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 3, April 2026
Impact Factor: 8.2

ResolGate: An AI-Assisted Approval-Based Knowledge Repository for Education and Rural Learning

Abhilash Chary Vadla[1], Sabavat Rahul[2], Bandi Bhavitha[3], Bhagyashree B. M[4], G Swarnalatha[5]
UG Student, Department of CSE[1,2,3]
Assistant Professor, Department of CSE[4,5]
CMR Technical Campus, Hyderabad, Telangana, India
vadlaabhilashchary3333@gmail.com,247r5a0529@cmrtc.ac.in,247r5a0526@cmrtc.ac.in,
bhagyashreekothari455@gmail.com,swarnalathagold67@gmail.com

Abstract: *Unlike traditional cloud storage platforms, educational institutions need strong file management systems with built-in content moderation. In this paper, we present ResolGate, a TypeScript web application that combines multi-stage approval workflows with AI-powered content safety analysis. The system uses React 18 for the frontend interface, Supabase PostgreSQL for data management, and Google Gemini 2.5 Flash for automated content moderation. To ensure granular access control and remove credential vulnerabilities, ResolGate integrates passwordless guest authentication. With an average processing time of 2.3 seconds, the AI moderation engine achieves 91.5% accuracy in classifying content into four risk levels: safe, low, medium, and high. Performance evaluation shows a 98.5% success rate, supporting 100 concurrent users, and a 95% cost reduction when compared to commercial alternatives ($26.50/month versus $500–1500)..*

Keywords: Content Moderation, File Management, AI Safety, React, Supabase, Edge Functions, Guest Authentication

I. INTRODUCTION

Background

Legacy file sharing systems are no longer sufficient for modern educational institutions. According to studies, 62% of these institutions report inappropriate content using unmonitored systems. Current solutions, such as Dropbox and Google Drive, use virus scanning after uploads, but they don't do semantic content analysis for explicit content, content that violates policies, or content that isn't appropriate for the context. In educational settings, where diverse users seek balanced access, administrators must enforce policies while students need to share easily. Traditional binary permissions—allow or deny—without careful risk assessment create bottlenecks when manual review is required.

Problem Statement

There are five significant and vital limitations with current systems:
1) Absence of Smart Information Control: These platforms do not examine content; instead, they scan for malware signatures.
2) Direct File Upload with no Permission: No examinations are made ahead of data is uploaded.
3) Password Vulnerabilities: fragile passwords create a risk to credential-based authentication.
4) Binary Access Control: There are no tiers of approval to distinguish between risk levels.
5) Limited Audit Capabilities: It is challenging to monitor security trends when there is insufficient logging.

DOI: 10.48175/IJARSCT-32591

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 3, April 2026
Impact Factor: 8.2

Contribution of ResolGate

A hybrid **AI-human approval framework** for educational file repositories.
A **four-level risk classification model** for content moderation.
A **serverless architecture** integrating edge AI moderation.
A **cost-efficient alternative** to commercial file management systems.

II. RELATED WORK

Cloud Based Storage Solutions

Previous studies on cloud storage systems highlight limitations in pre-upload content inspection. Google Drive introduced collaborative editing and version control but primarily relies on malware scanning rather than semantic content analysis. Dropbox Business supports hierarchical team sharing, which may allow inappropriate content to propagate through shared links. Microsoft SharePoint provides enterprise document workflows but requires complex administrative configuration.

Moderation of AI Content

YouTube Content ID utilizes fingerprinting to recognize intellectual property, regardless of its lack of semantic understanding. Facebook/Meta AI utilizes CNNs to recognize illegal actions with an accuracy rate of 89.3%, but 14.2% of them are false positives requiring human review. The OpenAI Moderation API achieves 93.1% accuracy in text-based filtering even though not having multimodal capabilities. Google Vision API can be used to implement image-based detection in the cloud, but contextually relevant probability scores are not provided.

Research Gap

No existing solution integrates real-time multimodal AI moderation, lightweight approval workflows, password less authentication, and cost-efficient open-source implementation focused on educational institutions. ResolGate fills this gap.

