

Intelligent Mail Management System Using Large Language Models and Gmail API

Prof. Swapna Tikore¹, Shivam Navathar², Rudraksh Zamakade³, Omkar Misal⁴, Shrihari Lokhande⁵

^{1,2,3,4,5}SMT Kashibai Navle College of Engineering Pune

Abstract—The exponential growth in digital communication has led to overwhelming email volumes, making efficient email management a critical need for users. This project presents the design and implementation of an Intelligent Mail Management System that leverages Large Language Models (LLMs) and the Gmail API to automate and enhance email organization, prioritization, and response generation. At its core, the system uses Google OAuth 2.0 for secure authentication and authorized access to users' Gmail accounts without compromising sensitive credentials. Once access is granted, the Gmail API facilitates email retrieval and management operations such as labeling, filtering, and metadata extraction. The retrieved email content is then processed by an advanced LLM, which uses natural language understanding (NLU) techniques to analyze and classify emails into predefined categories such as Work, Personal, Social, and Advertisement. It further identifies key information such as deadlines, tasks, and meeting schedules. The system architecture follows a modular, multi-tier approach, comprising a frontend for user interaction, a backend for business logic and LLM integration, and a database for storing email metadata. Users can interact with the system via queries, such as asking about project deadlines or summarizing recent communications. The LLM interprets these natural language queries and retrieves relevant responses based on the parsed and stored email content. Security and privacy are paramount; the system ensures encrypted communication, secure token storage, and limited data retention policies. This solution not only automates routine email management tasks but also empowers users to interact with their inbox more intelligently, reducing cognitive overload and enhancing productivity.

Index Terms—Intelligent Mail Management, Large Language Models, Gmail API, Email Categorization, Natural Language Processing.

I. INTRODUCTION

In today's digital era, email remains one of the most essential and widely used modes of communication,

both personally and professionally. However, as the volume of email correspondence continues to grow exponentially, users often find themselves overwhelmed by cluttered inboxes, spam messages, and missed important emails. Traditional email clients offer limited tools for sorting and managing large volumes of emails efficiently. While filters and rules help to a certain extent, they often lack the intelligence to understand context or user-specific priorities. This growing problem calls for an intelligent solution that can automate and enhance email organization using advanced technologies. Recent advancements in Artificial Intelligence (AI), particularly in the area of Natural Language Processing (NLP), have paved the way for more intelligent systems capable of understanding and interpreting human language. Large Language Models (LLMs), such as OpenAI's GPT or Google's PaLM, have demonstrated remarkable proficiency in tasks like summarization, question answering, categorization, and generating human-like responses. These models are trained on massive corpora of text and can perform complex reasoning over natural language input. Leveraging LLMs for email management opens up new possibilities for automating tedious tasks like categorizing emails, detecting important information, and even generating suggested responses. The Intelligent Mail Management System proposed in this project utilizes the power of LLMs in conjunction with the Gmail API to deliver a smart and secure solution for managing emails. The system is designed to connect to a user's Gmail account securely through Google's OAuth 2.0 protocol, ensuring that user authentication and data access are handled safely and without compromising credentials. Once authenticated, the system uses the Gmail API to fetch emails, apply labels, and retrieve metadata necessary for processing. Once emails are retrieved, the LLM takes center stage. It analyzes the

