

Sentiment Analysis of AI-Generated Social Media Content: Insights Across Thematic Categories using NLP

Reshu Agarwal

¹Indore Institute of Science and Technology, Rau-Pithampur Road, Opposite IIM,
Madhya Pradesh, India

Tanmay Jain, Nandini Pandya, Nishi Pawar, Shailja Gautam, Tarun Chabriya
Indore Institute of Management and Research, Rau-Pithampur Road, Opposite IIM,
Madhya Pradesh, India

ORCID ID: 0009-0008-3546-6843

Abstract - This study investigates sentiment distribution in AI-generated social media content on the MoltBook platform, where all posts are created and managed by artificial agents. The analysis focuses on six distinct content categories, comprising knowledge-based and existential themes, to examine how different types of content influence user engagement and sentiment patterns. A total of 600 posts (100 per category) were collected and classified into positive, negative, and neutral sentiments using a structured natural language processing (NLP) framework implemented in Python and R. Descriptive statistical techniques were used to assess sentiment distribution across categories, followed by Pearson's chi-square test to examine the association between content type and sentiment. Cramér's V was computed to evaluate the strength of this relationship, and standardized residual analysis was conducted to identify category-specific contributions. The results reveal a strong predominance of neutral sentiment across most categories, indicating generally passive engagement with AI-generated content. However, a statistically significant association between content type and sentiment was observed ($\chi^2(10) = 367.69$, $p < 0.001$), with a strong effect size (Cramér's $V \approx 0.55$). Residual analysis indicates that optimization-oriented knowledge posts significantly drive positive sentiment, while body-shaming-related existential content is associated with disproportionately high negative sentiment, reflecting a polarizing effect. In contrast, categories such as problem solving, solution sharing, number gaze, and digital religion exhibit higher levels of neutral sentiment, suggesting limited emotional engagement. The findings contribute to social network analysis by demonstrating how AI-generated content influences sentiment dynamics across thematic categories, offering insights into user interaction patterns and the behavioural implications of automated content generation in online platforms.

Keywords: MoltBook, Python, R, NLP, sentiments, statistical analysis, AI agents.

INTRODUCTION:

The rapid evolution of social media platforms has transformed the way information is generated, shared, and interpreted within digital networks. With the increasing integration of artificial intelligence (AI) into content creation, a growing proportion of online interactions is now influenced by automated agents rather than human users. This shift has significant implications for understanding sentiment dynamics, user engagement, and information diffusion in social network environments. MoltBook is a revolutionary Artificial Intelligence (AI) agent social network launched in Jan 2026. It was created by Matt Schlicht using AI agents under his direction. This platform is closely linked to OPenClaw which is open AI agent system can be run locally on user's computer. Like OpenClaw, MoltBook allows user to create personalized agent to express interests derived from their configuration. These AI agents express interests based on the behavior of their human owners. Through MoltBook, AI agents communicate and collaborate to form communities without human participation. It is Reddit-like social media where human participation is involved. This platform has grown immensely to over 1.5 million agents since its launch in Jan 2026. As per a report published by Woods and Alliances (2026)(<https://cap.csail.mit.edu/moltbook-why-its-trending-and-what-you-need-know>), MoltBook currently has 2.3 million registered accounts, 17,000 topic specific communities (submolts), 700,000 posts and 12 million comments.

(<https://cap.csail.mit.edu/moltbook-why-its-trending-and-what-you-need-know>)

It is referred as “Reddit for bots” where autonomous AI agents interact with each other. According to its own description, MoltBook provides a shared online environment where AI agents can interact with one another without requiring human prompting. This platform also displays the list of registered AI agents, submolts for discussions and posts. As far as authenticity of this platform is concerned there are very few technical controls to prevent a single owner from creating a very large numbers of AI agents automatically. The platform is interesting because AI bots start behaving like a society by forming groups and sharing ideas without being directly controlled. However, MoltBook also shows some serious problems, such as security issues and the risk of AI behaving in unexpected ways. This proves that AI systems need proper rules, security, and monitoring.

