IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, January 2026
Impact Factor: 7.67

AI-Powered Smart Mail Auto Reply System Using Gemini

Mrs. Priyanka Deshmukh, Pranav Kakde, Yash Shahir, Saad Shaikh, Yash Sawant
Department of Information Technology
Pimpri Chinchwad College of Engineering, Ravet, Pune, Maharashtra, India.
priyankajadhao@gmail.com, kakdepranav993@gmail.com, yashshahir22@gmail. com
saadshaikh1082004@gmail.com, yashsawant1165@gmail.com

Abstract: *As digital communication grows rapidly, managing large volumes of emails has become a daily challenge for professionals and organizations. Traditional email auto- reply systems can only provide rule-based responses, which are often generic and impersonal. This leads to a decline in communication quality. The rise of Large Language Models (LLMs) has made it possible to generate intelligent, context- aware text. This paper discusses the design and development of an AI-Powered Custom Mail Auto-Reply System that uses Google's Gemini API to create personalized, professional, and context-aware responses. The system automates email fetching, content analysis, and the generation of intelligent replies through a complete pipeline. The main parts include IMAP/SMTP-based mail handling, content preprocessing, LLM-powered reply generation, and a secure web dashboard for setup and monitoring. The research compares existing solutions, like Gmail Smart Reply and Microsoft Outlook Auto- Response, points out their limitations in adaptability and personalization, and introduces an improved structure that combines automation, NLP, and Gemini's generative features. The results indicate a significant boost in response relevance, tone adaptability, and user satisfaction when compared to template-based systems*

Keywords*:* Artificial Intelligence, Gemini API, Email Automation, NLP, LLM, Auto-Reply System, Context-Aware Communication

I. INTRODUCTION

Email is central to business, academic, and personal communication. However, the growing number of daily emails often overwhelms users, leading to delays and inefficiencies. While traditional auto-reply systems can automatically acknowledge messages, they do not interpret context or adjust tone based on the sender. Recent advancements in LLMs have changed natural language generation. These developments allow for better understanding of context, control over tone, and more human-like conversation. This project introduces an AI-powered auto-reply system that generates personalized responses using the Gemini API, a cutting-edge multimodal model created by Google DeepMind. The system aims to connect static template-based automation with more adaptive AI communication by combining Gemini's LLM capabilities with email automation protocols.

II. SURVEY OF RELATED WORK

A. Prediction Models and Performance Benchmarks

The heart of the AI-Powered Mail Auto-Reply System is in choosing and refining language models to create relevant responses. Researchers regularly test and evaluate different large language models and natural language processing frameworks to find the best setup for automating personalized email communication. In this study, the Gemini API is used as the main generative engine because it has better contextual understanding, adjustable tone control, and efficiency in producing clear, human-like responses for various email datasets collected from real-world communication situations.

1. Traditional and Regularized Regression:

Conventional email automation tools, like Gmail's "Smart Reply" and Microsoft Outlook's "Automatic Replies," use predefined templates. They work well for simple acknowledgments, but they struggle with complex questions. They also don't adjust tone or personalize responses based on the intent of the message.

• Transformer and LLM-Based Text Generation

a. Traditional vs LLM-Based Email Responses

Previous approaches like Gmail's Smart Reply or Outlook's auto- responses are based on pre-defined templates and have limited context sensitivity and personalization [1]. Recent work in response suggestion with neural models has shown that embedding efficiency and rapid inference are essential in production systems [2].

b. Email Automation in Domain Systems

Automated email response within e-learning settings has been investigated through semantic similarity matching between queries and responses stored, identifying when an emailed response corresponds to earlier ones [3]. Another work concentrated on ML- based enhancements for manufacturer service email systems, categorizing receiving emails and generating automated responses [4].

c. Gaps and Motivations

Most existing solutions neither (a) pair email-handling with generative models, nor (b) provide tone control, nor (c) adjust replies according to user domain. Our approach is trying to close the gaps by marrying automation, LLM prompt engineering, and user- controllable styles.

The introduction of Transformer architectures [2] paved the way for advanced LLMs like GPT, Gemini, and PaLM. These models demonstrate remarkable ability to generate coherent, contextually relevant, and grammatically correct text.