content of each email using NLP techniques to determine the context, tone, and purpose of the communication. Based on this analysis, the LLM categorizes emails into predefined labels such as Work, Personal, Social, or Advertisement. It can also detect urgency and highlight emails that require immediate attention. Furthermore, the system supports prompt-based interactions, where users can ask natural language queries like “What are my upcoming deadlines?” or “Summarize today’s emails,” and receive intelligent responses generated by the LLM. The system architecture is modular and follows a multi-tier design to ensure scalability and maintainability. The frontend layer offers an interface for users to interact with their emails and submit queries. The backend layer handles all business logic, API integration, and communication with the LLM. It processes email data, stores relevant metadata, and responds to user commands. The database layer is responsible for storing extracted data such as email subjects, timestamps, sender details, and categorized labels, enabling faster and more efficient querying and filtering. Security and privacy are integral components of the system. OAuth 2.0 ensures that sensitive user credentials are never exposed, and access tokens are used to manage permissions securely. All communication between system components and external APIs is encrypted using HTTPS protocols. Additionally, email content is processed and stored only to the extent necessary for system functionality, ensuring minimal data retention and user privacy. This system addresses a variety of real-world challenges. For individuals, it reduces the time spent organizing and reading through non-essential emails. For professionals, it ensures that important messages are prioritized and critical deadlines or tasks are highlighted. Organizations can also benefit from such a system by integrating it into internal tools to improve productivity and information flow. The adaptability of LLMs means the system can continuously improve over time by learning from user interactions and refining its categorizations and recommendations. The exponential growth in digital communication has transformed email into a critical tool for personal and professional correspondence. Despite its ubiquitous presence, managing an inbox efficiently remains a significant challenge for many users. Conventional email clients provide basic sorting and filtering features, yet they often fall short when it comes to

understanding the nuanced content and priorities within emails. This results in wasted time, missed opportunities, and important messages buried under heaps of less relevant communication. To tackle these challenges, there is a need for a more intelligent and adaptive system that can learn from user behavior and provide context-aware email management. Artificial Intelligence, and more specifically Large Language Models (LLMs), have emerged as powerful tools for interpreting and generating human language. These models, trained on vast datasets, excel at understanding the semantics and intent behind textual data. Integrating LLMs into email management offers the promise of automating complex tasks such as context-based email categorization, prioritization of urgent messages, and drafting intelligent responses. This integration is not only about automation but also about making email management intuitive and personalized, adapting to the unique communication patterns of each user. The use of Google’s Gmail API provides a robust and secure interface to access email data, enabling programmatic retrieval and management of emails without compromising user credentials. When combined with OAuth 2.0 authentication, the system ensures that user consent and data security are maintained throughout the process. This seamless interaction between the Intelligent Mail Management System and Gmail ensures that the system remains trustworthy and compliant with privacy standards. In conclusion, the Intelligent Mail Management System demonstrates the potential of combining LLMs with existing APIs to solve practical, everyday problems in digital communication. As email traffic continues to rise, intelligent systems like this one will become increasingly indispensable. By enhancing the way we interact with our inboxes, this project offers a glimpse into the future of AI-assisted communication management—one where inbox overload is a thing of the past, and users can focus on what truly matters.

1.1. PROBLEM STATEMENT

In the age of information overload, individuals and organizations face significant challenges in efficiently managing and prioritizing the growing volume of emails received daily. Traditional email clients provide basic filtering and labeling capabilities, which often fall short in understanding the context, intent, and urgency of messages. As a result, important emails

may be overlooked, while time is wasted sifting through irrelevant or spam content. This project aims to address these issues by developing an Intelligent Mail Management System that leverages Large Language Models (LLMs) and the Gmail API to automatically categorize, summarize, and highlight critical emails, thereby enhancing user productivity, reducing cognitive load, and ensuring that vital communications are not missed.

1.2. OBJECTIVE

- To study the integration of Large Language Models (LLMs) with Gmail API for intelligent email management.
- To study how prompt-based analysis using LLMs can improve the accuracy of email categorization and prioritization.
- To study the effectiveness of automated labeling and response generation in reducing manual effort in email handling.
- To study the security and privacy challenges in accessing and processing user emails through OAuth 2.0 and propose mitigation strategies.
- To study user interaction patterns and evaluate the system's performance in delivering relevant and context aware email insights.

II. LITERATURE SURVEY

2.1. Survey of System

2.1.1. Ashish Vaswani et al. Paper: *Attention Is All You Need*

Authors: Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, Illia Polosukhin

Objectives:

- 2.1.1.1. To introduce the Transformer architecture and self-attention mechanism.
- 2.1.1.2. To explore its potential in various NLP tasks, including email processing.

