Indian Journal of Computer Science and Technology
Volume 5, Issue3 (September-December 2026), PP: 709-714.
www.indjcst.com
ISSN No: 2583-5300

Zeno AI: A Firebase-Based Approach to Automated Advertisement Generation and Publishing

Swapnil Subhash Harale[1], Sakib Samir Nadaf [2], Mohammad Siddik Paigambar Mulani[3], Samir Shafiq Mushrif [4], Athrav Vijay Chavan[5]
[1]Professor, Shri, Ambabai Talim Sanstha's Sanjay Bhokare Group of Institute, Miraj, Maharashtra, India.
[2,3,4, 5]Ambabai Talim Sanstha's Sanjay Bhokare Group of Institute, Miraj, Maharashtra, India.

To Cite this Article: *Swapnil Subhash Harale[1], Sakib Samir Nadaf [2], Mohammad Siddik Paigambar Mulani[3], Samir Shafiq Mushrif [4], Athrav Vijay Chavan[5] "Zeno AI: A Firebase-Based Approach to Automated Advertisement Generation and Publishing", Indian Journal of Computer Science and Technology, Volume 05, Issue 03 (September-December 2026), PP: 709-714.*

Copyright: ©2026 This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution License; Which Permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: *Small businesses, individual creators and early-stage brands need advertisements for festive sales, product launches and promotions, yet often lack designers, copywriters or a video-production budget. Producing one advertisement normally involves several separate steps: branding, creative design, copywriting and publishing. This is time-consuming, costly and often results in inconsistent branding. This paper presents the development of Zeno AI, an AI-based advertisement generator built on the Firebase platform (Authentication, Firestore, Cloud Storage and Cloud Functions). The platform offers a single guided workflow in which users can generate a brand logo, an AI character or model, a matching background scene and an advertisement script in one fast generation cycle, review each element, assemble the approved elements into a finished image or short video, and publish or schedule it to linked social media accounts. Previous work in this field has largely focused on isolated design tools or standalone image generators, lacking an integrated, affordable pipeline from brief to published advertisement. Our project aims to reduce the effort of advertisement production by combining generative AI with human-in-the-loop review and direct publishing. The paper outlines the design considerations, development methodology, testing procedures and potential impact of the proposed solution, and discusses directions for future research.*

Key Word: *Advertisement Generation, Generative AI, Text-to-Image, Content Generation, Firebase, Cloud Functions, Social Media Publishing, Web Application.*

I.INTRODUCTION

Social media platforms such as Instagram, Facebook, YouTube, LinkedIn and X (Twitter) have become primary channels for advertising, branding and customer outreach. For a small business or creator to remain visible, fresh and professional-looking advertisements are needed regularly, often at short notice for a festive sale, a flash offer or a local event. However, hiring a freelance designer or agency is expensive, while doing the work manually, by finding stock photos, writing captions separately and editing everything together in a design tool, is slow, fragmented and error-prone.

This paper addresses these challenges by presenting a comprehensive overview of our project: the development of Zeno AI, an AI-based advertisement generator application built on Firebase. Our endeavor stems from the observation that logo design, character creation, background selection, copywriting and publishing are usually handled by separate tools. Generic AI image generators can produce a single picture, but they do not understand advertising structure: they do not automatically pair a suitable character, a matching background and a persuasive script, and they do not connect the final creative to the business's social media accounts.

Previous work in this area has mainly involved either template-based design tools that still require manual editing, or generative AI tools that produce isolated assets but cannot assemble or publish a complete advertisement. While both have improved parts of the workflow, they often lack integration, brand consistency and an affordable unified experience. The purpose of our project is twofold: firstly, to create a centralized and user-friendly platform on which a user can move from an idea to a shareable advertisement with minimal manual editing. Secondly, we aim to use a serverless Firebase backend together with generative AI services to provide features such as logo generation, character generation, background generation, script and caption writing, automated ad assembly, direct publishing and performance analytics.

