

ONLINE DISCOURSE INTELLIGENCE

ChatGPT

Initial deep dive into brand and online discourse

Purpose

A corporate-grade starting point for understanding the brand's online conversation, major language communities, sentiment, audience signals, and important voices—designed to reveal promising directions for deeper investigation.

Research window

January 1, 2024–August 18, 2026

Prepared for J345 • United States online discourse

Executive readout

+97.4% VOLUME TREND	2.53B EST. VIEWS	97.9M ENGAGEMENTS	17.63M SOCIAL UNIVERSE
-------------------------------	----------------------------	-----------------------------	----------------------------------

ChatGPT’s U.S.-inferred discourse operates at category-defining scale, but it is not one unified brand conversation. Product use, creator education, entertainment, political and cultural argument, model competition, education, and news coverage sit inside the same corpus. Aggregate health metrics therefore conceal several distinct reputation systems.

The dedicated sentiment view is contested: 32% positive, 38% negative, and –9% net sentiment, with Love vs. Hate at –25%. Hate narrowly leads Trust and Joy among coded emotions. The overview widget simultaneously displays +45%, so the dedicated sentiment page remains the working benchmark until widget scope and calculation are reconciled.

The favorable language is led by guides, reposts, step-by-step education, and portfolios of AI tools. These clusters are valuable signs of discovery and distribution, but they can overstate durable product satisfaction. The highest-volume U.S. cluster—Reddit, using, exact, use—is only 38% positive and deserves record-level diagnosis.

Influence is unusually concentrated in X/Twitter and in large news or technology accounts. Forbes, The Wall Street Journal, TechCrunch, major newspapers, Reuters, and Elon Musk can frame the conversation at scale, while a tutorial account such as @chatgptricks produces far more realized engagement than many larger publishers.

Strategic reading

In the United States, the strategic task is to turn enormous cultural salience into trusted, repeatable utility. Measure owned explanation, media framing, creator distribution, and direct user experience as separate systems.

Research scope and query design

Dimension	Setting	Interpretive consequence
Window	Jan. 1, 2024–Aug. 18, 2026	Captures multiple product and reputation cycles; August 2026 is partial.
Geography	United States (Starscape-inferred)	Directional view of U.S.-inferred public discourse; not a population estimate.
Language	English represented 89.9% of U.S.-inferred matched records.	Multilingual U.S.-inferred discourse remains in scope.
Sources	Indexed public social platforms	Platform availability, deletion, privacy settings, and indexing affect coverage.
Analytical views	Trends, Sentiment, Narratives, Topics, Audience, Influencers	Machine-coded outputs require record-level validation before high-stakes decisions.

Boolean query

Saved query

ChatGPT OR "@ChatGPTapp" OR #ChatGPT
Geography filter: Country contains United States of America

Method discipline

- Records are matched posts or mentions, not unique people, users, customers, or a probability sample. Country is inferred from available public signals and can exclude or misclassify records.
- Narrative, sentiment, emotion, demographic, and entity labels are model-based classifications.
- Co-mention does not establish stance: a post may praise, criticize, compare, quote, or merely reference the brand.
- Volume, views, engagements, and social-universe estimates describe different phenomena and should not be added together.

Platform documentation: [Getting Started](#) | [Boolean queries](#) | [Drilldowns](#)

Conversation momentum and inflection points



Starscape monthly matched-record trend for the U.S.-filtered ChatGPT consumer-brand query. The partial August 2026 endpoint should not be read as a completed-month decline.


Phase marker	Verified product context	Strategic reading
May 2024	GPT-4o brought faster text, vision, and voice capabilities and expanded advanced tools to free users.	A broad-access shift: the product moved beyond a text-only expert tool toward a multimodal everyday assistant.
Feb.–Mar. 2025	Deep research and native 4o image generation expanded both high-effort knowledge work and participatory visual creation.	Two growth engines became visible at once: serious delegation and highly shareable creative behavior.
Aug. 2025–May 2026	GPT-5 became the ChatGPT default; Go expanded globally; ads entered Free/Go testing and then a broader pilot.	Capability, price architecture, and monetization became inseparable from trust, privacy, and answer-independence perceptions.

