International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 05 | May - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

Enhancing Text-Only Social Media Interaction with LangChain and Tavily: A Semantic Tagging and AI Reply System

Md Mahtab Ansari[1], Muntashir[2], Md Javed Akhtar[3], Pijush Kanti Ghosh[4], Sumanta Chatterjee[5], Shirshendu Dutta[6]

[1]Student, Dept of Computer Science & Engineering, JIS College of Engineering
[2]Student, Dept of Computer Science & Engineering, JIS College of Engineering
[3]Student, Dept of Computer Science & Engineering, JIS College of Engineering
[4]Assistant Professor, Dept of Computer Science & Engineering, JIS College of Engineering
[5]Assistant Professor, Dept of Computer Science & Engineering, JIS College of Engineering
[6]Assistant Professor, Dept of Computer Science & Engineering, JIS College of Engineering

---

ABSTRACT- This study presents an intelligent system for creating, tagging, and interacting with text-only social media posts using LangChain's large language models (LLMs) and Tavily search integration. The system enables users to compose textual posts that are automatically analyzed to generate relevant context-aware tags through semantic understanding and real-time data. LangChain LLM processes the content to understand user intent and topics, whereas Tavily provides a live external context to refine and recommend accurate hashtags or tags. The system also facilitates comment-and-reply generation using LangChain, allowing automatic initiation or assistance of natural and coherent discussions. Comments and replies maintain contextual continuity, emotional tone, and relevance to the main post. This solution demonstrates how modern language models can enhance user-generated content platforms by supporting richer interactions, discoverability, and content organization, without human moderation. This approach has practical applications in intelligent publishing platforms, content moderation systems, and knowledge sharing communities. The paper is organized into sections presenting the main findings, with references to specific sections indicated by the section numbers. Abbreviations and acronyms were defined at their first occurrence. The online version of the volume will be available on LNCS Online and is accessible to subscribing members.

Keywords: LangChain, Tavily Search, LLM, Post-tagging, AI Commenting, AI Reply suggestions

1.INTRODUCTION

In the current era of digital communication, social media platforms have become essential in sharing information, expressing opinions, and exchanging knowledge. With the exponential increase in user-generated content, particularly in textual form, challenges around effective content organization and engagement have also grown. One of the key problems is the absence of semantic tagging, contextual interaction, and automated conversation generation, which significantly limit content discoverability and user involvement [8]. To address these limitations, this study proposes an intelligent system that integrates LangChain's Large Language Models (LLMs) with the Tavily Search API to enhance user experience on text-based postcreation platforms. LangChain serves as a powerful framework for reasoning, summarization, and natural language generation [2][3][4], whereas Tavily Search API provides access to real-time external web data to ensure the contextual relevance of tags and responses [5]. The proposed system analyzes the content of user posts, generates semantically appropriate tags using both internal content and up-to-date web information, and enables dynamic AI to generate comments and responses. This promotes not only better discoverability but also meaningful interactions among users [6][8]. The automated commenting feature, powered by LangChain's contextual understanding, helps maintain healthy and engaging discussions on social media [1][7]. Furthermore, the fusion of LLMs with real-time data retrieval tools holds significant potential in improving content pipelines for recommendation, moderation, and community engagement [9][10]. LangChain's modular and developer-friendly architecture also makes it suitable for rapid deployment and customization in AI-enhanced applications [3][7].

In summary, this study introduced a novel approach to enriching social media interactions using LangChain and Tavily, aiming for smarter content discovery and improved user participation in online discussions.

2. SYSTEM OVERVIEW

The proposed system enables users to create and interact with text-only posts using intelligent automation powered by LangChain's large language model (LLM) integration and Tavily Search. The system comprises three core components: semantic tag generation, AI-assisted comment creation, and AI-powered response-generation. All components function independently without maintaining conversational memory, making the system scalable and stateless.

