© JUN 2026 | IRE Journals | Volume 9 Issue 12 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1718522

Advantage Gen – An AI-Powered Automated Social Media Campaign Studio

YASH BADGUJAR[1], SHRIPAD S. BHIDE[2]
[1]*MCA Student, MCA Department, P.E.S. Modern College of Engineering, Pune, India*
[2]*Project Guide, MCA Department P.E.S. Modern College of Engineering, Pune, India*

Abstract- ***The rapid growth of digital marketing has created an urgent demand for intelligent, automated content generation tools. Conventional social media marketing workflows rely on multiple disconnected tools—graphic design platforms, copywriting utilities, hashtag planners, and brand management systems—that collectively impose high time, cost, and coordination burdens. We present AdVantage Gen, an AI-powered Automated Social Media Campaign Studio that integrates image generation, caption synthesis, hashtag optimisation, brand overlay, and platform-specific formatting within a single unified pipeline. The system accepts a natural-language prompt and autonomously produces complete, platform-ready marketing assets using multi-modal Artificial Intelligence. The stack comprises React for the frontend, Node.js/FastAPI for the backend, Hugging Face Diffusion Models for image generation, and Google Gemini for text synthesis. Evaluation demonstrates a 100% functional test pass rate, campaign assembly within 10–15 seconds, and a feature profile superior to representative incumbent platforms.***

Keywords—Generative AI; Social Media Automation; Diffusion Models; Natural Language Processing; Marketing Technology; Prompt Engineering; Full-Stack Web Development; Google Gemini; Hugging Face.

I. INTRODUCTION

In today's digital era, social media marketing has become one of the most critical channels for brand promotion and customer engagement. Businesses operating on platforms such as Instagram, LinkedIn, Facebook, and Twitter require a continuous stream of high-quality campaign assets—images, captions, and hashtags—tailored to each platform's specifications. Producing these assets manually involves sequential coordination among graphic designers, copywriters, and marketing managers, making campaign creation slow, expensive, and difficult to scale [1].

The rapid advancement of Artificial Intelligence—particularly Diffusion-based image synthesis [2] and Large Language Models (LLMs) [3]—has created an unprecedented opportunity to automate this pipeline. Generative models can now produce high-fidelity visual content from textual prompts, while LLMs generate persuasive marketing copy with contextual awareness. Together, these technologies enable a shift from multi-tool, multi-day campaign production to a single-prompt, near-instantaneous workflow.

This paper presents AdVantage Gen, an end-to-end AI Campaign Studio that integrates multi-modal generative AI within a cohesive full-stack web application. The remainder of the paper is structured as follows: Section II reviews related work; Section III describes the system architecture and methodology; Section IV details implementation; Section V presents results; Section VI discusses applications and advantages; and Section VII concludes with future directions.

II. RELATED WORK

A. Conventional Marketing Tools

First-generation tools such as Canva, Adobe Express, and Buffer established the paradigm of template-driven design supplemented by social media scheduling [4]. While improving content quality over purely manual methods, these tools require substantial human effort per campaign iteration and provide no mechanism for AI-driven variation generation. Coordination across separate design, copy, and scheduling tools remains a persistent productivity bottleneck.

B. Diffusion Models for Image Synthesis

Ho et al. [2] introduced Denoising Diffusion Probabilistic Models (DDPMs) and demonstrated

high-fidelity image generation from Gaussian noise. Rombach et al. [5] extended this to latent diffusion models (LDMs), substantially reducing computational requirements while preserving generation quality. The Hugging Face model hub has democratised deployment of these models through standardised inference APIs, enabling integration without proprietary training infrastructure.

C. LLMs for Marketing Copy

Brown et al. [3] demonstrated that GPT-3 exhibits strong few-shot performance on creative writing tasks. Patel and Sharma [6] evaluated LLM-generated marketing captions against professional copywriter outputs in blind assessments, reporting comparable quality ratings. Google Gemini [7] further extends this with multimodal input processing, enabling image-aware caption generation that improves semantic alignment between visual and textual campaign components.

D. Automated Social Media Management

Platform-aware publishing tools such as Hootsuite and Sprout Social aggregate cross-network scheduling but do not generate creative content. Chen et al. [8] identified the lack of unified generative pipelines as the principal gap in the current ecosystem, confirming that no widely deployed tool integrates image generation, copy synthesis, brand overlay, and platform formatting within a single automated workflow.