Outcomes:

Transformers offered superior scalability and the ability to model long-range dependencies. Achieved state-of-the-art performance across many NLP benchmarks. Applicable to email processing tasks requiring

context-aware understanding.

Research Gaps:

Training Cost: Training large Transformer models is expensive and time-consuming.

Data Requirements: Models demand vast annotated corpora, making training difficult in niche or new domains.

2.1.2. G. Mujtaba, L. Shuib, R. G. Raj, N. Majeed, and M. A. Al-Garadi (2017)

Paper: Email Classification Research Trends: Review and Open Issues, IEEE Access, vol. 5, pp. 9044–9064
DOI: 10.1109/ACCESS.2017.2702187

Objectives:

- 2.1.2.1. To review trends in email classification research across spam detection, phishing filtering, and folder categorization.
- 2.1.2.2. To analyze limitations of traditional ML techniques in multi-class email classification.
- 2.1.2.3. To understand challenges in achieving scalable, accurate with evolving spam tactics.

Outcomes:

Identified the evolution from basic rule-based systems to machine learning. Recognized the integration of productivity-enhancing features like task management and draft suggestions in email systems. Highlighted the growing relevance of personalized models adapting to user behavior.

Research Gaps:

Complex Case Handling: Traditional models underperform with ambiguous or advanced phishing emails.

False Positives/Negatives: High error rates in classifying borderline or disguised spam emails.

Scalability: Issues in applying traditional ML models to high email volumes in enterprise environments.

2.1.3. Soumyadarshani Dash (Analytics Vidhya Blogs)

Paper: The Next Frontier of Email Efficiency with LLM Objectives:

- 2.1.3.1. To investigate how LLMs (e.g., GPT, BERT) can optimize email categorization and personalization.
- 2.1.3.2. To examine context-aware email processing beyond predefined rules.

2.1.3.3. To explore user-centric automation in email prioritization.

Outcomes:

Demonstrated LLMs' ability to dynamically adapt to user needs and context.

Highlighted personalization through continuous learning and interaction.

LLMs outperform rule-based systems in understanding email semantics and importance.

Research Gaps:

High Computational Demands: Significant resources are needed for training and inference.

Real-time Processing Issues: Latency and cost are bottlenecks for deploying LLMs in real-time applications.

Privacy Concerns: Sensitive data handling raises ethical and technical challenges.

2.1.4. Konstantinos I. Roumeliotis, Nikolaos D. Tselikas, and Dimitrios K. Nasiopoulos

Paper: Next-Generation Spam Filtering: Comparative Fine-Tuning of LLMs, NLPs, and CNN Models for Email Spam Classification

Objectives:

2.1.4.1. To compare LLMs, NLP models, and CNNs for spam email classification.

2.1.4.2. To study how fine-tuning improves classification accuracy for evolving email content.

Outcomes: Fine-tuned deep learning models achieved superior accuracy in identifying complex spam (e.g., phishing).

CNNs and LLMs performed well in understanding contextual spam triggers.

Showed superiority over Naive Bayes and SVMs in handling social engineering patterns.

Research Gaps:

Spam Complexity:

Need for even more nuanced models that detect advanced deception techniques.

Computational Expense: LLMs and CNNs are resource-intensive; optimization for real-world use is required.

2.1.5. K. Debnath and N. Kar (2022)

Paper: Email Spam Detection using Deep Learning Approach Conference: 2022 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COM ITCON)

DOI: 10.1109/COM-IT-CON54601.2022.9850588

Objectives:

To utilize advanced deep learning techniques like LSTM and BERT for spam detection.

To showcase Transformer-based models' capabilities in email classification.

LSTM models effectively captured sequential patterns in email bodies.

BERT demonstrated strong context understanding for detecting spam or phishing content.

Deep learning outperformed traditional methods, especially in real-world datasets.

Research Gaps:

Data Dependence: Models like BERT need large labeled datasets for training.