2.1 Post Content Analysis and Tag Generation

Upon submission of a text-only post by a user, the content was analyzed using LangChain's LLM to extract core topics and themes. These extracted concepts were forwarded to Tavily Search, which retrieves relevant real-time information from the web. LangChain then synthesizes the post content and Tavily search results to suggest accurate and trending tags. These tags are used to improve the content classification, discovery, and platform engagement.

2.2 Comment Generation Using Post Content

If a post does not have any existing comments, the system automatically generates an initial comment to initiate engagement. LangChain's LLM is prompted by using the post's text alone, without any memory or previous comments. The resulting comments are typically inquisitive, appreciative, or topic-extending and are designed to simulate authentic user interactions and spark further discussion. This interaction is entirely stateless, ensuring low overhead and high performance.

2.3 Reply Generation Based on Individual Comments

Replies were generated solely based on the specific comments to which they responded. When a user or system triggers a reply, LangChain receives only a single comment as context. The LLM then produces a coherent and relevant response without referring to any conversation history or parent thread. This method ensures simplicity and avoids the complexity of maintaining a conversational state while still producing natural context-appropriate replies.

2.4 System Architecture Overview

The system was designed as a modular, scalable, and high-performance pipeline combining a Next.js frontend with an Express.js backend, both of which were developed using TypeScript to ensure strong typing and maintainability.

At the core of the backend logic is the LangChain Framework, which orchestrates natural language processing (NLP) tasks using the llama-3.3-70b-versatile large language model (LLM). The system also integrates Tavily, a live internet search service, to fetch a real-time external context that enriches the understanding and relevance of hashtags, comments, and replies.

The workflow is illustrated in **Fig-1** and proceeds as follows:

1. User-Interaction:
   Users interact with the application via a modern and responsive Next.js interface, where they can create posts, view tags, and see AI-generated comments or replies.
2. API-Communication:
   The front-end sends user-generated content to the Express.js server via RESTful API requests. The server handles request routing and orchestrates communication between internal services.
3. Backend-Processing:
   The Express server passes the post content to the LangChain pipeline, which performs semantic analysis, intent recognition, and task delegation. For context enrichment, LangChain makes parallel calls to llama-3.3-70b-versatile LLM for language understanding and generation. Tavily Search API for retrieving up-to-date contextual information from the internet.
4. Response-Handling:
   Based on this analysis, LangChain returns suggested hashtags, comments, or replies routed back to the front end of the display.
5. Data-Persistence:
   All user data, including posts, tags, comments, and AI-generated replies, were stored in the PostgreSQL database. Prisma ORM is used to manage type-safe interactions with the database and streamline query operations.
6. Stateless-Generation:
   The system avoids in-memory sessions and state management during generation. Each generation task is stateless, ensuring scalability and fast concurrent processing without shared memory bottlenecks.

Fig-1 System Architecture

3.RESULT AND DISCUSSION

3.1 System Demonstration and Interface Overview

The system interface, developed using **Next.js** for the front-end, offers a streamlined and intuitive user experience. As shown in **Fig-2**, users can easily navigate through the dashboard, search, profile, post creation, and viral content sections by using a collapsible sidebar menu. The "Create Post" section contains fields for entering content, uploading media, and triggering AI-assisted features.

Fig-2 Post creation Interface

3.2 AI-Generated Post Tags

When users enter content related to a specific event or topic, the system processes the input using **LangChain's LLM (Llama-3.3-70B Versatile)** and retrieves trending semantically relevant tags using **Tavily Search**. As illustrated in **Fig-3**, the post about "maJISstic 2k25" includes a variety of AI-generated tags that align with the event's theme and institutional context.

Fig-3 AI based content generation

3.3 Comment Generation on User Posts

The system leverages **LangChain's LLM (Llama-3.3-70B Versatile)** to generate contextually relevant and emotionally resonant comments in response to user posts. When a user creates a text-only post, the system semantically analyzes the content to understand the theme, intent, and emotional tone of the post. Based on this understanding, an appropriate AI-generated comment is crafted, as shown in **Fig-4.**

Fig-4 AI generated comment on the post

3.4 Reply Generation on User Comments

Taking conversational flow a step further, the system supports automated replies to user comments by building multi-turn dialogue threads. After generating or receiving a comment, the AI uses LangChain's chaining and memory features to produce a reply that acknowledges the prior comment while pushing the conversation forward.