The design goal of Zeno AI is speed. The platform generates the character, background and script together in a single coordinated cycle, so that the user reviews one complete draft advertisement rather than several disconnected assets. A like-or-regenerate checkpoint keeps the user in control, so that nothing is rendered into its final form or published without explicit approval.

The remainder of this paper is organized as follows: the Literature Review summarizes related work; the Modules section describes the main components of the system; Design Considerations and Development Methodology explain how the platform was planned and built; Testing and Validation describes the verification procedures; and the Conclusion discusses the impact, limitations and future scope of the project.

II. LITERATURE REVIEW

The literature on digital advertising and AI-assisted content creation offers valuable insight into the challenges, trends and best practices of the field. Studies on social media usage show that organizations benefit from consistent engagement, audience understanding and timely communication [1]. Template-based design tools such as Canva popularized the idea of a guided design canvas with drag-and-drop brand assets. Zeno AI studies this interaction model but replaces manual template editing with brand-matched creative elements generated automatically from a short text brief.

Recent progress in generative modelling has transformed automated image creation. Denoising diffusion probabilistic models showed that high-quality images can be synthesized by learning to reverse a gradual noising process [2], and latent diffusion models made high-resolution text-to-image synthesis computationally practical [3]. These models provide the conceptual basis for the Character Generator and Background Generator modules, which produce a photorealistic or stylised image from a structured prompt describing subject, pose, setting and style.

In natural language processing, the Transformer architecture introduced attention-based sequence modelling and became the foundation of modern language models [4]. Large generative models showed that a single model can produce fluent text for many tasks with only a few examples in the prompt [5]. A model conditioned on product attributes, tone and target audience can therefore produce headlines, body copy and captions comparable in structure to human-written advertising copy, which motivates the prompt design of the Script Generator module. General treatments of artificial intelligence and deep learning are given in [10] and [14].

Industry literature on generative advertising creative describes tools that produce headline, image and call-to-action variants automatically and then rank or A/B test them. Zeno AI adapts this idea at a smaller, single-brand scale suitable for a diploma-level implementation.

From the engineering perspective, REST-based architectural principles guide the design of scalable web services through which the client, server and third-party platforms communicate [6]. Firebase documentation describes a serverless, real-time NoSQL architecture suited to applications that must store user accounts, structured campaign data and binary media while scaling automatically [7]. Publishing to social platforms relies on OAuth 2.0 authorization [8] and the official publishing APIs of each platform, such as the Meta Graph API and the YouTube Data API [9].

Despite these advancements, challenges remain. These include maintaining the quality and originality of AI-generated media, avoiding depictions of real identifiable people, clearly labelling synthetic content, complying with platform API policies and rate limits, handling access tokens securely, and keeping generation latency and cost under control. In summary, the literature indicates that no single existing consumer tool combines logo generation, AI character generation, background generation, ad-copy generation and direct multi-platform publishing in one guided workflow aimed at fast advertisement production for small businesses. The contribution of Zeno AI is this end-to-end orchestration rather than any single generative model.

III. MODULES

Homepage and Onboarding Module: Introduces the platform, guides a new user through account creation and first brand-profile setup (business name, industry, colours, tone of voice, logo preference), and offers a template gallery of layouts such as square posts, stories and short videos.

Authentication and User Management Module: Handles sign-up, login and session management using Firebase Authentication with email/password or social sign-in. Password storage and verification are handled by the managed identity service, and each user session is identified by a Firebase ID token used for all subsequent API calls.

Social Account Integration Module: Users connect Instagram, Facebook, YouTube, LinkedIn and X through OAuth 2.0 and the respective publishing APIs [9], [17], [18]. Access tokens are stored encrypted and never exposed to the client. Users can view, refresh or revoke linked accounts, and publish or schedule a finished advertisement per platform.

AI Brand and Logo Generator Module: Produces several logo variations from the business name, industry, preferred colours and a style tag (minimal, playful, luxury, bold, tech and so on). The selected logo and its colour palette are stored in the brand profile so that later generations stay visually consistent.