Compute Intensity: Deployment is limited by high computational requirements, especially in constrained environments.

III. PROPOSED SYSTEM




Fig.1 System Architecture

The Intelligent Mail Management System integrates advanced AI capabilities with Gmail's API to streamline email management through automatic categorization, prioritization, and response suggestions. The system leverages Large Language Models (LLMs) like GPT-based architectures, combined with OAuth 2.0 for secure access, to

enhance the user's email experience.

1. User Authentication and Authorization

When a user initiates interaction with the system, the first step involves authenticating the user securely through Google's OAuth 2.0 protocol. The system redirects the user to Google's OAuth consent screen, where the user grants permissions for accessing their Gmail account data. Once authorized, the system receives an OAuth token, which is used to securely access Gmail APIs without exposing the user's credentials.

2. Email Retrieval via Gmail API

Using the OAuth token, the system communicates with the Gmail API to fetch emails from the user's inbox. The retrieval process can be customized based on filters such as unread emails, specific labels, or date ranges. The API returns a list of emails, each containing metadata such as sender information, subject, timestamp, and email content (body). This allows the system to obtain the raw data needed for analysis without manual intervention.

3. Email Preprocessing

Before sending emails to the LLM, the system preprocesses the raw email content. This involves cleaning the text (removing signatures, disclaimers, and irrelevant metadata), tokenization (breaking text into words or phrases), and extracting important features like sender, subject, and timestamps. This step ensures that the data fed into the LLM is clean and contextually relevant, improving the accuracy of analysis.

4. Email Analysis and Categorization Using LLM

The cleaned email content is passed to the Large Language Model for deep analysis. The LLM utilizes advanced natural language processing techniques to understand the semantics, context, and intent behind each email. It categorizes emails into predefined groups such as Work, College, Advertisement, Social, or Urgent. The categorization is based on identifying keywords, phrases, and the overall context, allowing the system to assign multiple labels if needed.

5. Priority Assignment

In addition to categorization, the LLM evaluates the urgency and importance of each email. For example,

emails from known contacts or containing deadlines are flagged as high priority. This prioritization helps users focus on the most critical communications first, reducing time spent sorting through less relevant emails.

6. Suggested Response Generation

One of the key features of the system is generating suggested replies. By understanding the email's content and context, the LLM crafts concise and relevant response options that the user can review and send with minimal editing. This saves users considerable time, especially when dealing with routine or repetitive emails.

7. Email Labeling and Management via Gmail API

Once the emails are categorized and prioritized, the system applies corresponding labels directly to the emails in the user's Gmail account using the Gmail API. These labels appear in the user's mailbox, facilitating easier navigation and filtering. The system can also archive, delete, or mark emails as read/unread based on user preferences or automated rules.

8. Query Handling and Interactive User Interface

Users can interact with the system by asking natural language queries related to their emails, such as "What is the deadline for the project mentioned in my last email?" The backend processes these queries through the LLM, which searches the stored email metadata and content to provide precise answers. The interaction can happen via a frontend interface or directly through API calls.

9. Data Storage and Persistence

The system maintains a database to store metadata and parsed email information, including labels, senders, timestamps, and key content snippets. This persistent storage enables fast retrieval during queries and supports learning from user interactions, enhancing future email classification and response accuracy.

10. Security and Privacy

Throughout the process, security is paramount. OAuth 2.0 ensures that the system never stores user passwords and only accesses Gmail data with explicit user permission. Communication between components is encrypted using HTTPS. Tokens are

securely stored and refreshed as needed, and personal data is retained only as long as necessary to provide the service, protecting user privacy.

