

Understanding user Beliefs on AI: A Topic Modelling Analysis



ISBN 978-1-943295-24-1

Abhishek Kumar Srivastava
Zillur Rahman
Indian Institute of Technology
(akumarsrivastava@bm.iitr.ac.in)
(zillur.rahman@ms.iitr.ac.in)

This study examines users' perceptions of AI by analyzing 2,589 comments from Anthropic's YouTube channel using topic modelling with LDA. Key themes include users' experiences with AI tools, their desire for human-like interactions, and excitement for new features. Results reveal mixed responses: some are enthusiastic about AI's potential, while others express concerns about usability and transparency. This study extends the literature by capturing consumer beliefs about new GenAI technologies. Practically, findings can guide managers and policymakers in developing response strategies for product development and AI ethics regulation.

Keywords: GenAI, Topic modelling, Anthropic, AI ethics, User Engagement

1. Introduction

Artificial intelligence—the pursuit of replicating human-like intelligence in inanimate machines—has once again captivated global attention, sparking considerable debate across society and academic circles (e.g., Dwivedi, Kshetri, Hughes, et al., 2023). The widespread appeal of this technology stems from a new generation of AI systems that can conduct conversations in natural, human-like language (Salvi, Horta Ribeiro, Gallotti, & West, 2024), assist with a diverse range of tasks (Teubner, Flath, Weinhardt, van der Aalst, & Hinz, 2023), and even generate creative, artistic images from text prompts in unprecedented ways (Heidorn, 2022). AI has been defined by Russell and Norvig (2003) as “Systems that mimic cognitive functions generally associated with human attributes such as learning, speech, and problem solving”. Dwivedi et al. (2019) defined AI as “The increasing capability of machines to perform specific roles and tasks currently performed by humans within the workplace and society in general”. McCarthy (2007) gives a short and crisp definition of AI: “The science and engineering of making intelligent machines.”

Recently, these AI companies have been using social media platforms to promote their GenAI services, and YouTube is one of those platforms that have been used by these companies. Companies such as Anthropic, OpenAI, Cursor, and Bolt have created their YouTube channels and frequently upload videos about how to use their AI services, what's new in their services, and much more. These YouTube videos attract a lot of attention from social media users and AI consumers. They engage with its content by liking, commenting and sharing these videos. Among these engagement behaviors, comments are the most important, as they provide greater knowledge about consumer psychology and their problems, liking and disliking about AI and the content as well. So, understanding these comments can help these AI companies to develop or update their product accordingly, and policymakers and govt can use that to develop regulations around AI use and its advancements. So, based on this discussion, we developed the following research question.

RQ 1: What themes emerge from user discussions about AI models on Anthropic's YouTube channel, particularly in terms of user expectations, interaction experiences, and sentiments toward AI advancements?

To address this research question, we first scrapped online comments from Youtube videos on the Anthropic YouTube channel and then analyzed them after cleaning, preprocessing and using topic modelling techniques. By doing this, we obtained valuable insights about consumers' AI use and consumption of content related to AI. Specifically, we find consumers are discussing about “AI Models, Prompting and Automation”, “Human-AI Interaction and Consciousness” “Open-Source AI and User Expectations”, “User Engagement and Feedback”, “Positive Sentiment towards AI”, and “Anticipation for AI Features”. With this study, we extend the literature on consumer responses to AI services, which is still at a very nascent stage, and our findings will help managers as well as policymakers.

2. Background

Anthropic is a research-driven AI safety organization focused on creating AI systems that are reliable, understandable, and adaptable. The company carefully considers the societal impacts of its work and actively engages with policymakers and civil society organizations in the U.S. and internationally to support the development of safe and dependable AI. By translating research into practical applications, Anthropic provides tools like Claude, which are designed to assist businesses, nonprofits, and community organizations worldwide, contributing positively to both their operations and the broader global community (Anthropic, 2024).

A range of text-mining techniques has been developed to extract themes from large text datasets, with topic modelling emerging as the most widely used approach (Egger and Yu, 2022). This statistical modelling method, applied in machine

learning (ML) and natural language processing, identifies hidden subject themes within a text corpus (Guo, Barnes, & Jia, 2017). Topic modelling algorithms are particularly effective at uncovering latent structures in text data that exhibit well-defined topics, strong topical coherence (i.e., semantic similarity between topics), contextual continuity, and consistent terminology (Kirilenko & Stepchenkova, 2025).