AI Character Generator Module: Creates the human model, mascot or product-hero figure for the advertisement from the product category, target audience, pose or expression and brand tone. Multiple options are produced per request. Characters are synthetic, content-safety filtering is applied to prompts and outputs, and the system avoids copying or impersonating any real person.

AI Background and Scene Generator Module: Produces the setting for the advertisement, such as a studio, outdoor lifestyle, festive, product-shelf or fully custom scene. Perspective, lighting and colour tone are kept compatible with the chosen character so that the composited result looks coherent.

AI Script and Copy Generator Module: Writes a headline, supporting body copy, a call-to-action and platform-specific captions and hashtags, based on the offer description, tone, platform and promotional details such as discount, deadline or location. Multiple variants can be generated, and the copy can be liked or regenerated independently of the visuals.

Fast Ad-Generation Orchestration Module: The coordination layer of the system. A single Generate Ad action triggers Cloud Functions that call the character, background, script and, if needed, logo generators in parallel and return a combined draft once all elements are ready. This module retries failed sub-generations and tracks generation cost against the user's credit balance.

Ad Assembly and Rendering Module: Selects a template matching the target platform's aspect ratio, composites the approved character over the approved background, overlays the logo, headline and call-to-action using the brand's colours and fonts, optionally adds simple motion and caption timing for short videos, and renders a final PNG, JPEG or MP4 file stored in Firebase Storage. Prior versions are retained so the user can revert.

Review, Approval and Publish Module: Presents the assembled advertisement for a final check, then lets the user download it, publish it immediately or schedule it. Every attempt is recorded in a publish log with status, platform post ID and timestamp.

Dashboard Module: Shows campaigns with their status (draft, generating, awaiting review, assembled, published), credit usage, recent activity, a quick-create entry point and a brand switcher for users with several brands.

Analytics Module: Tracks the number of generations, the regenerate rate per element type and the average time to first draft, as well as successful versus failed publish attempts. Where a platform API allows, it retrieves basic engagement figures such as views, likes and clicks for a published advertisement.

IV. DESIGN CONSIDERATIONS

The design of Zeno AI was guided by several key considerations, including usability, speed, scalability, security, brand consistency and responsible AI use.

