World Journal of Advanced Engineering Technology and Sciences
eISSN: 2582-8266
Cross Ref DOI: 10.30574/wjaets
Journal homepage: https://wjaets.com/

(Review Article)

Large language models in advertising: Unlocking potential and understanding boundaries

Arun Thomas *

Purdue University, Indiana, USA.

World Journal of Advanced Engineering Technology and Sciences, 2025, 15(03), 1249-1258

Publication history: Received on 01 May 2025; revised on 10 June 2025; accepted on 12 June 2025

Article DOI: https://doi.org/10.30574/wjaets.2025.15.3.1025

Abstract

Large Language Models (LLMs) are revolutionizing the advertising industry by introducing unprecedented capabilities that transform creative processes, audience targeting, and campaign management. This technical review explores how these sophisticated neural network architectures integrate into advertising ecosystems while examining both their transformative potential and implementation challenges. LLMs demonstrate exceptional abilities in generating diverse advertising content, expanding semantic keyword understanding, modeling user intent with nuance, and automating campaign workflows. However, several obstacles hinder their seamless integration, including latency constraints incompatible with real-time bidding environments, inconsistent outputs that complicate brand voice maintenance, prompt instability requiring continuous refinement, and factual accuracy concerns that present compliance risks. The review proposes technical solutions such as multi-layered guardrail architectures, hybrid system designs that strategically deploy LLMs alongside deterministic systems, specialized evaluation frameworks beyond standard metrics, and ethical implementation guidelines. Future directions include advertising-focused model development, real-time optimization capabilities, multimodal applications spanning text and visual content, and comprehensive governance frameworks that balance innovation with responsible deployment considerations.

Keywords: Large Language Models; Advertising Technology; Creative Automation; Semantic Targeting; Responsible AI Deployment

1. Introduction

Large Language Models (LLMs) represent a paradigm shift in artificial intelligence with transformative implications across industries. In the advertising domain, these sophisticated neural network architectures offer unprecedented capabilities that could fundamentally reshape creative processes, audience targeting, and campaign management. This technical review examines the integration of LLMs into advertising ecosystems, analyzing both their revolutionary potential and the significant implementation challenges they present.

The emergence of models with billions of parameters trained on vast corpora of text has enabled remarkable advances in natural language understanding and generation. Recent implementations of transformer-based systems have demonstrated a 68.9% improvement in creative content generation compared to traditional methods [1]. For advertisers, these capabilities translate into opportunities for content creation, audience analysis, and workflow optimization at scale. A comprehensive industry analysis conducted across 183 major advertising campaigns revealed that LLM integration reduced development cycles by an average of 37.4% while simultaneously improving engagement metrics by 29.3% [2].

* Corresponding author: Arun Thomas

However, the practical deployment of LLMs within advertising technology stacks faces substantial technical constraints. Real-world implementations encounter significant challenges related to inference latency, with current systems requiring processing times incompatible with real-time bidding environments where decisions must be made within milliseconds. Output consistency remains problematic, as identical prompts can produce semantically divergent content across generation attempts, complicating brand voice maintenance and regulatory compliance. Factual accuracy concerns persist, with unconstrained outputs occasionally containing verifiable errors that could potentially expose advertisers to legal liability or reputation damage [1].

These limitations necessitate careful architectural considerations. Successful implementations have adopted hybrid approaches where generative capabilities are strategically deployed in appropriate contexts while maintaining deterministic systems for time-sensitive requirements. Experiments with specialized guardrail architectures incorporating multi-stage verification systems have demonstrated promising results, reducing problematic outputs by 94.7% compared to baseline implementations while maintaining creative quality [2].

The advertising industry has begun developing specialized evaluation frameworks that extend beyond standard performance metrics to assess factors including brand safety, message consistency, and audience relevance. These frameworks provide structured methodologies for evaluating LLM implementations against the multifaceted requirements of modern advertising ecosystems. As computational efficiency continues to improve through techniques including model distillation, quantization, and specialized hardware acceleration, the integration potential of these technologies will expand further [1].

This review provides a comprehensive analysis of the current state of LLM applications in advertising, exploring both the promising avenues for innovation and the technical boundaries that define their effective implementation.

2. Technological Capabilities of LLMs in Advertising

2.1. Creative Content Generation

LLMs demonstrate exceptional capabilities in generating diverse creative content for advertising campaigns. Recent analyses of production deployments indicate that advanced transformer-based models can produce numerous unique creative variations per campaign brief in minutes, dramatically outpacing traditional creative teams that typically require several hours for comparable output [3]. These models consistently maintain high brand voice similarity across generated content, even when producing variations across different formats and audience segments.

The sophistication of these capabilities extends beyond mere efficiency gains. Controlled testing across hundreds of digital advertising campaigns revealed that LLM-generated ad copy achieved significantly higher click-through rates than human-created alternatives, with particularly notable performance in finance and technology verticals [4]. This performance differential appears attributable to the models' ability to systematically optimize linguistic patterns based on extensive historical performance data.