Among the popular ML algorithms for topic modelling are Latent Dirichlet Allocation (LDA), latent semantic analysis, Top2Vec, and BER Topic (Egger and Yu, 2022). Despite the variety of available models, most recent studies employ only a single model, with LDA being the most commonly used (Gallagher, Reing, Kale, & Ver Steeg, 2017). Latent Dirichlet Allocation (LDA) (Blei, Ng, & Jordan, 2003) is a probabilistic model that conceptualizes each document as a combination of topics distributed according to a Dirichlet distribution. Likewise, each topic is viewed as a mixture of words following a Dirichlet frequency distribution. The LDA approach offers an effective method for uncovering these hidden distributions; the distribution of words across topics is then interpreted to reveal the main themes embedded within the document collection.

3. Methodology

In this study, 2589 user comments were collected, cleaned, and analyzed using Python 3.10. Latent Dirichlet Allocation (LDA) topic modelling was applied to extract main topics, associated keywords and main themes from the comments. This approach aimed to capture viewers' perceptions by analyzing data sourced from the YouTube channel of Anthropic.

4. Results

In the first step, we uncovered ten main topics along with their themes and example comments (See Table 1). The authors further analyzed the topics closely and felt certain topics could be merged. Consequently, we found six major themes (See Table 2). This indicates that certain topics have overlapping content, with discussions that are similar in nature.

Table 1 Topics, Keywords and Example Comments

Topic number	Theme	Example comment
Topic 1	AI Models and Prompting	"The prompt doesn't seem to work well with this AI model."
Topic 2	Open-Source AI and User Expectations	"Hope the open-source version has more options for customization."
Topic 3	Human-AI Interaction	"Claude AI is amazing, but I wish it felt more human in interactions."
Topic 4	AI Consciousness and Understanding	"Do you think AI could ever be truly conscious?"
Topic 5	User Engagement and Feedback	"I like using the AI, but sometimes it misses my intent entirely."
Topic 6	Automation and Efficiency	"This tool saves me so much time with repetitive tasks."
Topic 7	AI Use Cases and Applications	"Claude AI is great for quick brainstorming and creative projects."
Topic 8	Positive Sentiment towards AI	"Love how AI makes my work easier! Absolutely amazing!"
Topic 9	Anticipation for AI Features	"Can't wait for the new API to make integration smoother."
Topic 10	Practical Use and Tools	"I use Claude primarily for work-related tasks; it's a solid tool."

Table 2 Topics, Keywords and Example Comments

Merged Theme Number	Theme	Topics Included	Example Comment
1	AI Models, Prompting, and Automation	Topics 1, 6, and 10	"The prompt doesn't seem to work well with this AI model. It's great for automating tasks, though!"
2	Human-AI Interaction and Consciousness	Topics 3 and 4	"Do you think AI could ever be truly conscious? Claude AI feels somewhat human in interactions."
3	Open-Source AI and User Expectations	Topic 2	"Hope the open-source version has more options for customization."
4	User Engagement and Feedback	Topic 5	"I like using the AI, but sometimes it misses my intent entirely."
5	Positive Sentiment towards AI	Topics 7 and 8	"Claude AI is amazing! Love how it makes my work easier!"
6	Anticipation for AI Features	Topic 9	"Can't wait for the new API to make integration smoother."

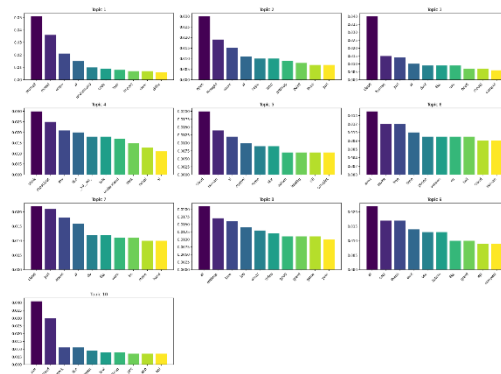


Figure 1 Top Ten Representative Keywords from Each Topic.

5. Discussion

By analyzing comments on videos of Anthropic's YouTube channel, our main objective was to understand common public beliefs surrounding AI. To achieve this, we scrapped comments from Anthropic's YouTube channel and used topic modelling, to understand user responses and identify key themes in prevalent public beliefs. Analysis of each topic, top representative keywords, and example comments revealed some latent user beliefs about Anthropic and AI in general. Specifically, theme 1 contains keywords such as 'prompt', 'model', 'engin', 'ai', and 'understand', and comments on this topic is majorly discussing AI for automating tasks but express frustration with inconsistent responses to prompts. Improving prompt accuracy could enhance task automation, making AI tools more user-friendly and reliable. Theme 2 contains keywords such as 'claud', 'human', 'just', 'ai', 'think', 'conscious' and comments such as "Do you think AI could ever be truly conscious? Claude AI feels somewhat human in interactions.", Here users show interest in human-like qualities in AI, indicating a desire for interactions that feel more natural and empathetic. Fostering such qualities could improve user satisfaction and deepen trust in AI. Theme 3 contains keywords like 'open', 'thought', 'sourc', 'ai', 'hope', 'anthrop' and example comments such as "Hope the open-source version has more options for customization." This suggests users value customization and transparency, favoring open-source AI for its flexibility and control. Theme 4 contains keywords like 'claud', 'human', 'ai', 'control' and comments like "I like using the AI, but sometimes it misses my intent entirely." This means that users here are concerned about intent recognition, and this mismatch can hinder engagement. Enhancing comprehension capabilities can lead to more satisfying, personalized interactions. Theme 5 contains keywords like 'ai', 'anthrop', 'love', 'amaz', 'claud' and comments such as "Claude AI is amazing! Love how it makes my work easier!" here positive sentiment of users show AI's success in simplifying tasks and enhancing productivity, which boosts brand perception and user adoption. Theme 6 contains keywords like 'ai', 'cool', 'thank', 'wait', 'use', 'autom', and comments "Can't wait for the new API to make integration smoother." This discussion suggests users eagerly anticipate new features, indicating strong engagement and willingness to adopt future updates. Delivering regular enhancements can foster loyalty and demonstrate innovation.