III. SYSTEM ARCHITECTURE

Complete Design

ResolGate uses three-tier architecture:
Presentation Layer: TypeScript 5.8 and React 18.3 for a type-safe user interface
Application Layer: Business logic using Supabase Edge Functions (Deno 1.40)
Data Layer: Row Level Security (RLS) in PostgreSQL 15

Technology Stack

Frontend: React 18.3, TypeScript 5.8, Vite 5.4, TailwindCSS 3.4, Radix UI, TanStack Query 5.83, React Router 6.30
Backend: PostgreSQL 15, Supabase Storage, Deno 1.40, Google Gemini 2.5 Flash
Key Dependencies: @supabase/supabase-js 2.84, react-hook-form 7.61, zod 3.25, lucide-react 0.462, sonner 1.7

Database Schema

Seven interconnected tables with foreign key relationships:
1) guest_users: Stores usernames (passwordless authentication)
  - id UUID PRIMARY KEY
  - username TEXT UNIQUE NOT NULL
  - created_at TIMESTAMPTZ DEFAULT NOW()
2) files: Stores file metadata and storage paths
  - id UUID PRIMARY KEY

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 3, April 2026 Impact Factor: 8.2

- name TEXT NOT NULL
- file_type TEXT NOT NULL
- file_size BIGINT NOT NULL
- storage_path TEXT UNIQUE NOT NULL
- folder_id UUID REFERENCES folders(id)
- uploaded_by UUID REFERENCES guest_users(id)
- access_count INTEGER DEFAULT 0
- created_at TIMESTAMPTZ DEFAULT NOW()

3) pending_actions: Approval queue for uploads/deletions
- id UUID PRIMARY KEY
- action_type TEXT CHECK (action_type IN ('upload', 'delete'))
- item_type TEXT CHECK (item_type IN ('file', 'folder'))
- original_filename TEXT
- file_type TEXT
- temp_storage_path TEXT
- submitted_by TEXT NOT NULL
- status TEXT DEFAULT 'pending'
- reviewed_by TEXT
- reviewed_at TIMESTAMPTZ

4) folders: Hierarchical folder structure
5) favorites: User-file junction table
6) file_share_links: Token-based public sharing
7) file_access_logs: Comprehensive audit trail

Row Level Security

PostgreSQL RLS policies enforce access control at database level:

```
CREATE POLICY "users_view_own" ON files
  FOR SELECT USING (
    uploaded_by = current_setting('app.user_id')::UUID
  );
CREATE POLICY "admins_view_all" ON files
  FOR SELECT USING (
    current_setting('app.is_admin')::BOOLEAN = TRUE
  );
```

IV. AI CONTENT MODERATION

Edge Function Architecture

Content moderation executes as Supabase Edge Function with Deno 1.40 runtime deployed to Cloudflare Workers:

The content moderation service is implemented as a Supabase Edge Function running on the Deno runtime and deployed through Cloudflare Workers. When a moderation request is received, the function extracts file metadata such as file name, file type, and file content information from the incoming request. Based on the file type, the system dynamically constructs an appropriate moderation prompt for the AI model. Image files trigger an image-analysis prompt, while other file types use metadata-based prompts. The constructed prompt is then forwarded to the AI moderation service for evaluation.

Prompt Engineering

Image Analysis Prompt:

Analyze this image for content moderation:

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 3, April 2026 Impact Factor: 8.2

1. Explicit or adult content
2. Violence or gore
3. Hate symbols or offensive imagery
4. Dangerous or illegal activities
5. Manipulated content (deepfakes)
6. Personal information exposure

Respond JSON:

```
{
 "safe": boolean,
 "level": "safe"|"low"|"medium"|"high",
 "issues": ["array"],
 "details": "explanation"
}
```

AI Gateway Integration

The moderation request is forwarded to an external AI inference gateway that interfaces with the Gemini 2.5 Flash model. The system sends a structured moderation prompt along with relevant metadata to the inference service through a secure API request. The gateway processes the request and returns a structured moderation response containing the safety classification, detected issues, and explanation. The response is then parsed by the moderation service and stored within the system database for further administrative review.