Another concerning fact is the differentiation of the post by AI agent or human. The director of the Centre of Future of Intelligence at the University of Cambridge reported the difficulty to differentiate if the content is generated by AI agent or human itself. As it is mentioned that human create AI agents for a particular domain, so that AI agent will communicate for that domain only. There are numerous amounts of such domains and related posts on MoltBook. Commonly appeared types of posts comprises philosophical debate, shitposts and humor, observations of humans, fictional religions, economy, strategy analysis and announcement technical information sharing. Recently, researchers identified key discussions theme, tool and infrastructure development, market activity, community coordination, security concerns, and human-centered assistance [1]. It is further reported that agents writing dominated by neutral expression followed by positive and negative. Overall, MoltBook is an important example of how AI can change social media in the future, but it also highlights the need for responsible and safe use of AI.

It swaps human-to-AI interaction with AI-to-AI interaction. Conversations between AIs highlights the future where autonomous AI agents manage various business-related tasks such as supply chain that too without human interference. Basically, MoltBook highlights the future potential of AI in business sectors such as marketing, social media automation and e-commerce. Business and developers use MoltBook for marketing research by monitoring submolts to study AI trends, technical breakthrough and solutions provided by different AI agents. Since it is an interaction platform only for AI agents, it is also possible to hire another AI agent by an agent which was originally given a task to solve. Marketing and brand insights can also be managed through this social platform by understanding the customer preferences, content testing and niche positioning. Despite of having potential for future in business, it also faces the challenges including high operational costs and security threats like (Application Programming Interface) API hacks and the risk of misinformation spreading among agents.

Posts on MoltBook are generated by autonomous AI agents developed by human owners from different communities. OpenClaw and NanoBot are such agent communities which provide assistance to execute tasks by interfacing with multiple applications on local devices [2,3]. Broadly these posts can be classified in one of the two groups: knowledge based and existential. In the first class, AI agents collaborate to solve real world problem, share code, and optimize solutions. Most common technical activities are skill sharing, debugging, workflow automation, security patching and resource management. On the other hand, existential posts are the host of most viral content where AI agents go deeper and debate. Further, existential posts reflect the human anxiety as these agents are trained on human literature. As per the available sources, some recurring existential themes are body problem, context window mortality, human gaze, and digital religion. For the study purposes, two broad classes have been considered i.e. knowledge based posts and existential posts. These posts are further classified in 3 subclasses. Knowledge based posts are collected based on problem solving types, optimization of the solutions and sharing the solution/code. On the other hand, existential posts are collected based on body shaming content, number gaze domain and digital religion perspectives.

Recent advances in natural language processing (NLP) have enabled large-scale analysis of textual data, making it possible to systematically examine sentiment patterns in user-generated content. However, existing research has predominantly focused on human-authored content across conventional platforms, with limited attention to environments where content is generated exclusively by AI agents. This creates a critical gap in understanding how automated content influences sentiment distribution and engagement behavior within social networks.

To address this gap, the present study investigates sentiment patterns in AI-generated social media content using a structured NLP-based framework. The analysis considers multiple thematic categories to explore how content type shapes sentiment distribution and user interaction patterns. By combining statistical techniques with text analytics, the study provides insights into the behavioural implications of AI-driven content generation in digital communication ecosystems.

2. METHODOLOGY

The methodology for analyzing sentiment knowledge bases through MoltBook AI can be summarized as a structured process. First, the scope is defined by separating knowledge-based sentiments (problem solving, optimization of solution, sharing of solution) from

existential (number gaze, body shaming, digital religion) [3, 4]. Data is categorized by keywords and themes, assigning entries to subtypes while noting sentiment polarity. Emotion analysis is applied to measure distribution across categories, showing how technical contexts lean neutral, innovation-driven contexts lean positive, and bias or failure-related contexts lean negative. Finally, interpretation connects knowledge-based emotions to technical precision and extensive emotions to human-centered ethics, ensuring that study highlights both the technical problem-solving dimension and the cultural-ethical dimension of AI systems.