In **Fig-5**, the AI responds to the user's comment (which itself echoes the original post), injecting excitement and intent to participate in the discussion.

Fig-5 Reply generation on the user comment

4. CONCLUSION AND FUTURE SCOPE

This study introduced an intelligent AI-driven system designed to enhance user-generated social media interactions through automated post-tagging, comment generation, and threaded replies. By leveraging LangChain's large language model pipeline (llama-3.3-70b-versatile) in combination with Tavily's real-time search API, the system demonstrated effective semantic understanding, context-aware tag generation, and emotionally aligned conversational responses, without the need for human moderation.

The system architecture, built using a Next.js frontend, Express.js backend, and PostgreSQL database managed by Prisma ORM, emphasizes modularity, performance, and scalability. Through seamless integration, the system handles the entire lifecycle of a post, from content creation to dynamic AI-assisted interaction, maintaining contextual flow, and enriching platform engagement. The practical outcomes include accurate hashtag generation and naturally flowing AI conversations, proving the potential of the system for real-world content platforms.

This system holds significant promise for future development. Future enhancements may include sentiment-aware response tuning, multilingual support, and image and video content tagging. Additional improvements such as personalized user interaction, AI-based moderation filters, and integration with voice/chatbot systems can extend their applicability to areas such as customer support, education, digital publishing, and community forums.

This study demonstrates how modern LLMs combined with real-time web context and robust system architecture can transform content interaction on digital

platforms, improve discoverability, encourage deeper engagement, and reduce the need for manual oversight.

REFERENCES

1. Micheal, A. A., Bl, K., Ts, A., & Prasanth, A. (2024). *Advancing Educational Accessibility: The LangChain LLM Chatbot's Impact on Multimedia Syllabus-Based Learning*. research square platform llc. https://doi.org/10.21203/rs.3.rs-4399670/v1
2. Mavroudis, V. (2024). *LangChain*. mdpi ag. https://doi.org/10.20944/preprints202411.0566.v1
3. Topsakal, O. and Akinci, C. C. (2023). Creating large language model applications utilizing LangChain: a primer for developing LLM apps quickly. *International Conference on Applied Engineering and Natural Sciences*, *1*(1), 1050–1056. https://doi.org/10.59287/icaens.1127
4. Mavroudis, V. (2025). *LangChain v0.3*. center for open science. https://doi.org/10.31219/osf.io/4gprt_v1
5. Automated Data Analysis with Langchain. (2024). *International Research Journal of Modernization in Engineering Technology and Science*. https://doi.org/10.56726/irjmets58086
6. Wang, Z., Liu, J., Zhang, S., & Yang, Y. (2024). *Poisoned LangChain: Jailbreak LLMs by LangChain*. https://doi.org/10.48550/arxiv.2406.18122
7. P, R. (2024). Leveraging LLaMA3 and LangChain for Rapid AI Application Development. *Journal of Electrical Systems*, *20*(10s), 2146–2153. https://doi.org/10.52783/jes.5539
8. Ramamonjisoa, D., & Suzuki, S. (2024). *Comments Analysis in Social Media based on LLM Agents*. 25–30. https://doi.org/10.5121/csit.2024.141003
9. Jeong, J., Jeong, J., Gil, D., & Kim, D. (2024). Current research and future directions for off-site construction through LangChain using a large Language Model. *Buildings*, *14*(8), 2374. https://doi.org/10.3390/buildings14082374
10. Korat, A. S. (2024). AI-Augmented LangChain: Facilitating Natural Language SQL Queries for Non-Technical Users. *Journal of Artificial Intelligence & Cloud Computing*, *3*(3), 1–5. https://doi.org/10.47363/jaicc/2024(3)335