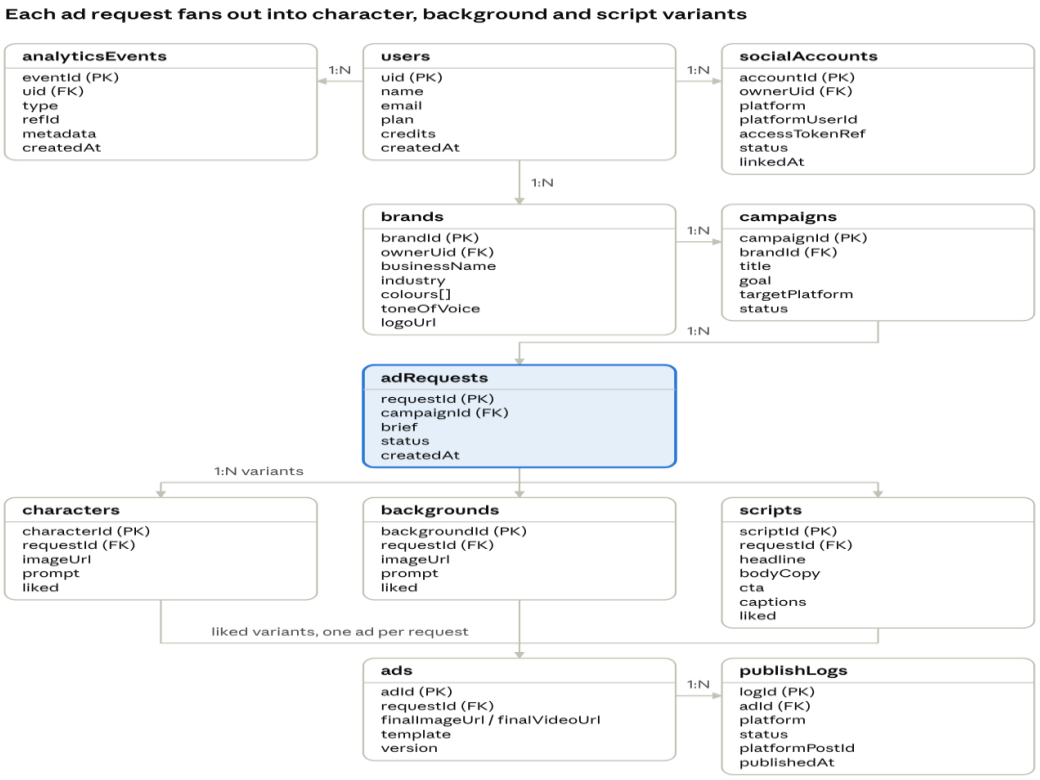


Fig. 1. ER diagram of the Zeno AI Firestore collections; 1:N marks a one-to-many relationship.

ER diagram · 11 Firestore collections

Usability: A primary objective was an intuitive, mobile-first interface that allows non-designers to produce a complete advertisement from a short text brief in very few steps. The like/regenerate control on every element and a combined draft preview were prioritized.

Speed: Character, background and script generation are issued as parallel Cloud Function calls so that all elements are ready at nearly the same time. Regenerating a single element does not require redoing the whole creative.

Scalability: The serverless Firebase backend scales automatically from a small number of student users to a larger user base without a self-managed server. Firestore's document model suits the fast-changing, media-heavy data of a generative application, while Cloud Storage holds the generated images and videos.

Security: User credentials are handled by Firebase Authentication, social-media tokens are stored encrypted and referenced rather than embedded in documents, API keys for AI services are kept on the server and never exposed to the client, and Firestore security rules restrict each user to their own data.

Brand Consistency: Every generation is conditioned on the stored brand profile (logo, colours, tone and industry), so that advertisements remain visually and verbally consistent across campaigns and platforms.

Responsible AI Use: Generated characters are synthetic and not intended to represent any real person. Rendered advertisements carry an AI-generated marker where required, generated scripts are treated as drafting assistance to be checked for accuracy and platform advertising policies, and a human review step is required before any advertisement is finalised or published.

Data Model: Data is organised into Firestore collections for users, brands, social accounts, campaigns, ad requests, characters, backgrounds, scripts, ads, publish logs and analytics events. A user owns one or more brands; a user links social accounts, and a brand has campaigns; a campaign contains ad requests; each ad request produces character, background and script variants, of which the liked ones are used to assemble one advertisement with possibly several rendered versions. The schema follows standard database design principles [13].

V. DEVELOPMENT METHODOLOGY

The development of the platform follows an iterative, modular and cloud-native approach, allowing continuous feedback and refinement throughout the process.

Requirements Gathering: Requirements were derived from the needs of small businesses, individual creators and event promoters, and were prioritized and translated into modules and feature specifications.

Firebase Implementation: Firebase was chosen because it combines authentication, a real-time NoSQL database, object storage and serverless functions in one managed platform. Firebase Authentication and the brand-profile module are built first, together with the Firestore and Storage data layer. The mobile client is built with Flutter [11], [15], and the overall process follows standard software-engineering practice [12].

AI Integration: The Logo, Character, Background and Script generators are integrated behind Cloud Functions as integration-ready components accessed through established generation APIs. Carefully designed prompts incorporate the brand profile, product, audience, tone and target platform, and generated outputs are filtered for content safety before being returned to the user. The original contribution of the project is the orchestration pipeline, the data architecture, the brand-consistency logic and the publishing integration, not a new generative model. Advertisement scripts are generated through the Gemini API [16].

Ad Assembly Development: The like/regenerate review step and the compositing and rendering engine are built and tested end-to-end with sample briefs.

Social Publishing Integration: Account linking and publishing are integrated and tested against each platform's sandbox or test environment.

Deployment: The final step involves deploying the client and Cloud Functions on Firebase, configuring environment variables and secure storage for keys, and testing the complete journey from sign-up to brief entry, generation, review, assembly and publishing.