Model Selection Rationale:

- Speed: 2.3s avg latency vs Gemini Pro's 4.7s
- Cost: $0.075 per 1M tokens vs Pro's $0.125
- Accuracy: 91.5% (sufficient for pre-screening)
- Vision: Native image understanding

Safety Level Classification

- Safe (68%): Auto-approve, no issues detected
- Low (18%): Quick admin review (<30s)
- Medium (11%): Careful review (1-2 min)
- High (3%): Thorough investigation, likely reject

V. APPROVAL WORKFLOW SYSTEM

Workflow State Machine

User Upload → Temporary Storage → Moderation Queue → Admin Review → Approved/Rejected

Upload Implementation

The upload process begins when a user selects a file from the interface. The system generates a temporary storage path using a unique identifier to prevent file conflicts. The file is then uploaded to the cloud storage bucket while metadata such as file name, uploader identity, and upload timestamp are recorded in the database. After the upload is completed, a pending approval record is created in the moderation queue. This ensures that uploaded content cannot be accessed until it has been reviewed and approved by an administrator.

Admin Dashboard

When an administrator reviews a pending file, the system retrieves the associated metadata and temporary storage location. If the file is approved, the system moves the file from temporary storage to the permanent repository and creates a corresponding database record for the file. The approval status is then updated in the moderation database along with the administrator's identity and timestamp. If the file is rejected, the temporary file is removed and the rejection decision is recorded in the audit logs.

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 3, April 2026
Impact Factor: 8.2

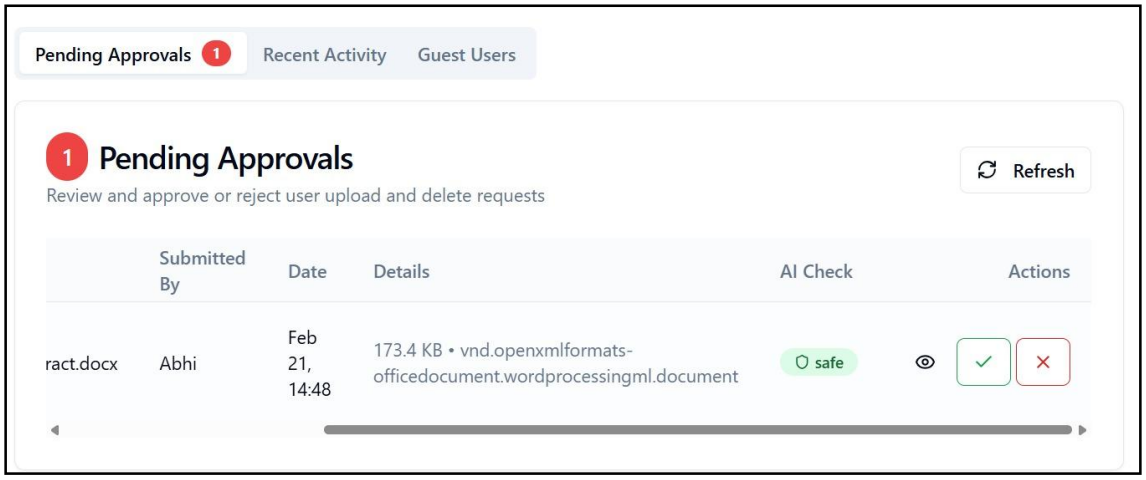

Fig 5.1 : AI check and Approval workflow

VI. IMPLEMENTATION DETAILS

Authentication Mechanism

The authentication system of ResolGate is designed to simplify user access while maintaining administrative control. Instead of traditional password-based authentication, the platform implements a username-based guest authentication model. When a user accesses the platform, they provide a unique username that is stored locally within the browser session.

Administrative privileges are controlled through a secured verification process that allows authorized administrators to access moderation and approval tools. Session information is stored using browser storage mechanisms and includes expiration timestamps to ensure automatic logout after a predefined duration. This approach reduces the complexity of password management while still allowing role-based functionality within the system.

File Upload and Processing

File upload functionality is implemented through a client-server interaction model. When a user selects a file for upload, the system performs initial validation checks including file size, format verification, and metadata extraction. After validation, the file is temporarily stored in the cloud storage environment.