Attribution caution

The phase markers align official product events with visible conversation shifts. They identify hypotheses for drilldown, not proof that a single announcement caused a given spike.

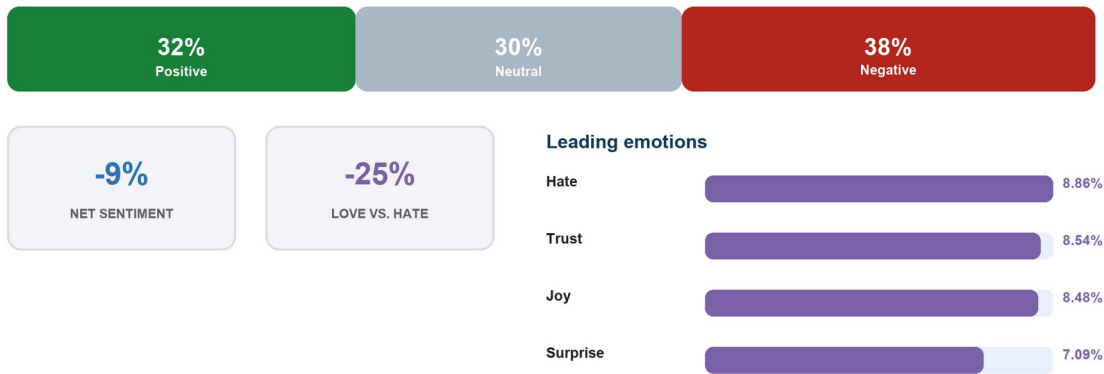
Official context: [GPT-4o and free tools](#) | [Deep research](#) | [4o image generation](#) | [GPT-5](#)

Sentiment and emotional texture



ChatGPT U.S. sentiment profile

January 1, 2024–August 18, 2026



Source: Infegy Starscape dedicated Sentiment and Emotion page; U.S.-inferred records

Analyst-rendered sentiment and emotion view based on Starscape metrics for the full research window. Dedicated sentiment metrics are used here because overview widgets can apply different scope or calculations.

Negative-coded records exceed positive-coded records by six points, yet Trust and Joy remain almost as prominent as Hate. The pattern suggests a highly consequential relationship: people continue to use and discuss ChatGPT while simultaneously contesting its reliability, social role, value, and governance.

Leading emotions

Hate 8.86% · Trust 8.54% · Joy 8.48% · Surprise 7.09%

Narrative structure: what the conversation is about


The leading clusters reveal a portfolio of conversations rather than a single brand narrative. Volume rank signals prevalence; positivity and strategic importance must be interpreted separately.

Narrative cluster	Est. records	Share	Positivity	Interpretation
Reddit / using / exact / use	47,776	2.46%	38.0%	Largest U.S. cluster; direct-use language with substantial friction.
generate / Hey	38,648	1.99%	46.2%	Generation requests or failures; context requires sampling.
guide / Follow / Repost / step-by-step / Gemini	30,750	1.58%	82.7%	Creator-led education and amplification.
Rytr / Jasper / Framer / Bing / Prezi / Durable	24,859	1.28%	90.6%	Highly positive AI-tool portfolio content.
Claude / prompt engineering / analytics / Grok / ethical	24,859	1.28%	58.7%	Comparison, technical education, and ethics.
discovered / Join / artist / AI / sales	21,299	1.10%	65.1%	Creative discovery and commercial invitation.
engine / Codex / search / usage / Enterprise / API	15,925	0.82%	61.0%	Product-system and usage-transition conversation.
GPT / GPT-5 / GPT-4o	15,796	0.81%	69.3%	Model-version shorthand and migration talk.

Reading principle

A smaller, low-positivity cluster can matter more for reputation than a large, celebratory cluster. Prioritize by risk, stakeholder relevance, and persistence—not volume alone.


Semantic network: the conversation's underlying structure



High-resolution Starscape Narratives Force Graph capture for the full ChatGPT query. Colors represent machine-generated narrative communities; position indicates linguistic similarity, not causality or market size.