VI. TESTING AND VALIDATION

Testing and validation are integral parts of the development process, aimed at ensuring the reliability and robustness of the platform.

Unit Testing: Individual components such as authentication flows, prompt builders, cost and credit calculations and template selection are tested in isolation across different scenarios.

Integration Testing: Integration tests validate the interaction between modules, including Cloud Functions, Firestore reads and writes, Storage uploads, generation API responses and the publishing workflow.

Generation Quality and Latency Testing: Each generator is tested for output quality, consistency with the brand profile and time to first draft, including the behaviour of parallel generation and retries when a sub-generation fails.

Assembly Correctness Testing: Rendered advertisements are checked for correct layout, aspect ratio, legibility of overlaid text, logo placement and output file integrity for both image and video formats.

Publishing Reliability Testing: Publishing and scheduling are tested against each platform's sandbox, including token refresh, revoked access and failed publish handling.

Usability and Security Testing: Sample users from the target groups evaluate ease of use and quality of generated advertisements, while access rules are checked to confirm that users can only read and write their own data.

VII. SYSTEM ARCHITECTURE

The system follows a cloud-native, three-tier architecture. The User interacts with the web or mobile client (dashboard, brief form, review screen, analytics). Requests are authenticated through Firebase Authentication and sent to Cloud Functions, which communicate with three main components: Cloud Firestore and Cloud Storage for data and generated media, the AI generation layer (logo, character, background and script generators), and the Ad Assembly service. After the user approves the result, the publishing layer calls the Social Media APIs to publish or schedule the advertisement and later retrieve engagement data. Management-oriented data structures support moderation and credit control, and a dedicated admin dashboard is planned as a production extension.

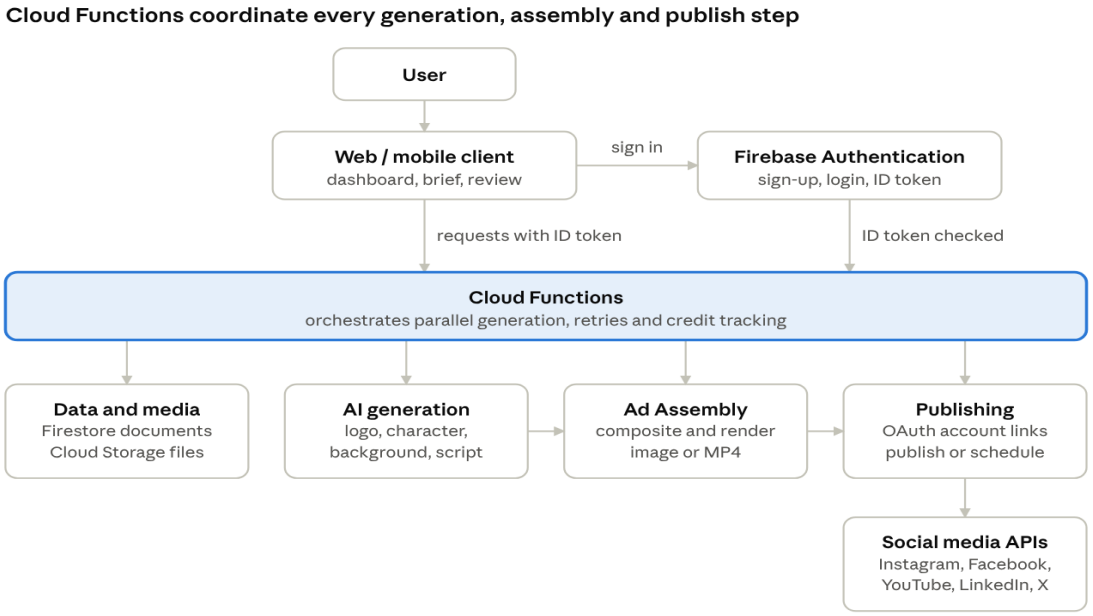


Fig. 2. System architecture of Zeno AI.