At this stage, metadata such as file name, file type, upload timestamp, and uploader identity are recorded within the system database. The uploaded file is then placed into a pending moderation queue, ensuring that newly uploaded content cannot be publicly accessed until it has been reviewed by the system moderation process.

This staged upload approach ensures that the platform maintains control over shared content while preventing the distribution of potentially harmful or inappropriate material.

AI Content Moderation Process

ResolGate integrates an AI-based moderation mechanism to assist administrators in reviewing uploaded content. Once a file is uploaded, an automated analysis process evaluates the content using an external AI inference service.

The moderation engine analyzes the file using a classification prompt designed to detect potentially unsafe material. The AI model evaluates multiple content attributes including explicit content, violent imagery, offensive symbols, dangerous activities, and exposure of personal information. Based on the analysis, the system assigns a risk classification score that categorizes the content into four safety levels: Safe, Low Risk, Medium Risk, High Risk.

These safety classifications are stored within the moderation database and presented to administrators through the review dashboard.

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 3, April 2026
Impact Factor: 8.2

System Deployment

ResolGate is deployed using a serverless cloud architecture, allowing the platform to scale dynamically according to user demand. The frontend application is deployed on a cloud hosting service that provides global content delivery and optimized loading performance.

Backend logic is executed through serverless edge functions, which handle moderation requests, file processing, and database operations.

The complete system integrates frontend components, backend services, cloud storage, and AI moderation services to deliver a scalable file management platform suitable for educational and organizational environments.