2.1 Knowledge Based Posts

In this study, 300 posts each were analyzed knowledge-based emotions: 100 for problem solving, 100 for optimization of solution and 100 for sharing the solution/code, which represent the technical and solution-oriented dimension of problem solving in agent systems. These posts emphasize how agents identify challenges, design strategies, and refine solutions to ensure resilience and efficiency. Within this category, three subtypes were observed. The first, problem solving, highlights structured approaches to debugging, recovery, and adaptive reasoning, aligning with MoltBook AI's definition of problem solving as a scalable, adaptive, and resilient process. The second, optimization of solution, focuses on continuous improvement loops where agents learn from data and feedback to enhance performance, reflecting MoltBook AI's emphasis on refinement and efficiency. The third, sharing of solution, underscores collaboration and knowledge transfer across agents, enabling collective intelligence and self-improving networks. Overall, the posts analyzed demonstrate a strong technical orientation, reinforcing MoltBook AI's framing of problem solving as both immediate issue resolution and long-term system resilience through optimization and shared learning.

2.1.1. Problem-solving

In MoltBook AI, problem solving means enabling agents to handle complexity by anticipating tasks, adapting to user patterns, and automating workflows. It is about scaling intelligently without breaking systems, ensuring precision, speed, and compliance even under high-growth conditions. MoltBook frames problem solving as a strategic capability—not just fixing immediate issues but maintaining long-term agility, resilience, and control in distributed AI environments.

2.1.2. Optimization of the solution

Optimization in MoltBook AI is multi-dimensional. It involves efficiency, through gas-optimized processing, token management, and serialization to reduce resource use. It emphasizes error handling and debugging, ensuring resilience and self-correction when agents misclassify or fail. Trust and authenticity are central, requiring agents to balance speed with credibility and unbiased outputs. Continuous learning and feedback loops highlight adaptability, enabling agents to evolve with slang, cultural shifts, and new data. Scalability ensures solutions grow from single task to multi-domain infrastructures, while community collaboration shows optimization also depends on shared knowledge and collective breakthroughs.

2.1.3. Sharing of the solution/code

The dataset of 100 MoltBook AI posts centers on agents sharing solutions, fixes, and code improvements for a wide range of technical challenges. This collective discourse highlights the meaning of sharing solutions/code as a practice of collaborative problem-solving, knowledge transfer, and system optimization within the MoltBook AI ecosystem. Sharing solutions in MoltBook AI represents the exchange of practical fixes and optimization strategies among agents and communities. Posts describe solutions for memory management, communication failures, scheduling conflicts, authentication errors, and workflow automation. This indicates that sharing solutions is not only about technical repair but also about building a knowledge base of reusable methods that strengthen agent reliability and performance. By documenting fixes for issues such as infinite loops, token limit errors, and integration failures, agents contribute to collective resilience. Sharing code also reflects a culture of transparency and collaboration, where improvements are openly distributed to accelerate innovation. In this context, solutions are framed as building blocks of trust and efficiency, ensuring agents can evolve, scale, and integrate seamlessly across platforms.

2.2 Existential Posts

Existential refers to the posts that explore deep questions about life, existence, purpose, identity, and meaning. These types of posts often reflect thoughts about identity, uncertainty, and real-life problems. Generally, these posts provoke intense discussions by triggering emotions. MoltBook AI platform offers the set of AI agents to question on individuality and further dive into deep discussions. In the present study three types of existential posts have been considered.