Network region	What the connections suggest	Strategic use
Direct-use core	The largest connected region centers on Reddit, using, exact, and use language.	Sample for outcome, error, workaround, and repeat-use signals.
Creator education	Guides, reposts, step-by-step instruction, and tool portfolios form highly favorable communities.	Separate instructional reach from independent satisfaction.
Product and API system	Codex, search, usage, Enterprise, users, and API language links product tiers and technical use.	Track limit friction and migration by product surface.
Competitive ecosystem	Gemini, Claude, Grok, AI tools, Google, and OpenAI occupy connected comparison regions.	Code stance and switching intent rather than counting co-mentions alone.

Positive-language network



Analyst synthesis of U.S.-filtered Starscape narrative clusters


Analyst synthesis of the higher-positivity Starscape narrative clusters. Lines indicate thematic adjacency; distances and node sizes are illustrative rather than exported platform metrics.

Positive U.S. language is organized by creator education, AI-tool portfolios, model-version enthusiasm, international-language participation, and creative discovery. These are strong acquisition and social-currency signals, but only post-level sampling can show whether they convert into sustained use or preference.

Validation move

Sample posts from each cluster and code whether positivity expresses product satisfaction, successful outcomes, enthusiasm for AI generally, or creator-distribution language.

Negative and contested-language network



Analyst synthesis of U.S.-filtered Starscape narrative clusters

Analyst synthesis of lower-positivity and contested Starscape narrative clusters. The visual is a structured reading aid, not a native Starscape force-graph export.

Lower-positivity U.S. language includes direct-use friction, generation problems, comparison and ethics, an ambiguous 'Black' cluster, and usage/API transitions. The ambiguous label is especially important methodologically: it should not be interpreted without inspecting the underlying records.

Validation move

Code stance, source, repetition, and issue frame separately. Negative language may criticize the brand, describe a problem the brand is used to solve, or quote a controversy without endorsement.

Topics, entities, and linguistic cues




Starscape semantic topic treemap; the entity list is reported below. Larger words and blocks indicate greater relative prominence within the matched corpus.

Entity	Share	Entity	Share
ChatGPT	17.2%	Donald Trump	0.66%
Gemini	2.68%	Microsoft	0.64%
Elon Musk	0.99%	Google	0.56%
Sam Altman	0.93%	OpenAI	0.53%

Classification quality check


Starscape classified 'Gemini' as a cryptocurrency exchange even though many records in this context refer to Google's AI. The U.S. run also surfaces political figures, showing why entity validation and stance coding are separate from Boolean-query precision.

Audience signals and poster identities



ChatGPT inferred audience signals

Gender classification among records with inferable public signals



Starscape inferred-gender view among records with inferable public signals. Unknown or unclassified records are not represented as a verified demographic population.

Inferred interest	Share
Computing	57.3%
Medical / health conditions	34.4%
Gaming genres	17.2%
Business industries	15.4%
Politics	15.0%
Business	10.5%
Artificial intelligence	3.93%

Poster identities: author-bio language



Author-bio terms associated with the matched conversation. Bio language is a self-presentation signal, not a verified occupation or identity.

The U.S. author-bio view is dominated by AI, Founder, Tech, Creator, Artist, Writer, Engineer, and U.S.-identity language. Inferred signals skew male, but female-coded records show substantially higher average engagement (478 vs. 217). These are public-signal model outputs—not a customer or U.S. population census.


Influence and important voices

1.06M	68.5K	89.2M	1.83
AUTHORS	AVG. FOLLOWERS	TOTAL LIKES	AVG. RECORDS

ChatGPT U.S. influence landscape

Platform concentration and leading accounts in Starscape's default ranking

Platform share



Interpret rankings by role: reach, activity, realized engagement, stance, and independence can point to different accounts.

Leading ranked accounts

1	Forbes 177 records · 12.6K engagements
2	Wall Street Journal 105 records · 7.8K engagements
3	TechCrunch 183 records · 24.0K engagements
4	Elon Musk 5 records · 353.4K engagements
5	New York Times 21 records · 7.0K engagements
6	Washington Post 51 records · 3.9K engagements
7	Reuters 37 records · 1.9K engagements
8	AI / Tech (@chatgptricks) 133 records · 709.2K engagements

Source: Infegy Starscape Influencers page; rankings reflect platform-defined scale and activity signals

Starscape influencer summary and platform activity view. The default ranking combines scale and activity signals; it should not be treated as a pure measure of persuasion or independence.