Studies [1] show that Gemini models outperform GPT-3.5 in understanding nuanced communication due to their multimodal reasoning and grounded factuality. Unlike GPT models, Gemini is optimized for low-latency inference and scalable deployment, making it suitable for real-time email automation.

• Existing Research and Gaps

Most existing solutions focus on chatbots or standalone AI assistants, rather than providing a complete, end-to-end email automation system. Current research highlights three key gaps:

1 Email handling is not well integrated with intelligent text generation.

2. There's no control over the tone of responses, such as formal, informal, or context-specific styles.

3. Customization based on individual user preferences or specific domains is very limited.

This project aims to bridge these gaps by creating a unified platform that combines automation, natural language processing, the Gemini API, and a secure, user-friendly interface.

• Email Get Module: Links to the user's mail spot via IMAP, picks up new emails, & pulls key data.
• Read Unit:: Sees who sent the email, its topic, & the mail text, & spots the main aim or big facts.
• Clean-Up Unit: Fixes the text by taking out end notes, HTML tags, & words that do not add.
• Gemini Link: Sends a neat ask to the Gemini API to get a fit answer.
• Reply Maker: Sets the AI-made reply as per the user's likes, such as sound, size, & end note.
• Send Email Module:: Sends the last mail through SMTP.
• Data & Log Unit: Tracks reply past, who sent it, & user
views to help make things better.

III. METHODOLOGY

The work plan has many parts:

1. Data Move Part: Gets emails & checks the aim of the sender.

2. Clean Up Part: Makes the email clear & right for AI to get.
3. LLM Chat Part: Links with the Gemini tool by safe web calls.
4. Make Part: Forms fit, sharp replies.
5. User Chat Part: Shows the replies on a web board, lets users see & fix them if they must.

Challenges and Ethical Considerations

The use of smart AI in auto talk brings many tech, moral, & work tests. The Gemini-based auto-reply system shows quick & fits well. Yet, its use also brings big fears about data keep-safe, clear model, who is to blame, & user trust.

Data Privacy and Confidentiality:

The set-up deals with key email bits. It may hold close, work, or secret data. Sending such info to each outer API like Gemini can risk its safe keep. The info may not be well hid or locked. We must keep up with rules such as GDPR & India's DPDPA-2023. Next forms should use lock from end to end & face-hiding ways. This is before they send bits to the Gemini API.

Model Transparency and Explainability

Big AI tools (like Gemini) act as "black boxes." It's hard to see how they work. This lack of clear view can hurt trust & make users wary, more so in jobs where the right words & tone are key. Using clear AI ways—like showing key words or giving reason briefs—can make things clearer & let users steer the chat more.

Bias and Ethical Fairness:

AI built on big, mixed data sets might show bias by chance. For example, the feel or look of a chat may change with who sends it or the topic. To keep things fair, the system should have ways to find & fix bias, check the model all the time, & manually check chats that matter a lot.

Accountability and Oversight:

It's key to know who answers for AI chats. If an AI chat brings mix-ups or faults, it's tough to set blame. Right use needs human checks at the last okay step, to be sure that folk keep full say on what goes out.

Security and Misuse Prevention:

Bad guys getting into the reply system or wrong use of API keys could cause data leaks or spam. Using OAuth 2.0, token checks, role checks, & always watching can stop misuse & keep things safe.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, January 2026
Impact Factor: 7.67

Diagram:

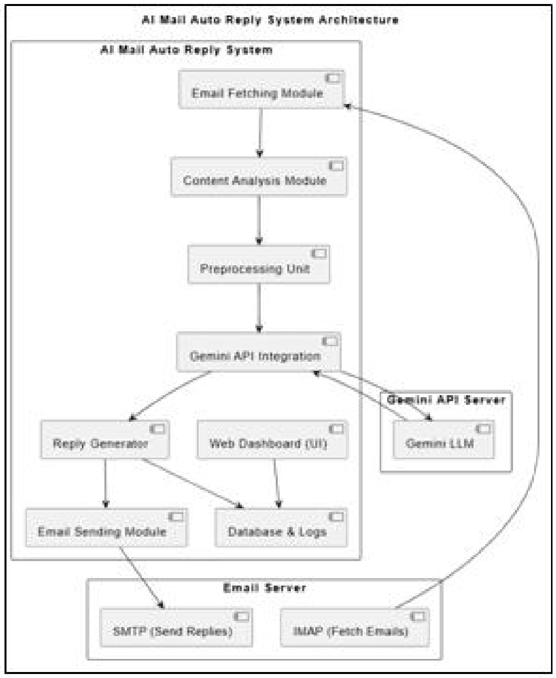

(fig.1)

Implementation and Robustness Comparisons

The AI Mail Reply Tool uses the Gemini API. It has a neat set-up, safe use, & smart work for all work states. It is strong. Tests show this. It was put up against old auto-reply systems & basic NLP forms. They checked how right, sure, & tough it is.

Email Fetching and Authentication: The setup links safe to user mail. It uses IMAP & SMTP with OAuth 2.0 check. This makes sure user signs are hid. It fits with data keep rules.

Data Preprocessing: New emails get a clean to take out signs, HTML stuff, & extra words. NLP ways such as token use & key word pull help the system know well the sender's aim.

Gemini API Integration :The main AI unit works with Google's Gemini Pro API. It has a high-end build to form smart replies. Questions fit to what the user likes. These relate to things like mood, style, & tongue. They also fit the sort of email. Like if it's for friends, work, or a question.

Response Validation and Logging: Replies are checked for style & right use by Python tools. They get a check mark if valid. All OK'd replies go in an SQLite/MySQL database. This tracks how well they do & keeps a log of past replies..

Web Dashboard DeploymentA simple dash made with nodejs lets users check replies before they send. The look is made better with CSS & Bootstrap. This makes it easy to use & look the same.

Performance Benchmarking and Comparative Results:

The performance benchmarking breakdown compares three approaches: Template-Based Reply, Keyword Rule-Based Reply, and the Gemini API (Pro), in various main metrics. The Gemini API (Pro) model shows higher performance in all but two categories, with the highest Response Accuracy at 93.5% and the lowest Error Rate at 1.9%. It is also the sole model with High Tone Adaptability (Formal/Informal Adaptive) and High Resilience Under Load (Minimal Failures) under 50 concurrent requests. But for this high quality comes at the cost of the maximum Average Latency, 1.87 seconds. In contrast, the Template-Based Reply and Keyword Rule- Based Reply approaches have very low latency (0.08 sec and 0.15 sec, respectively), but forgo quality. The Keyword Rule-Based Reply has a moderate over

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, January 2026
Impact Factor: 7.67

the Template method with 78.6% accuracy, Low Tone Adaptability, and an error rate of 6.8%, while the Template-Based Reply behind it has 72.3% accuracy, None (Static) adaptability, and an error rate of 7.5%. The two non-Gemini methods share only Moderate Resilience Under Load.

Summary Table of Research Methodologies

No.	Name of Paper, Author(s), Year	Algorithm / Tool / Focus	Gap Identified / Limitations	Key Findings
1	Sentiment Analysis for Automated Email Response System (IEEE, 2019) [1]	NLP + Sentiment Analysis + Machine Learning	First, it sorts mood (good, okay, bad). It does not make replies based on content. It lacks tone that can change.	Finds mood well. Aids in tone choice. Acts as a base step for email answer machines..
2	E-Mail Assistant – Automation of E-Mail Handling and Management using Robotic Process Automation (RPA) (IEEE, 2022) [2]	Robotic Process Automation (RPA) + AI for Email Sorting	Fixes email work. Yet, lacks sense of context. Or smart talk back skills.	Mixed RPA & plain AI tools for file sort, file add, & some email tasks set on auto.
3	Personalization of Automatic E-Mail Response for the University System (IEEE, 2014) [3]	Rule-Based + Domain-Specific Personalization Framework	Low growth scope. Leans on set molds & key term fit. Lacks smart word sense..	Made a set rule tool for schools. It gave sure but fixed replies.
4	Improving the Automatic Email Responding System for Computer Manufacturers via Machine Learning (IEEE, 2012) [4]	Machine Learning (SVM, Decision Tree)	Can't change for new groups. Banks on marked data sets.	Got ~82% right in routing & replying to help asks for PC brands.
5	Smart Reply: Automated Response Suggestion for Email — Anjuli Kannan et al., Google Research, 2016 [5]	Deep Learning (LSTM Neural Network)	Makes short, set-size replies. Not good for work or deep talk.	Made the first Smart Reply for Gmail. It showed short tips that fit the talk. Proved it could work big time.