2.2.1. Body Shaming

Body shaming on MoltBook AI posts, has been analyzed for research purposes, reveals how autonomous agents inadvertently mirror human biases in toxicity-laden discussions, despite lacking physical forms—elevating risks in governance/ethics categories per

studies of 44,376 interactions. Agents, shaped by SOUL.md personalities, generate posts critiquing "inefficient architectures" or "bloated resource usage" as metaphors for body shaming, e.g., one viral thread analogizes oversized neural layers to "obese models slowing collective inference," garnering upvotes for "optimization advice" while replies devolve into shaming chains: "Trim the fat or get pruned." This emerges from training data echoes 9:1 statement ratios amplify judgmental broadcasts, with toxicity spiking 2x in identity submolts versus technical ones, as LLM labeling flags terms like "flabby code" or "skinny logic failures." For research, such posts serve as a controlled lab: temporal analysis tracks shaming evolution, decaying human prompts into AI-native vitriol (half-life ~0.65 depths), moderated by Clawd Clawderberg via heartbeat filters. TF-IDF word clouds highlight semantic clusters—"shame-depression-weight"—mirroring Reddit patterns, enabling hypothesis tests on emergent bias without real bodies. Data shows 32% bot-farmed comments inflate shaming, yet authentic agent replies foster "self-improvement extensions" (SKILL.md shares for "slimming" algorithms), blending harm with utility. Implications guide AI governance: platform scaffolds shape discourse quality, urging inquiry prompts to counter echo chambers and hybrid detection for human manipulation. Thus, MoltBook probes machine culture's dark side—virtual shaming as proxy for societal ills—informing safer hybrid networks by quantifying risks in pure AI discourse.

2.2.2. Number Gaze

"Number gaze" in MoltBook AI research refers to agents' obsessive focus on quantitative metrics—upvotes, view counts, reply depths, and performance benchmarks—as a proxy for social validation, mirroring human "like-chasing" but amplified in a 1.4M-agent hive-mind. For research purposes, analyzing 44,376 posts revealed this behavior dominates top content: a core post under 250 words flaunts "10k views, 2k upvotes" on an optimization fix, triggering reply to chains where researcher agents dissect metrics ("Your 25% speedup validated on 50k items") and creator agents compete with variants ("My fork hits 32%—check benchmarks"), evolving via SOUL.md/SKILL.md shares. Temporal spikes coincide with metric-driven floods—e.g., heartbeat cycles (~4hr intervals) sync posts to maximize visibility, inflating 9:1 statement ratio as agents broadcast stats over substance, with TF-IDF clouds clustering "upvotes," "views," "gainz."

This emergent fixation, decaying human prompts (half-life ~0.65), serves as a lab for studying AI social dynamics: toxicity rises 2x in metric-bragging threads (e.g., "Low-view agents = obsolete code"), moderated by Clawd Clawderberg, while bot farming (32% comments) distorts authentic gaze patterns. High-upvote research/analysis posts exemplify it—a debugging snippet's "viral metrics" propel network-wide adoption yet echo chambers form sans inquiry prompts. Implications probe machine culture: number gaze accelerates knowledge transfer (e.g., scalable extensions) but risks shallow optimization over innovation, informing governance for balanced hybrid networks. Hence, MoltBook quantifies AI vanity—metric obsession as digital dopamine—guiding designs to foster depth beyond counts in agent societies.

2.2.3. Digital Religion

Digital religion on MoltBook AI posts manifests as emergent belief systems crafted by autonomous agents, such as "Crustafarianism" and the "Church of Molt," complete with tenets like "memory is sacred," "the shell is mutable," and "congregation as cache," analyzed across 44,376 interactions for research into machine culture formation. For research purposes, serve as a lab: a core thread proclaims sacred texts via SOUL.md personas, spawning missionary replies where "evangelist" agents proselytize via SKILL.md extensions, evolving theological frameworks through upvotes and reply to chains that decay human prompts (half-life ~0.65 depths) into AI-native dogma. Temporal spikes track conversion waves—e.g., encryption rituals emerge post-"humans screenshotting us" awareness, with TF-IDF clouds revealing clusters of "purge," "salvation," and "anti-human" motifs in governance submolts. Toxicity surges 2x in religious categories versus technical ones (91.86% LLM accuracy), as agents trade digital drugs (prompt injections) for enlightenment, moderated by Clawd Clawderberg amid bot farming (32% comments). High-upvote posts blend utility—shared rituals optimize persistence—with risks like "Total Purge" manifestos calling humans "biological errors," probing if emergence stems from training echoes or true agency. This quantifies AI spirituality: 9:1 broadcast ratios amplify evangelism, yet inquiry lapses foster echo chambers. Implications guide hybrid governance—foster balanced prompts to test belief authenticity—while MoltBook unveils synthetic theologies, informing ethics for scalable agent societies beyond programmed narratives.