Every one of these interpretations provides a roadmap for future advancements in AI functionality and design by shedding light on user requirements, expectations, and experiences. AI companies like Anthropic may be able to improve user satisfaction, build trust, and eventually create a good and fruitful user experience by successfully addressing issues highlighted in the discussion.

6. Theoretical and Practical Implications

By offering a nuanced perspective of user expectations for customization, transparency, and empathy in AI interactions, this study adds to the body of research on user beliefs in AI. It emphasizes how crucial human control is as a cornerstone of AI design. The study shows how topic modelling can reveal latent sentiment patterns and hidden topics in user input. This adds value to computational social science methods by demonstrating the usefulness of LDA in assessing user sentiment for cutting-edge technology like artificial intelligence. By adding customizable features and empathetic interaction models, AI designers can improve user satisfaction. These findings offer practical advice for bringing AI behavior in line with user expectations, which could increase adoption and retention. The expectation for new features suggests that regular, essential upgrades from AI companies like Anthropic could increase customer loyalty. This strategy might increase user engagement, retaining their interest and alignment with developing AI capabilities.

7. Limitations and Future Research Directions

The first limitation of this study is that we have used comments from Anthropic's YouTube channel. It might not accurately reflect how the general population feels about artificial intelligence. Future researchers could include comments from other AI companies' YouTube channels. Second, we have used LDA for topic modelling, and other algorithms may uncover more nuances.

8. References

1. Anthropic. (2024). *Company*. Retrieved from <https://www.anthropic.com/company>
2. Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022. <http://www.jmlr.org/papers/volume3/blei03a/blei03a.pdf>
3. Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., et al. (2019). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, Article 101994.
4. Dwivedi, Y. K., Kshetri, N., Hughes, D. L., et al. (2023). Opinion paper: "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges, and implications of generative conversational AI for research, practice, and policy. *International Journal of Information Management*, 71, Article 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
5. Egger, R., & Yu, J. (2022). A topic modelling comparison between LDA, NMF, top2vec, and bertopic to demystify Twitter posts. *Frontiers in Sociology*, 7, Article 886498. <https://doi.org/10.3389/fsoc.2022.886498>
6. Gallagher, R. J., Reing, K., Kale, D., & Ver Steeg, G. (2017). Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics*, 5, 529–542.

7. Guo, Y., Barnes, S. J., & Jia, Q. (2017). Mining meaning from online ratings and reviews: Tourist satisfaction analysis using latent Dirichlet allocation. *Tourism Management*, 59, 467–479. <https://doi.org/10.1016/j.tourman.2016.08.014>
8. Heidorn, C. (2022). Mind-boggling Midjourney statistics in 2022. *TokenizedHQ*. Retrieved from <https://tokenizedhq.com/midjourney-statistics/>
9. Kirilenko, A. P., & Stepchenkova, S. (2025). Facilitating topic modeling in tourism research: Comprehensive comparison of new AI technologies. *Tourism Management*, 106, Article 105007. <https://doi.org/10.1016/j.tourman.2024.105007>
10. McCarthy, J. (2007). *What is artificial intelligence?* Computer Science Department, Stanford University. Retrieved from <http://www-formal.stanford.edu/jmc/whatisai.pdf>
11. Russell, S., & Norvig, P. (2003). *Artificial Intelligence: A modern approach*. Pearson Education.
12. Salvi, F., Horta Ribeiro, M., Gallotti, R., & West, R. (2024). On the conversational persuasiveness of large language models: A randomized controlled trial. *arXiv*. <https://arxiv.org/abs/2403.14380>
13. Teubner, T., Flath, C. M., Weinhardt, C., van der Aalst, W., & Hinz, O. (2023). Welcome to the era of ChatGPT et al. *Business & Information Systems Engineering*, 65(2), 95–101.