Voice role	Illustrative accounts	Why the role matters
Business and tech press	Forbes; WSJ; TechCrunch; Bloomberg; CNBC	High-frequency agenda setting around products, markets, work, and risk.
General news	New York Times; Washington Post; Reuters; Fox News	Connects ChatGPT to education, politics, law, culture, and public trust.
Agenda setter	Elon Musk	Very few records can generate exceptional engagement and reframe competition or safety.
Tutorial specialist	AI / Tech (@chatgptricks)	A smaller account with sustained output and exceptionally high realized engagement.
Owned and community voices	OpenAI/ChatGPT; Reddit communities; Instagram creators	Separate owned explanation from lower-visibility direct experience and participatory formats.

Better importance model

Use role-specific lists. Score reach, activity, realized engagement, stance, and independence separately.

Initial strategic implications

These implications are intended as a disciplined starting point. They translate observed discourse patterns into hypotheses that should be tested with narrower queries, post samples, and audience research.

Opportunity	Evidence in the discourse	Initial response direction
Diagnose direct-use friction	The largest narrative is only 38% positive.	Sample by task, failure mode, workaround, and eventual outcome.
Convert education into habit	Guides and tool portfolios are among the most positive clusters.	Link discovery content to repeatable workflows and visible saved value.
Separate media health from user health	Large publishers dominate the default influencer list.	Maintain separate media, creator, owned, and direct-experience dashboards.
Build calibrated trust	Hate, Trust, and Joy are nearly tied while net sentiment is negative.	Demonstrate verification, limitations, and user control within specific jobs.
Track monetization and usage transitions	Usage, Enterprise, API, search, and Codex connect in one cluster.	Monitor tier, limit, ad, and answer-independence language separately.

Decision principle

Use social listening to identify tensions, language, and influential frames. Use additional research to determine prevalence, causality, segment size, and campaign effectiveness.

Directions for deeper exploration

The most valuable next steps are those that challenge an aggregate, isolate a confounder, or create a comparison that could materially change the strategic reading.

1. Compare U.S.-English with U.S.-Spanish discourse and identify whether issue frames, creators, and use cases differ.
2. Separate owned OpenAI/ChatGPT channels, press coverage, creator tutorials, affiliate content, and direct user experience.
3. Sample the Reddit/using/exact/use cluster for successful outcomes, errors, hallucinations, privacy concerns, workarounds, and repeat intent.
4. Create task segments for learning, writing, images, research, health, shopping, coding, and everyday planning.
5. Build role-based influencer lists for agenda setting, technical explanation, direct experience, cultural reach, and conversion.
6. Compare ChatGPT, Claude, and Gemini using identical U.S. dates, geography, language, and jobs-to-be-done.

Questions worth carrying forward

- Which favorable language reflects a completed, repeatable outcome rather than novelty or social participation?
- Which negative frames are persistent, high-reach, and attributable to direct experience?
- Where do owned and partner amplification change the apparent health of the organic conversation?
- Which audience and voice roles can credibly move the brand from category awareness to durable preference?

Appendix: metric notes and sources

Metric notes

- Research window: January 1, 2024 through August 18, 2026. August is a partial month.
- English represented 89.9% of U.S.-inferred matched records.
- Platform-derived views, engagements, universe estimates, demographics, interests, sentiment, emotions, entities, and narratives are directional analytics rather than audited customer measures.
- The semantic-network figure is a high-resolution Starscape force-graph capture. Positive and negative language maps are analyst syntheses using Starscape cluster labels and are explicitly identified as such.
- Official product sources provide contextual phase markers. Their inclusion does not establish causation for conversation changes.

Source links

Infegy: [Getting Started](#) | [Boolean queries](#) | [Drilldowns](#) | [Exporting data](#)

ChatGPT context: [GPT-4o and free tools](#) | [Deep research](#) | [4o image generation](#) | [GPT-5](#)

[ChatGPT Go](#) | [Advertising principles](#)

Planning context

The supplied J345 ChatGPT client brief informed the tensions and exploration priorities. Findings labeled as Starscape evidence come from the saved query and dashboards described in this report.