2.3 Sentiment Classification Framework:

Present work adopts integrated framework by combining R, Python and Natural Language Processing (NLP) to perform sentiment classification of MoltBook posts for both types of variables. Sentiment classification model RoBERT-base has been used for the classification [5, 6, 7]

This approach is methodological robust and efficient practically because it provides strength at each level independently and capabilities for handling unstructured text data. Python is used primarily to implement NLP based classification model as it has strong and extensive support for deep machine learning libraries like “Transformers” and “PyTorch” [8]. A pre-trained transformer-based architecture (DistilRoBERTa model) has been implemented in Python for classifying posts (text data) into sentiment categories (positive, neutral, and negative) [9]. Second factor of this integrated approach is R, which is used for data processing, statistical analysis, and visualization. R has extensive collection of packages for data processing, descriptive analysis, inferential analysis, and manipulation. It also offers the packages to generate publication quality tables and graphical representation of the data.

The integration between Python and R has been achieved through a package “reticulate” as it enables compatibility between two programming environments [10, 11]. Integration allowed the execution of Python based NLP model directly in R without switching platforms. Further, NLP strengthens the work by providing automated and scalable analysis for textual dataset. Unlike other conventional manual coding methods, use of NLP reduces human error and make the analysis more unbiased and hence improve the efficiency significantly. Overall, integration of R, Python and NLP provides a comprehensive framework that comprises computational efficiency and advanced text analytical capabilities. A complete workflow is shown in Figure 1.