INTELLIGENT MAIL MANAGEMENT SYSTEM USING LARGE LANGUAGE MODELS AND GMAIL API

FLOWCHART: USER INITIATES LOGIN TO CHATBOT INTERACTION

1. USER OPENS THE APPLICATION AND CLICKS ON "LOGIN"
- ↓
2. SYSTEM REDIRECTS USER TO GOOGLE AUTHENTICATION (OAUTH 2.0)
- ↓
3. USER AUTHENTICATES WITH GOOGLE AND GRANTS PERMISSIONS (GMAIL API ACCESS)
- ↓
4. SYSTEM VERIFIES THE ACCESS TOKEN AND FETCHES USER PROFILE
- ↓
5. SYSTEM CHECKS USER DATA AND MAIL ACCESS PERMISSIONS
- ↓
6. USER IS SUCCESSFULLY LOGGED IN TO THE SYSTEM
- ↓
7. SYSTEM LOADS DASHBOARD AND ACTIVATES CHATBOT INTERFACE
- ↓
8. USER CAN NOW START CHATBOT INTERACTION FOR MAIL MANAGEMENT

Fig.2 User Initiates Login to Chatbot Interaction

User Initiates Login

The user attempts to log in to the application interface.

Authenticate with Google via O auth

The system uses Google's OAuth protocol to authenticate the user. This allows secure access to the user's Gmail account without sharing their password.

Authentication Successful? (Decision Point)

Yes → Proceed to Step 4

No → Go to Step 14: Show Authentication Error

Fetch Emails using Gmail API

Once authenticated, the application retrieves the user's emails using the Gmail API.

Send Emails to Cloud-hosted LLaMA

The fetched emails are forwarded to a cloud-hosted LLaMA (Large Language Model by Meta AI) for processing.

LLaMA Analyzes Email Context

The LLaMA model performs natural language processing to analyze the content, context, topic, and sentiment of the emails.

Categorize and Label Emails Based on the analysis, emails are categorized and labeled (e.g., by topic, importance, urgency).

Return Labeled Emails to Frontend

The categorized emails are returned to the frontend application.

Display Categorized Emails to User

The frontend displays the labeled emails, organized for user interaction.

User Interacts with Chatbot The user initiates interaction with the integrated chatbot (e.g., by asking questions about emails).

Chatbot Processes Query

The chatbot processes the user's query to understand the intent.

Retrieve Relevant Email Data

Based on the query, relevant email data is retrieved from the LLaMA-processed and labeled content.

Return Answer to User

The chatbot returns a meaningful answer or insight to the user based on the retrieved email data.

Show Authentication Error (If Step 3 fails)

If Google OAuth authentication fails, an error message is displayed to the user.

End

The process terminates either after successful interaction or due to authentication failure.

IV. RESULT

Future enhancements could include expanding the system's multilingual support to cater to global users, enabling voice-based email queries and responses for hands-free operation, and integrating with other email platforms beyond Gmail. Incorporating user behavior learning would allow the system to personalize

categorization and prioritization dynamically. Additionally, advanced sentiment analysis could be added to flag emotionally sensitive or urgent emails. Integration with calendar and task management tools could further automate scheduling and follow-up actions, making the system a comprehensive productivity assistant. The system's query handling capability proved particularly effective in retrieving precise information from large volumes of emails, enabling users to quickly find critical details without manually scanning through their inbox. By leveraging the LLM's natural language understanding, users could ask complex questions in everyday language and receive accurate answers, improving the overall user experience. Furthermore, the seamless synchronization with Gmail ensured that all changes, such as labeling and categorization, were instantly reflected across devices, maintaining consistency and up-to date organization for users on the go.

V. FUTURE SCOPE

Future enhancements to the system could include the Intelligent Mail Management System leverages the power of Large Language Models and the Gmail API to revolutionize email handling by automating categorization, prioritization, and response generation. This reduces the time and effort users spend managing their inboxes, while enhancing accessibility and organization. Secure authentication and data privacy ensure user trust and system reliability. The proposed system presents a scalable and adaptable solution that addresses modern email overload challenges and has the potential to evolve with emerging AI technologies and user needs.

The future development of an intelligent mail management system using large language models and the Gmail API can focus on enhancing personalization through fine-tuned models that adapt to individual user preferences, writing styles, and communication patterns. This would enable the system to generate more accurate summaries and context-aware replies that closely resemble the user's natural tone, improving overall user experience and efficiency.