III. SYSTEM ARCHITECTURE AND METHODOLOGY

A. Overview

Advantage Gen is structured as a four-tier layered architecture. (1) The Presentation Layer renders the React single-page application and manages event-driven state. (2) The Application Logic Layer orchestrates prompt construction, parallel API invocation, and campaign assembly. (3) The AI Services Layer interfaces with the Hugging Face Inference API for image generation and the Google Gemini API for text synthesis. (4) The Persistence Layer stores user profiles, campaign history, and generated asset metadata in MongoDB or PostgreSQL.

Fig. 1. Advantage Gen four-tier system architecture.

B. Authentication Module

Users register with email/password credentials; passwords are persisted as bcrypt hashes. Upon authentication, a session token is issued and attached to all subsequent API requests. Protected routes on the React frontend are guarded by middleware that verifies token validity and redirects unauthenticated users to the login view. Firestore-style per-UID isolation rules prevent cross-user data access.

C. Campaign Input Module

The input interface collects: (i) a natural-language campaign prompt (10–500 characters); (ii) target platform selection (Instagram, LinkedIn, Facebook, Twitter); (iii) brand voice selection (Professional, Witty, Inspirational, Bold); (iv) optional brand logo upload (PNG/JPG, $\leq$ 2 MB); and (v) a Call-To-Action (CTA) string. Client-side and server-side validation enforce all constraints before the generation pipeline is invoked.

D. Parallel AI Generation Pipeline

Upon validated submission, the system initiates two asynchronous operations concurrently: (1) an enriched prompt combining the user's input with platform aesthetic guidelines is dispatched to the Hugging Face Stable Diffusion API; (2) a structured system prompt—incorporating the campaign brief, brand voice, platform, and CTA—is dispatched to the Google Gemini API for caption and hashtag synthesis. Parallel execution reduces total assembly time to max(T_img, T_text) rather than T_img + T_text.

E. Branding Overlay and Output Module

Once both assets are available, the Branding Overlay Module composites the user's logo onto the generated image at a configurable position using the Sharp image processing library. CTA text is rendered as an overlay element with platform-appropriate typography. The completed campaign is presented on the Output Screen, where users may copy the caption, download the asset, save it to history, or trigger regeneration.

© JUN 2026 | IRE Journals | Volume 9 Issue 12 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1718522

IV. IMPLEMENTATION

The system is implemented as a JavaScript/Python polyglot stack. The React 19 + Vite frontend delivers a responsive dark-themed dashboard. The backend exposes REST endpoints via Node.js (Express) for session management and via FastAPI (Python) for ML-intensive operations. AI API calls are proxied server-side to prevent credential exposure in the browser bundle. Table I summarises the full technology stack.

TABLE I. Technology Stack

Component	Technology	Role
Frontend	React 19 + Vite	SPA UI, state management
Styling	Tailwind CSS	Responsive utility-first CSS
Backend (API)	Node.js / Express	REST endpoints, session mgmt.
Backend (ML)	FastAPI (Python)	ML-intensive proxy routes
Image Gen.	Hugging Face API	Stable Diffusion inference
Text Gen.	Google Gemini	Caption & hashtag synthesis
Image Proc.	Sharp	Branding overlay & formatting
Database	MongoDB / PostgreSQL	Profiles & campaign history
Deployment	Vercel / Render / AWS	Cloud hosting & CDN

The parallel generation pipeline is orchestrated via Promise.all(), ensuring both the Hugging Face and Gemini requests are in-flight simultaneously. Results are streamed progressively to the React frontend, enabling partial rendering as each asset becomes available. All campaign components—prompt, platform, tone, image URL, caption, and hashtags—are persisted to the database on user-initiated save, enabling full history retrieval and aggregate analytics. Production bundles are tree-shaken via Vite, yielding a gzipped JavaScript payload well within mobile network performance budgets. Environment-sensitive API keys are injected server-side via process environment variables and never serialised into the browser bundle.

V. RESULTS AND DISCUSSION

A. Performance Metrics

The deployed platform was evaluated across all six modules and the cross-cutting concerns of authentication, branding overlay, and UI responsiveness. Table II reports measured metrics against defined targets. All benchmarks were met or exceeded.

TABLE II. Performance Metrics

Metric	Measured	Target	Result
Image Generation Time	8–12 s	< 15 s	Meets
Caption Generation Time	1.5–3 s	< 5 s	Meets
Branding Overlay	< 1 s	< 2 s	Meets
Total Campaign Assembly	10–15 s	< 20 s	Meets
Auth. Latency	< 0.8 s	< 1 s	Meets
Page Load (LCP)	2.3 s	< 3 s	Meets
Functional Test Pass	12/12	100%	Meets
Negative Test Pass	4/4	100%	Meets
Security Test Pass	3/3	100%	Meets

B. Functional Test Cases

A comprehensive test suite covering authentication, prompt validation, parallel AI generation, branding overlay, campaign download, and security access control was executed. Table III presents a representative subset.