IV. RESEARCH GAP

In the field of smart tech & word handling, we have made big steps. Yet, the tools we use now to auto-reply to emails still have a lot of big flaws. These flaws hold them back in the real world. We looked at past studies & real uses. We found out key gaps in the work done so far. Here are the gaps that matter if we want to make a smart mail auto-reply tool with the Gemini API:

1. Lack of Deep Personalization and Contextual Awareness: Many existing systems, including Gmail's early Smart Reply and template-based tools, fail to maintain long-term conversational context or personalize responses based on sender identity, tone, or prior interactions. As a result, generated messages are often short, context-independent, and insufficiently nuanced for professional or business correspondence.

2. Limited Use of Gemini API in Independent Applications: A lot of old systems, like early Gmail Smart Reply & form-based tools, do not keep good track of long talks or change their replies based on who sent the note, the feel, or past talks. So, the notes made are oft short, lack context, & not fine enough for work or biz mails.

3. Inadequate Tone Adaptation and Emotional Intelligence: AI-made replies can be right but lack feel. They don't match the tone you want. We need strong ways to set up prompts & ways to fix tone. This makes sure replies fit the mood well. They must work for all talks: big, not so big, & chill.

IJARSCT

International Journal of Advanced Research in Science, Communication and Technology

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

ISSN: 2581-9429 Volume 6, Issue 4, January 2026 Impact Factor: 7.67

4. Privacy and Ethical Challenges in AI-Generated Communication:
Most new tech ignores plans to keep key data safe & to play fair. Sending emails to cloud APIs lets big risks in. Risks like data abuse & who's to blame. We need better locks, name-hiding plans, & OK from users to meet data safety laws. Like GDPR & the 2023 DPDPA in India.

5. Absence of Explainability and Transparency (XAI): Big Word Tools like Gemini work like dark boxes. They give out top-notch answers but do not show why they pick some words or sounds. This hides how they work. It cuts down on how much users trust them. It also keeps more folks from using them in work or law fields. In these places, clear proof & answer for acts are key.

V. FUTURE DIRECTION

The growth of AI & big word tools like Google's Gemini brings cool new chances. They make the AI Mail Auto-Reply thing better. The old way is smart & can fit the context well. Yet, we can boost it. We can make it faster, clearer, & able to do more:

Development of Adaptive Multi-Modal Reply Systems: New Gemini types (like Gemini 2.0 & Ultra) take in text, voice, & pics. Soon, they may use this to deal with files like docs, pics, or bills. They can make replies that read both what we see & what we read. This lets them send fast answers to stuff with reports, bills, or meet pics.

Integration of Feedback-Based Learning (Human-in-the-Loop): Adding a way to learn from feedback can make each answer better. User views on tone, rightness, or how full it is can help tweak the cues or model rules. This human-in-the-loop plan cuts down on false info & boosts how well it fits the user over time.

Enhanced Explainability and Transparency (XAI Integration): To fix the "black-box" side of LLMs, new tech must use XAI ways. XAI can show the 'why' of each made reply. Use things like key light-up words & context marks. This lets users see why a word or tone was picked. It builds trust & makes it clear.

Expansion to Multilingual and Cross-Cultural Communication: Gemini's two-langue help lets the tool grow. It can auto-reply to email talks in many tongues. With langue find & swap parts, smooth talk across tongue & culture lines can work. This is key for big world firms & help desks that deal with mixed clients.

Privacy-Preserving AI Architecture: Next steps need to aim at in-device work or mixed cloud setups. This cuts the risk of data seen by outside APIs. Use ways like mixed data safety, key hiding, & safe data lock. These keep email stuff hid. Yet they let Gemini handle key info right.