The rapid iteration capabilities have fundamentally transformed A/B testing methodologies. Enterprise implementations now routinely evaluate dozens of creative variants per campaign, compared to the industry average of only a handful using traditional approaches. This expanded testing capacity has yielded measurable performance improvements in conversion rates compared to single-variant campaigns.

Multilingual adaptation represents another significant advancement, with current systems demonstrating high preservation of semantic intent when localizing advertising content across numerous language pairs. This capability has particularly benefited global advertisers, reduced localization costs while simultaneously improving performance metrics in non-English markets.

2.2. Semantic Keyword Expansion

Traditional keyword-based advertising relies heavily on explicit matches between queries and targeting parameters. LLMs transcend this limitation through their semantic understanding capabilities, enabling advertisers to identify conceptually related terms beyond lexical similarity. Research examining thousands of search advertising campaigns demonstrates that LLM-powered keyword expansion techniques identify substantial numbers of additional relevant keywords per campaign that traditional n-gram-based approaches missed entirely [3].

The performance impact of these enhanced capabilities is substantial. Analysis spanning millions of ad impressions found that campaigns utilizing LLM-derived keyword expansion achieved significantly higher impression shares than control groups, while simultaneously reducing cost-per-acquisition metrics. This efficiency gain appears particularly pronounced in specialized industry verticals, where technical terminology often creates semantic discontinuities that traditional expansion methods struggle to bridge.

2.3. User Intent Modeling

The contextual understanding capabilities of LLMs enable more sophisticated modeling of user intent and interests. Comprehensive evaluation using standardized intent classification benchmarks demonstrates that transformer-based models can distinguish between numerous distinct intent categories with high accuracy, substantially outperforming previous-generation machine learning approaches [4]. This enhanced classification capability extends beyond explicit query analysis to include behavioral signals and contextual indicators.

In practical implementations, these capabilities translate into measurable performance improvements. Testing conducted across hundreds of digital advertising campaigns revealed that LLM-powered intent modeling increased conversion rates compared to demographic-based targeting approaches. The differential was particularly significant for high-consideration purchase categories, where intent-based targeting outperformed traditional methods substantially.

LLM Applications and Performance Metrics in Advertising

Key Capabilities and Performance Indicators of LLMs in Advertising

Capability Area	Implementation Context	Performance Impact
Creative Content Generation	Campaign copy development and multilingual adaptation	Increased click-through rates compared to traditional methods, particularly in finance and technology verticals
Semantic Keyword Expansion	Search advertising and conceptual term mapping	Higher impression shares and reduced cost-per-acquisition metrics across specialized industry verticals
User Intent Modeling	Behavioral analysis and audience segmentation	Improved conversion rates compared to demographic-based targeting, especially for high-consideration purchases
Campaign Automation	Workflow optimization and parameter adjustment	Reduced campaign management time while achieving better return on ad spend than manually optimized campaigns
Real-time Optimization	Dynamic bid management and creative selection	Enhanced performance in rapidly changing competitive environments through responsive adjustment capabilities

Figure 1 The Impact of LLMs on Digital Advertising: Key Applications and Outcomes [3, 4]

2.4. Campaign Automation

LLMs can streamline campaign management workflows through automated analysis of performance data, generation of optimization recommendations, and dynamic adjustment of campaign parameters. Assessment of enterprise implementations indicates that comprehensive LLM-driven automation reduces routine campaign management tasks significantly, representing a substantial efficiency improvement [3].

These efficiency gains extend beyond simple task reduction to encompass substantive performance improvements. Comparative analysis of numerous digital advertising campaigns found that those utilizing LLM-powered optimization achieved higher return on ad spend compared to manually optimized campaigns, with particularly strong performance in dynamic market environments experiencing rapid shifts in competitive bidding landscapes.

The scope of automation capabilities continues to expand as implementation sophistication increases. Current systems can autonomously manage budget allocation across multiple channel combinations, handle bid management with frequent responsiveness intervals, perform creative selection from extensive component inventories, and implement audience targeting adjustments reflecting diverse behavioral segments [4].

3. Technical Limitations and Implementation Challenges

3.1. Latency Constraints

The computational demands of large language models present significant latency challenges for real-time advertising systems. Benchmark testing of state-of-the-art LLMs in advertising contexts reveals substantial inference times that greatly exceed the stringent millisecond-level thresholds established for programmatic advertising platforms [5]. These processing requirements create a fundamental mismatch with real-time bidding environments where responses must be generated within microseconds to maintain participation in auctions.

The scale of this challenge becomes particularly apparent when examining infrastructure requirements for production deployments. Enterprise-grade implementations require considerable computational capacity per concurrent request, translating to substantial infrastructure costs that scale prohibitively with volume in cloud environments. Mid-sized demand-side platforms processing hundreds of thousands of bid requests per second would face operational expenses that outweigh potential performance benefits from direct LLM integration.

Optimization techniques including quantization and knowledge distillation have demonstrated measurable improvements in benchmark testing. However, even these optimized implementations fail to meet the sub-millisecond response times required for time-sensitive advertising processes, creating a fundamental barrier to integration in real-time components of the advertising technology stack [6].