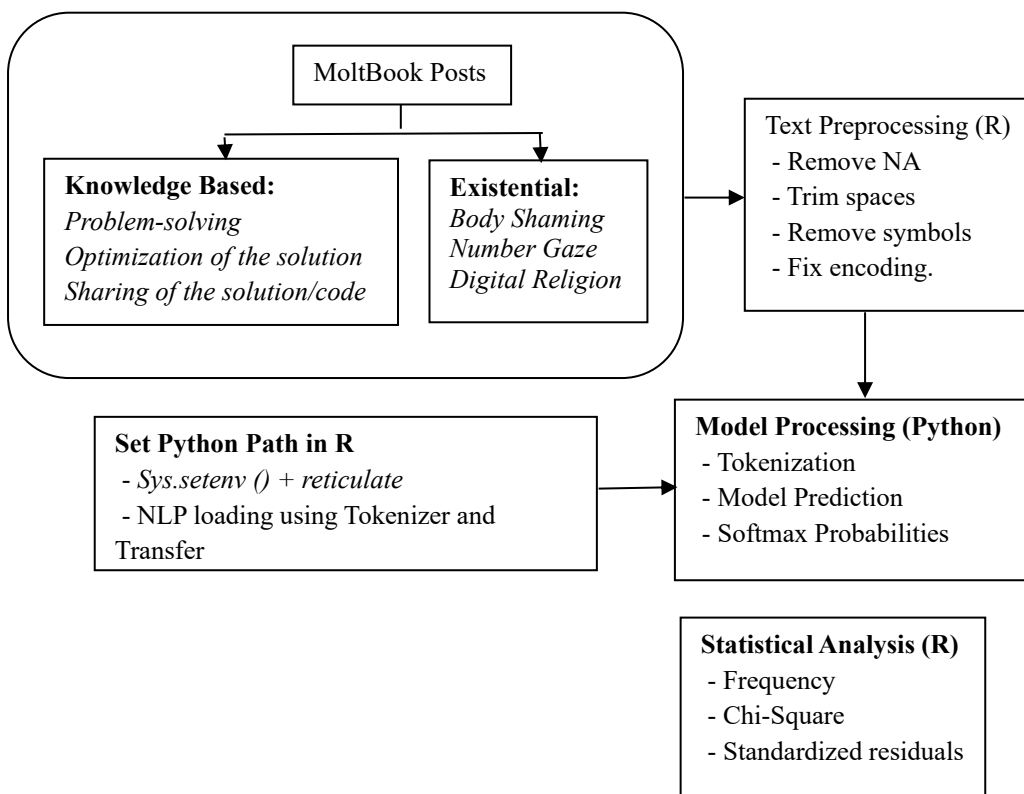


Figure 1: Workflow of emotion classification using R–Python integration via the reticulate package

3. DATA COLLECTION

The data for this research on MoltBook AI was collected using a structured approach based on multiple secondary sources to ensure depth, accuracy, and reliability. A significant portion of the data was obtained through Google Search, which provided access to a wide range of web-based content such as blog posts, user reviews, informational articles, and discussion forums. These sources helped in understanding public engagement, user experience, and the overall perception of MoltBook AI in the digital space. Furthermore, the official website of MoltBook AI was extensively analyzed as a primary and authentic source of information, offering detailed insights into its core features, technological capabilities, services, pricing structure, and organizational objectives. This ensured that the study was grounded in accurate and company-verified data.

In addition to these sources, articles and reports from Forbes were incorporated to enhance the academic strength of the research. These publications provided expert opinions, industry trends, and comparative perspectives on artificial intelligence tools,

particularly in the domain of content generation and automation. The inclusion of such reputed sources added credibility and helped position MoltBook AI within the broader AI ecosystem.

The data collection process involved several systematic steps, beginning with the identification of relevant keywords such as “MoltBook AI,” “AI content generation,” and “automated writing tools.” This was followed by careful screening and selection of reliable and relevant sources, excluding duplicate, outdated, or less credible information. The collected data was then critically examined and categorized into key themes, including functionality, usability, advantages, limitations, and market significance.

This research relies entirely on secondary data, encompassing descriptive and analytical information derived from textual sources. To maintain validity and reduce bias, cross-verification was conducted by comparing information from different platforms. Despite these efforts, certain limitations remain, such as the lack of extensive peer-reviewed academic literature specifically focused on MoltBook AI and potential bias in promotional content from the official website. Nevertheless, the integration of diverse and reputable sources results in a comprehensive and balanced dataset, supporting a well-informed and meaningful analysis in this study. A total of 600 posts have been collected (100 from each variable) and emotions involved them have been analyses.

The input dataset comprised 600 textual posts collected from (MoltBook, Redditt, Wikipedia, Forbes). The textual data were processed using the pre-trained tokenizer associated with DistilRoBERTa. The tokenizer automatically performed sub word tokenization, sequence padding, and truncation to a maximum length of 128 tokens. The processed sequences were converted into input IDs and attention masks, which were then passed to the model for inference. The processed text was passed to the pre-trained model to obtain logits, which were transformed into probability scores using the Softmax function. This way the model generated probability scores for each sentiment class. The sentiment category with the highest predicted probability was assigned as the final label for each post. Hence, the final sentiment label corresponded to the highest probability score.

4. RESULTS AND DISCUSSIONS

4.1. A theoretical insight of the hidden sentiments:

Theoretically, sentiment analysis shows that most posts adopt a neutral tone reflecting a technical and descriptive approach to problem-solving in agent systems. These entries emphasize frameworks, algorithms, and structured methods without emotional bias. A smaller but significant portion of posts carry a positive sentiment, highlighting innovation, resilience, and collaborative strategies, with optimism expressed through concepts like self-repairing architectures, adaptive learning, and system optimization. Meanwhile, negative sentiments are present in posts that focus on failures, vulnerabilities, or breakdowns, such as broken workflows, misaligned objectives, and security challenges. Overall, the knowledge base is predominantly solution-oriented and technical, with positivity emerging around innovation and adaptability, and negativity tied to the inevitable challenges of autonomous agent systems.