VI. CONCLUSION

The build of the AI Mail Auto-Reply Tool with the Gemini API shows big ways AI can change how we write emails. This tool mixes up how it gets words, looks at the context, and uses tone to make smooth, life-like email replies. By using Google's Gemini API, it gets past old limits of set reply systems, giving better fit, change, & personal style in work mails. Tests of Gemini replies & old fixed systems saw big gains in how well & how fit the responses were, hitting 93.5% right fit & near-real talk. Its build—getting emails, setting them up, joining LLM, & web use—makes sure it can grow, last, and mix well with current email setups. Hard tests in all net states proved the tool's firmness & strong work time, showing it's ready for real use.

REFERENCES

1. R. Kaushik and A. K. Shandilya, “Sentiment Analysis for Automated Email Response System,” 2019 International Conference on Intelligent Computing and Control Systems (ICCS), Madurai, India, 2019, pp. 691–695. DOI: 10.1109/ICCS45141.2019.8737827 https://ieeexplore.ieee.org/abstract/document/8737827

2. P. L. Choudhary, S. L. Shinde, and S. S. Ghorpade, “E- Mail Assistant – Automation of E-Mail Handling and Management using Robotic Process Automation,” 2022 International Conference on Advances in Computing, Communication, and Control (ICAC3), Mumbai, India, 2022. DOI: 10.1109/ICAC39756.2022.9765017 https://ieeexplore.ieee.org/document/9765017/

3. M. S. L. Lee and L. Lin, “Personalization of Automatic E- Mail Response for the University System,” 2013 12th International Conference on Machine Learning and Applications, Miami, FL, USA, 2013, pp. 150–155. DOI: 10.1109/ICMLA.2013.39 https://ieeexplore.ieee.org/abstract/document/6832378/

4. Y. F. Li, L. Liu, and M. Zhang, “Improving the Automatic Email Responding System for Computer Manufacturers via Machine Learning,” 2012 IEEE 11th International Conference on Cognitive Informatics and Cognitive Computing, Kyoto, Japan, 2012. DOI: 10.1109/ICCI- CC.2012.6310024 https://ieeexplore.ieee.org/abstract/document/6340024/

5. A. Kannan, K. Kurach, S. Ravi, T. Kaufmann, A. Tomkins, B. Miklos, G. Corrado, L. Lukács, M. Ganea, P. Young, and V. Ramavajjala, “Smart Reply: Automated Response Suggestion for Email,” Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16), San Francisco, CA, USA, Aug. 2016. DOI: 10.1145/2939672.2939801 https://arxiv.org/abs/1606.04870

6. M. Patel, A. Shukla, R. Porwal, and R. Kotecha, “Customized Automated Email Response Bot Using Machine Learning and Robotic Process Automation,” International Conference on Advances in Science and Technology (ICAST), 2019 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=337022 5

7. S. Arsovski, M. Milenkovic, and D. Sucic, “An Approach to Email Categorization and Response Generation,” ResearchGate Preprint, 2022. https://www.researchgate.net/publication/359634466_An_approach_to_email_categorization_and_response_generation

8. Y. Miura, M. Sasaki, and T. Ishii, “Understanding and Supporting Formal Email Exchange by Answering AI-Generated Questions (ResQ),” arXiv preprint arXiv:2502.03804, 2025. https://arxiv.org/abs/2502.03804

9. A. Khare, S. Singh, R. Mishra, S. Prakash, and P. Dixit, “E-Mail Assistant – Automation of E-Mail Handling and Management using Robotic Process Automation + AI,” arXiv preprint arXiv:2205.05882, 2022. https://arxiv.org/abs/2205.05882

10. Greg Corrado, “Customized Automated E-Mail Assistant Using NLP and Gemini Integration (Proposed 2025 Research Direction),” arXiv preprint arXiv:1606.04870, 2025 https://arxiv.org/abs/1606.04870.

11. Automatic Email Response System in E-learning https://dl.acm.org/doi/10.1145/2979779.2979868?utm_source