named one of the most promising cybersecurity startups of 2023 by both *Fortune* and *Business Insider*, and featured on *Fast Company*'s 2025 Most Innovative Companies list.

Turn Insight into Action with Alethea:

[Request a Demo](#)



Copyright © 2025 Alethea
alethea.com