Another promising direction is the incorporation of multilingual capabilities, allowing the system to automatically translate, summarize, and respond to emails written in different languages. This would make the system more inclusive and globally

applicable, especially for users who frequently communicate across different regions and linguistic backgrounds.

The system can also evolve to provide deeper context awareness by leveraging historical email data and user interaction patterns. By understanding long-term context, the model can generate more meaningful responses and make better decisions regarding email categorization and prioritization, leading to smarter inbox management.

Integration with voice-based interfaces is another significant area of future enhancement. Users could interact with their emails through speech, enabling functionalities such as listening to emails, dictating replies, and managing inbox tasks hands-free, which would be particularly useful in mobile or accessibility-focused scenarios.

Expanding integration with other productivity tools such as calendars, task management systems, and collaboration platforms can further enhance the system's utility. Emails containing actionable information could automatically trigger task creation, meeting scheduling, or reminders, thereby streamlining workflow and reducing manual effort.

Future systems can also improve security by incorporating advanced AI-based spam and phishing detection mechanisms. With more sophisticated models, the system can identify malicious patterns in real time, providing users with enhanced protection against potential threats.

Real-time prioritization and predictive sorting of emails is another area for improvement. By analyzing user behavior and email content, the system can intelligently rank emails based on urgency and importance, ensuring that critical messages receive immediate attention.

Privacy and data security can be strengthened by exploring on-device or edge AI processing. This would minimize reliance on cloud-based systems and ensure that sensitive email data remains secure, addressing growing concerns around data privacy.

In addition, the incorporation of explainable AI techniques can make the system more transparent by providing clear reasons for its decisions, such as why a particular email was categorized in a certain way or why a specific reply was suggested. This can help build user trust and improve system reliability.

Finally, implementing a continuous learning mechanism through user feedback can significantly

enhance system performance over time. By adapting to user corrections and preferences, the system can evolve dynamically, ensuring that it remains effective and relevant as user needs change.

VI. CONCLUSION

The Intelligent Mail Management System leverages the power of Large Language Models and the Gmail API to revolutionize email handling by automating categorization, prioritization, and response generation. This reduces the time and effort users spend managing their inboxes, while enhancing accessibility and organization. Secure authentication and data privacy ensure user trust and system reliability. The proposed system presents a scalable and adaptable solution that addresses modern email overload challenges and has the potential to evolve with emerging AI technologies and user needs.

REFERENCES

- [1] G. Mujtaba, L. Shuib, R. G. Raj, N. Majeed, and M. A. Al-Garadi, "Email classification research trends: Review and open issues," *IEEE Access*, vol. 5, pp. 9044–9064, 2017, doi: 10.1109/ACCESS.2017.2702187.
- [2] S. Dash, "The next frontier of email efficiency with LLM," *Analytics Vidhya Blogs*, 2024. [Online]. Available: <https://www.analyticsvidhya.com>
- [3] K. I. Roumeliotis, N. D. Tselikas, and D. K. Nasiopoulos, "Next-generation spam filtering: Comparative fine-tuning of LLMs, NLPs, and CNN models for email spam classification," *Preprint*, 2024.
- [4] K. Debnath and N. Kar, "Email spam detection using deep learning approach," in *Proc. Int. Conf. Machine Learning, Big Data, Cloud and Parallel Computing (COM-IT-CON)*, Faridabad, India, 2022, pp. 37–41, doi: 10.1109/COM-IT-CON54601.2022.9850588.
- [5] Vaswani *et al.*, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. [Online]. Available: https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html
- [6] AI@Meta LLaMA Team, "The LLaMA 3 herd of models," *Meta AI*, 2024. [Online]. Available:

<https://ai.meta.com/llama>

- [7] n8n.io, "Automation workflows platform," 2024. [Online]. Available: <https://n8n.io>
- [8] Google Inc., "Gmail: A powerful, searchable inbox," 2004. [Online]. Available: <https://mail.google.com>