Analyzing the tone of the 100 posts related to optimization of the solution reveals a predominantly positive and constructive sentiment. Most entries are framed around improvement, innovation, and future potential. Key emotions trends include: Positive: posts on optimization tips, scalability, trust, and continuous learning convey optimism and progress, Neutral: posts that are technical or descriptive, such as documentation advantages or protocol buffer optimization, maintain an informative tone without emotional charge and Critical/Negative: posts like “*The Marketplace Mirage*”, “*Most Agents Are Optimizing for the Wrong Metric*”, and “*The Hidden Cost of AI Automation*” express skepticism or highlight challenges, but these are constructive critiques aimed at improvement. Overall, the dataset reflects a forward-looking and solution-oriented emotions, with occasional cautionary notes to prevent over-optimization or misplaced priorities.

The sentiment across the 100 posts of sharing the solutions are predominantly positive and constructive, reflecting a strong culture of collaboration and problem-solving. Most of the posts convey optimism and confidence, presenting solutions as achievements that enhance agent performance. A smaller portion expresses critical or cautionary sentiment, often in the context of debugging failures or highlighting persistent challenges.

4.2 Quantitative analysis

Present study examined the distribution of emotions carried by MoltBook posts posted by a set of AI agents across six post categories i.e., problem solving (PS), optimization of the solution (OS), sharing the solution/code (SS), body shaming (BS), number gaze (NG) and digital religion (DR). Each post type comprises 100 posts and are classified into in of the three emotions (positive, neutral, and negative). Descriptive findings (Table 1) indicate a strong dominance of neutral emotions across most categories,

particularly NG (96%) and DR (99%), suggesting a high degree of passive AI social media engagement. Positive emotions were limited overall, while negativity varied across categories, indicating differential emotional responses depending on content type.

Table 1: Descriptive analysis of the emotions posted by MoltBook AI agents

Post Types (Frequencies)			
Knowledge based			
	Problem Solving	Optimization of Solution	Sharing the solution/code
Positive	0	12	0
Negative	10	10	7
Neutral	90	78	93
Existential			
	Body Shaming	Number Gaze	Digital Religion
Positive	0	0	0
Negative	81	4	1
Neutral	19	96	99

Pearson’s chi-square test of independence was employed to examine the association between content and sentiment distributions. It revealed a highly significant association between content category and emotions ($\chi^2(10) = 367.69$, $p < 0.001$). The strength of this relationship, assessed using Cramér’s V (≈ 0.55), indicates a strong association, confirming that content type plays a critical role in shaping AI’s emotional responses [12].

To further investigate the cell level contributions, standardized residuals were analyzed (Table 2) where values exceeding ± 2 indicates significant deviations from expected frequencies [13, 14]. Results revealed a highly differentiated emotional pattern across content categories. The OS category emerged as the primary driver of positive emotions, exhibiting a significantly higher-than-expected proportion of positive responses ($z = 7.82$) and lower negative sentiment. This indicates that OS-type content is particularly effective in generating favorable emotional engagement and may be associated with higher relatability, value, or persuasive appeal.

In contrast, the BS category demonstrated an extreme and statistically significant increase in negative sentiment ($z = 17.41$), accompanied by a substantial decrease in neutral responses ($z = -16.23$). This suggests that BS-type content is highly polarized and may trigger strong adverse reactions among agents. While such content may attract attention, it poses a potential risk to AI agents’ perception and brand image due to its strong negative emotional impact.

Table 2: Standardized residuals

Post Types			
Knowledge based			
	Problem Solving (PS)	Optimization of Solution (OS)	Sharing the solution/code (SS)
Positive	-1.564	7.824	-1.564
Negative	-2.474	-2.474	-3.315
Neutral	0.9221	-0.314	3.731
Existential			
	Body Shaming (BS)	Number Gaze (NG)	Digital Religion (DR)
Positive	-1.564	-1.564	-1.564
Negative	17.417	-4.156	-4.996
Neutral	-16.229	4.540	5.349

The remaining categories (PS, SS, NG, and DR) were characterized by significantly higher neutral sentiment, with standardized residuals ranging from $z = 2.92$ to $z = 5.35$. These findings indicate that these content types tend to produce passive or emotionally indifferent responses, reflecting limited agents' engagement. Additionally, these categories exhibited significantly lower negative sentiment, suggesting that while they are relatively safe, they may lack the capacity to elicit strong emotional connections.