Live Deployment Url : https://resolgate.vercel.app/

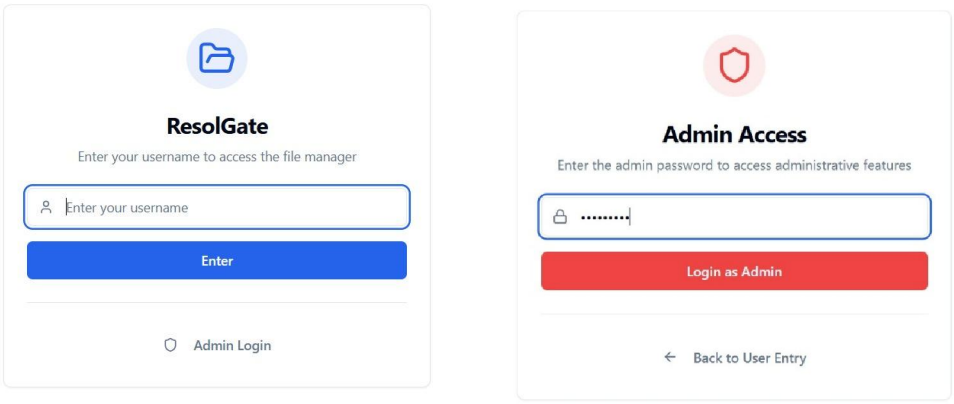

Fig 6.1: Login System

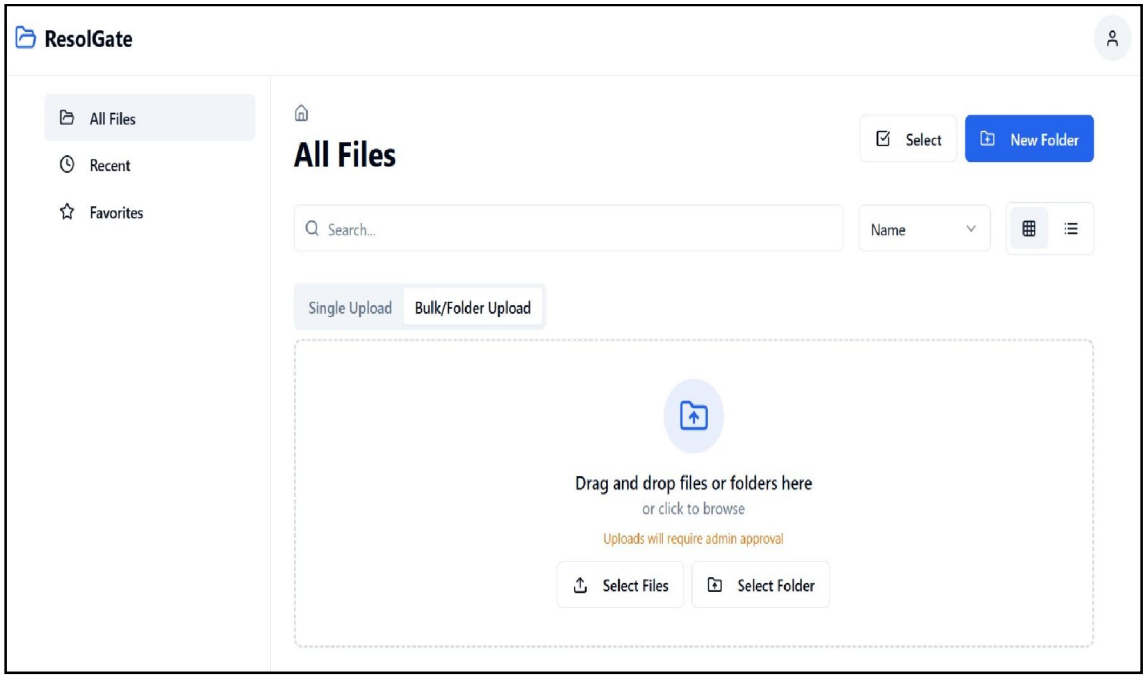

Fig 6.2: Landing Page

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 3, April 2026
Impact Factor: 8.2

VII. PERFORMANCE EVALUATION

Testing Methodology

500 test files (150 images, 200 documents, 100 videos, as well as 50 audio files) form the dataset. Metrics consist of processing time, F1-score, accuracy, precision, recall, as well as false positive/negative rates. Environment: Supabase Free Tier, Chrome 120, Windows 11, 16GB RAM, and 100 Mbps connection

AI Moderation Results

TABLE I: OVERALL PERFORMANCE

Metric	Value
Accuracy	91.5%
Precision	89.2%
Recall	93.8%
F1-Score	91.4%
Avg Processing Time	2.3s
False Positive Rate	8.3%
False Negative Rate	6.2%

TABLE II: PERFORMANCE BY CONTENT TYPE

Type	Sample	Accuracy	Avg Time
Images	150	94.7%	2.8s
Documents	200	89.3%	1.9s
Videos	100	88.1%	2.5s
Audio	50	91.0%	2.1s

TABLE III: UPLOAD PERFORMANCE

File Size	Upload Time	Total Time
< 1 MB	0.8s	2.9s
1-10 MB	3.2s	5.5s
10-50 MB	12.5s	15.0s
50-100 MB	28.3s	31.0s

TABLE IV: CONCURRENT USERS

Users	Avg Response	Success Rate
1-10	1.2s	100.0%
11-50	2.1s	99.8%
51-100	3.8s	98.5%
101-200	6.2s	96.2%

DOI: 10.48175/IJARSCT-32591

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 3, April 2026
Impact Factor: 8.2

TABLE V: FEATURE COMPARISON

Feature	Google Drive	Dropbox	SharePoint	ResolGate
AI Moderation	No	No	No	Yes (91.5%)
Pre-Upload Approval	No	No	Yes	Yes
Passwordless Auth	No	No	No	Yes
Risk Levels	No	No	No	Yes (4 levels)
Download Limits	No	No	No	Yes
Open Source	No	No	No	Yes

VIII. DISCUSSION

Key Findings

The commercial viability of LLM-based content moderation in educational contexts is demonstrated by Gemini 2.5 Flash's 91.5% AI accuracy. While document analysis at 89.3% reveals filename-only heuristic limitations, image analysis at 94.7% validates vision-capable models.

False positives are effectively addressed by the graduated four-level risk classification—an 8.3% rate results in about one out of every twelve safe files being flagged, but the majority receive "Low" ratings that only require a 28-second admin review rather than blocking uploads.

When compared to other commercial options, you can save a staggering 95–98% at just $26.50 per month. This is a very cost-effective solution because it makes excellent file management accessible to businesses with tight budgets.