System architecture · 4 services

VIII. CONCLUSION

In conclusion, Zeno AI represents a significant step toward simplifying advertisement production for small businesses and creators. Throughout this paper, we have outlined the design considerations, development methodology and key features of the platform, highlighting its potential to reduce the time, cost and skill required to move from an idea to a published advertisement.

The advantages of the system are manifold. By combining logo, character, background and script generation with automated assembly and direct publishing in one guided workflow, users can create, review and publish advertisements without switching between tools. The Firebase backend provides secure authentication, scalable storage and serverless orchestration, and the stored brand profile keeps every campaign consistent.

However, it is important to acknowledge the limitations of our approach. The quality of generated characters, backgrounds and scripts depends on the underlying generation services and prompts, and human review remains necessary for factual accuracy, legal claims and platform advertising policies. The platform also depends on third-party generation and social media APIs, whose policies, rate limits, pricing and access requirements may change. The current scope does not include a proprietary, production-grade image or video model, nor automated ad-spend management.

Looking ahead, there are several avenues for future work. These include higher-quality AI video with voice-over and music, A/B testing and ranking of generated variants, prediction of best posting times using machine learning, multilingual and regional-language advertisements, team collaboration and approval workflows, a dedicated admin dashboard, support for additional platforms, and automated ad-spend management. In summary, the proposed platform has the potential to become a valuable tool for individuals, small businesses and organizations seeking fast, consistent and affordable advertising.

Acknowledgements

We extend our sincere gratitude to all those who contributed to the realization of this project. First and foremost, we express our deepest appreciation to our project guide, Prof. Swapnil Subhash Harale, for his invaluable guidance, support and encouragement. We are also indebted to the faculty members of the Department of Artificial Intelligence & Machine Learning at Shri Ambabai Talim Sanstha's Sanjay Bhokare Group of Institute, Miraj, for their valuable insights and feedback during the development process. We thank our peers and the technical staff of the institute for their assistance and infrastructure support. Last but not least, we express our gratitude to our families and friends for their unwavering support and motivation.

REFERENCES

1. J. H. Kietzmann, K. Hermkens, I. P. McCarthy, and B. S. Silvestre, "Social media? Get serious! Understanding the functional building blocks of social media," Business Horizons, vol. 54, no. 3, pp. 241–251, 2011.
2. J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in Advances in Neural Information Processing Systems (NeurIPS), vol. 33, 2020, pp. 6840–6851.
3. R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2022, pp. 10684–10695.
4. A. Vaswani et al., "Attention is all you need," in Advances in Neural Information Processing Systems (NeurIPS), 2017.
5. T. B. Brown et al., "Language models are few-shot learners," in Advances in Neural Information Processing Systems (NeurIPS), 2020.
6. R. T. Fielding, "Architectural styles and the design of network-based software architectures," Ph.D. dissertation, Univ. of California, Irvine, 2000.
7. Google, "Firebase Documentation: Authentication, Cloud Firestore, Cloud Storage for Firebase and Cloud Functions," [Online]. Available: https://firebase.google.com/docs
8. D. Hardt, Ed., "The OAuth 2.0 Authorization Framework," RFC 6749, IETF, Oct. 2012.
9. Meta Platforms, "Instagram Platform and Graph API Documentation," [Online]. Available: https://developers.facebook.com/docs; Google, "YouTube Data API," [Online]. Available: https://developers.google.com/youtube/v3
10. S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ, USA: Pearson, 2020.
11. E. Windmill, Flutter in Action. Shelter Island, NY, USA: Manning Publications, 2019.
12. Sommerville, Software Engineering, 10th ed. Harlow, U.K.: Pearson, 2016.
13. A. Silberschatz, H. F. Korth, and S. Sudarshan, Database System Concepts, 7th ed. New York, NY, USA: McGraw-Hill, 2019.
14. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.
15. Google, "Flutter Documentation," [Online]. Available: https://docs.flutter.dev
16. Google, "Gemini API Documentation," [Online]. Available: https://ai.google.dev/gemini-api/docs
17. LinkedIn Corporation, "LinkedIn Marketing and Share API Documentation," [Online]. Available: https://learn.microsoft.com/en-us/linkedin.