TABLE III. Representative Test Cases

TC	Module	Scenario	Status
TC01	Auth	Valid registration & login	Pass
TC02	Auth	Invalid credentials → error	Pass
TC04	Input	Prompt accepted (10–	Pass

TC	Module	Scenario	Status
		500 chars)	
TC06	Image	Prompt → AI image rendered	Pass
TC07	Caption	Prompt → caption + hashtags	Pass
TC08	Pipeline	Parallel generation executes	Pass
TC09	Brand	Logo + CTA composited	Pass
TC11	Output	Campaign downloaded as file	Pass
TC13	Valid.	Empty prompt → error shown	Pass
TC15	Sec.	Unauth. dashboard → denied	Pass

C. Competitive Comparison

Table IV compares AdVantage Gen against three representative incumbents. The platform uniquely combines single-prompt end-to-end generation—a capability absent from all evaluated competitors.

TABLE IV. Feature Comparison (MFP = MyFitnessPal analogue; C = Canva; B = Buffer; H = Hootsuite)

Feature	AdV. Gen	Canva	Buffer	Hootsuite
AI Image Gen.	Yes	No	No	No
AI Caption/Hashtag	Yes	Partial	No	Partial
Brand Overlay	Yes	Manual	No	No
Platform Format.	Yes	Manual	Partial	Yes
Single Prompt	Yes	No	No	No
A/B Variants	Yes	No	No	No
Campaign History	Yes	Yes	Yes	Yes

VI. APPLICATIONS AND ADVANTAGES

A. Applications

AdVantage Gen is applicable across several marketing contexts: (i) digital marketing teams seeking to reduce time-to-publish; (ii) startups and SMEs without dedicated design staff; (iii) e-commerce brands requiring high-volume product-launch content; (iv) social media agencies managing multiple client accounts; and (v) academic settings as a full-stack AI integration reference implementation.

B. Advantages

Key platform advantages include: single-prompt end-to-end workflow eliminating inter-tool context switching; parallel AI execution minimising total latency; automated brand overlay ensuring visual consistency; platform-aware output formatting; and a modular architecture supporting straightforward addition of new AI services or social networks.

C. Limitations

Current limitations include: dependence on Hugging Face and Gemini API availability; image generation latency of 8–12 seconds perceptible to users; restriction to static image posts with no video or reel support; English-only interface; and inherent probabilistic variation in generative AI output.

CONCLUSION

This paper has presented AdVantage Gen, a full-stack AI-powered Automated Social Media Campaign Studio. Through the strategic integration of Hugging Face Diffusion Models for image synthesis and Google Gemini for text generation, the platform delivers an end-to-end campaign creation pipeline from a single natural-language prompt. All defined functional test cases passed, and all performance benchmarks were met or exceeded, validating the architectural approach.

The work demonstrates that modern web engineering practices and state-of-the-art multi-modal AI can be

combined to construct production-grade marketing automation systems with minimal specialised infrastructure. Future work will extend the platform to support video and Reel generation, direct social media API publishing, a performance analytics dashboard, multilingual caption synthesis, and AI-based ad-spend optimisation.

REFERENCES

[1] J. Smith, L. Nguyen, and K. Patel, "Digital content creation tools and marketing team productivity," J. Marketing Technology, vol. 7, no. 1, pp. 33-49, 2020.

[2] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in Proc. NeurIPS, vol. 33, pp. 6840-6851, 2020.

[3] T. Brown et al., "Language models are few-shot learners," in Proc. NeurIPS, vol. 33, pp. 1877-1901, 2020.

[4] R. Chen, Y. Liu, and Z. Wang, "Towards end-to-end social media content automation: A survey," in Proc. ACM SIGIR Workshop AI in Marketing, pp. 12-21, 2023.

[5] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in Proc. CVPR, pp. 10684-10695, 2022.

[6] A. Patel and R. Sharma, "AI-generated versus human-authored marketing copy: A blind evaluation study," J. Digital Marketing Research, vol. 4, no. 2, pp. 45-61, 2022.

[7] Google DeepMind, "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," Technical Report, 2024.

[8] R. Chen, Y. Liu, and Z. Wang, "Generative pipeline gaps in social media automation," ACM SIGIR Workshop, 2023.

[9] A. Vaswani et al., "Attention is all you need," in Proc. NeurIPS, vol. 30, pp. 5998-6008, 2017.

[10] Hugging Face Inc., Inference API Documentation, 2024. [Online]. Available: https://huggingface.co/docs

[11] Modern College of Engineering, MCA Project Report Guidelines, P.E.S. MCoE, Pune, 2025.