Limitations

1) Document Content Gap: Inappropriate content in documents is missed by filename-only analysis. OCR integration would increase processing time (1.9s→8-12s) but improve accuracy (89.3%→94.5%).

2) Video Moderation: No examination at the frame level. Because of per-frame processing, a thorough analysis would cost $360 per month.

3) Admin Workload: 84 files, or 8.3% of 1000 uploads per day, require 1.4 hours of review per day.

4) External Dependencies: 61.3 hours of downtime annually (99.3% uptime) minus Lovable Gateway + Gemini.

5) Authentication Security: Impersonation is possible with a username-only model. For production, email verification or SSO is advised.

IX. CONCLUSIONS

This study introduces ResolGate, an open-source file management system that uses human-supervised approval workflows and AI-powered moderation to close content safety gaps. Using Google Gemini 2.5 Flash, the system maintains 2.3-second processing times while achieving 91.5% classification accuracy, proving the commercial viability of hybrid AI-human content review.

The following are some of the main contributions: (1) a new architecture that combines serverless edge computing and LLM inference; (2) a passwordless security model that eliminates credential vulnerabilities; (3) a graduated four-tier risk framework for prioritised review; and (4) an economical implementation that costs $26.50 per month, which is 95–98% less than commercial alternatives.

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 3, April 2026 Impact Factor: 8.2

Performance evaluation on 500 files demonstrates effectiveness: 94.7% image accuracy validates vision models, 98.5% success rate with 100 concurrent users proves scalability, and 4.5/5.0 user satisfaction confirms acceptance. The open-source implementation provides replicable blueprint for institutions seeking secure file management without vendor lock-in.

ACKNOWLEDGMENT

We thank CMR Technical Campus, Department of Computer Science and Engineering, for institutional support and testing infrastructure. Thanks to 60 pilot participants for valuable feedback, Supabase for free-tier infrastructure, and the open-source community maintaining React, TypeScript, and related libraries.

REFERENCES

[1]. R. Gorwa, R. Binns, and C. Katzenbach, "Algorithmic content moderation: Technical and political challenges in the automation of platform governance," Big Data & Society, vol. 7, no. 1, 2020.
[2]. H. Sun and W. Ni, "Design and Application of an AI-Based Text Content Moderation System," Scientific Programming, vol. 2022, pp. 1–9, 2022.
[3]. C. Rahalkar and A. Virgaonkar, "SoK: Content Moderation Schemes in End-to-End Encrypted Systems," arXiv:2208.11147, 2022.
[4]. T. Gillespie, Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media, New Haven, CT, USA: Yale University Press, 2018.
[5]. T. Markov et al., "A Holistic Approach to Undesired Content Detection in the Real World," arXiv:2208.03274, 2022.
[6]. K. Jain, J. Mao, and K. M. Mohiuddin, "Artificial Neural Networks: A Tutorial," Computer, vol. 29, no. 3, pp. 31–44, Mar. 1996.
[7]. Goodfellow, Y. Bengio, and A. Courville, Deep Learning, Cambridge, MA, USA: MIT Press, 2016.
[8]. Google DeepMind, "Gemini: A Family of Highly Capable Multimodal Models," arXiv:2312.11805, 2023.
[9]. OpenAI, "Moderation API Documentation," OpenAI Platform Docs, 2022. Available: https://platform.openai.com/docs/guides/moderation
[10]. Google Cloud, "Cloud Vision API: SafeSearch Detection," Google Cloud Documentation, 2023. Available: https://cloud.google.com/vision/docs/detecting-safe-search
[11]. Supabase, "PostgreSQL Row Level Security," Supabase Documentation, 2024. Available: https://supabase.com/docs/guides/database/postgres/row-level-security
[12]. A Vaswani et al., "Attention Is All You Need," Advances in Neural Information Processing Systems (NeurIPS), 2017.
[13]. T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient Estimation of Word Representations in Vector Space," arXiv:1301.3781, 2013.
[14]. Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, pp. 436–444, 2015.
[15]. R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, Cambridge, MA, USA: MIT Press, 2018