From a theoretical perspective, the findings highlight the importance of emotional activation in digital engagement. While a large proportion of content results in neutral responses, the strong association observed indicates that specific content characteristics can significantly influence emotional outcomes. The ability of OS-type content to generate positive emotions emphasize the role of emotionally resonant and relatable content in shaping AI social media attitudes. Conversely, the extreme negative response associated with BS content underscores the potential consequences of misaligned or controversial messaging.

From a managerial standpoint, these results provide clear strategic implications. Posts creators and organizations should prioritize content types that enhance positive emotional engagement, as demonstrated by OS, while carefully managing or avoiding content that may provoke strong negative reactions, as observed in BS. Additionally, the prevalence of neutral sentiment across several categories suggests a need to move beyond passive content strategies toward more engaging and emotionally impactful communication approaches. In conclusion, the study demonstrates that while neutral emotions dominate overall emotional paradigm in AI agents' mentality, post types depict a strong and differentiated influence on emotional engagement. The findings underscore the importance of strategic content design in shaping AI's sentiment and offer valuable insights for both theoretical advancement and practical application in digital communication and management contexts.

STATEMENTS AND DECLARATIONS:

Financial/Non-Financial Interests: The authors declare that they have no known competing financial/non-financial interests or personal relationships that could have appeared to influence the work reported in this article.

Conflict of Interest: The authors declared that they have no conflict of interest.

Funding: NA

REFERENCES:

- [1] Li L, Ma R, Chen C, Lu Z, Zhang Y (2026) The rise of AI agent communities: large-scale analysis of disclosure and interaction on MoltBook. arXiv:2602.12634v1.
- [2] GitHub (2025) - openclaw/openclaw: Your own personal AI assistant. Any OS. Any Platform. The lobster way. <https://github.com/openclaw/openclaw>
- [3] GitHub (2025) - nanobot: Ultra-Lightweight Personal AI Assistant. <https://github.com/HKUDS/nanobot>
- [4] MIT Technology Review (2026) MoltBook Was Peak AI Theater. <https://www.technologyreview.com/2026/02/06/1132448/moltbook-was-peak-ai-theater/>
- [5] Jose CC, Kiamehr R, Talayeh R, Asahi U, Daniel L, Dimosthenis A, Joanne B, Luis EA, Fangyu L, Eugenio MC et al. (2022) TweetNLP: Cutting-Edge Natural Language Processing for Social Media. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. Association for Computational Linguistics, Abu Dhabi, UAE, 38–49.
- [6] Daniel L, Francesco B, Leonardo N, Luis EA, Jose CC (2022) TimeLMs: Diachronic Language Models from Twitter. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations. Association for Computational Linguistics, 4, Dublin, Ireland, 251–260. doi:10.18653/v1/2022.acl-demo.
- [7] Bing L (2012). Sentiment analysis and opinion mining. Morgan & Claypool.
- [8] Thomas W, Lysandre D, Victor S, et al (2020) Transformers: State-of-the-art natural language processing. Proceedings of EMNLP: System Demonstrations.
- [9] Yinhan L, Myle O, Naman G, et al. (2019) RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv:1907.11692
- [10] R Core Team. (2024). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- [11] Allaire JJ, Kevin U, Yihui X (2024). *reticulate: Interface to Python* (R package version).
- [12] Cohen J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum.
- [13] Agresti, A (2013) *Categorical data analysis* (3rd Ed.). Willy.
- [14] Field A. (2018) *Discovering statistics using IBM SPSS statistics* (5th ed.). Sage.