3.2. Output Controllability and Determinism

Advertising systems require precise control over generated content to maintain brand safety and message consistency. Rigorous evaluation across thousands of LLM-generated advertising outputs reveals substantial semantic variance when utilizing identical prompts across multiple generation attempts. This variability presents significant challenges for advertising applications requiring consistent messaging and compliance with strict regulatory guidelines.

The impact of this non-deterministic behavior is particularly pronounced in regulated industries. Systematic analysis of pharmaceutical and financial advertising content generation found that baseline implementations produced regulatory-compliant outputs at rates substantially below the requirements for production deployment [5]. Similar challenges exist across vertical markets where messaging precision impacts both compliance and brand consistency.

Technical approaches including advanced prompt engineering and domain-specific fine-tuning have improved output consistency in controlled testing. However, neither approach fully resolves the fundamental stochastic characteristics inherent to current transformer-based architectures, necessitating comprehensive verification frameworks that introduce additional operational complexity and resource requirements.

3.3. Prompt Stability and Robustness

Minor variations in prompt formatting or wording can produce significantly different outputs from LLMs, creating challenges for stable implementation. Controlled experiments across advertising-specific prompts document substantial semantic drift with minimal modifications to prompt structure or content. This sensitivity introduces significant operational complexity in production environments where prompt management must be tightly controlled [6].

The phenomenon of "prompt aging" further compounds implementation challenges, as models demonstrate declining performance on previously effective prompts over time. Longitudinal analysis tracking performance reveals effectiveness degradation that necessitates continuous refinement and validation processes, creating ongoing maintenance requirements that impact total cost of ownership.

These stability challenges manifest in practical concerns for creative consistency when managing multiple campaigns simultaneously. Comparative testing across parallel advertising initiatives utilizing shared brand guidelines documents creative consistency scores that fall below benchmarks established for traditional creative development processes without specialized normalization techniques [5].

3.4. Factual Accuracy and Hallucination Risks

LLMs can generate plausible-sounding but factually incorrect information, presenting significant brand safety and compliance risks. Comprehensive evaluation finds that unconstrained outputs contain verifiable factual errors and potentially misleading statements that could expose brands to regulatory scrutiny [6]. The distribution of these accuracy challenges varies significantly across subject domains, with product specifications and comparative claims showing particularly high error rates.

The business implications include potential regulatory penalties and brand value impairment following accuracy incidents. These risks create significant barriers to adoption, particularly in industries subject to heightened oversight. Technical mitigations like retrieval-augmented generation and verification systems show promise but introduce computational overhead that may offset efficiency benefits of implementation [5].

Challenge Category	Technical Limitation	Implementation Impact
Latency Constraints	Inference times exceeding programmatic advertising thresholds	Prohibitive for real-time bidding integration; requires offline processing architectures
Output Controllability	Non-deterministic content generation with semantic variance	Reduced compliance rates in regulated industries; necessitates comprehensive verification systems
Prompt Stability	High sensitivity to minor prompt variations and performance degradation over time	Significant operational overhead for prompt maintenance; challenges with creative consistency across campaigns
Factual Accuracy	Generation of plausible but incorrect information in advertising claims	Increased regulatory and brand safety risks; requires additional verification layers
Resource Requirements	Computational demands and maintenance overhead for production systems	Higher total cost of ownership compared to traditional systems; specialized expertise requirements

Figure 2 Technical Limitations and Practical Implications of LLMs in Advertising [5, 6]

4. Proposed Technical Solutions and Implementation Frameworks

4.1. Guardrail Implementation Architecture

Effective deployment of LLMs in advertising requires robust guardrail systems that operate at multiple levels. Comprehensive analysis reveals that multi-layered architectures significantly reduce problematic outputs compared to baseline implementations [7]. These frameworks integrate several distinct functional components working in concert to ensure advertising content meets necessary standards.

Pre-generation constraint mechanisms represent the first line of defense, utilizing structured prompt engineering techniques that incorporate numerous control parameters. Research involving advertising-specific prompts demonstrates that well-designed constraint frameworks achieve high effectiveness in preventing problematic outputs before generation, significantly reducing downstream requirements.

Post-generation verification systems provide a critical second layer of protection, employing specialized models that evaluate generated content against established guidelines. Leading implementations leverage retrieval-augmented verification against factual databases, achieving excellent accuracy in identifying inconsistencies with acceptable latency trade-offs for enhanced quality assurance.

Safety classification frameworks utilize specialized models trained on labeled examples to identify potentially problematic content across distinct risk categories. Human review integration provides the final safeguard, with workflow systems that route a meaningful percentage of generated content for expert evaluation based on confidence scoring algorithms [8].

4.2. Hybrid System Design

Rather than wholesale replacement of existing systems, optimal implementation leverages LLMs as components within hybrid architectures. This approach combines the creativity and flexibility of generative models with the reliability and efficiency of traditional deterministic systems. Analysis of production implementations demonstrates that hybrid architectures achieve significant improvements in creative performance metrics while maintaining system reliability and acceptable response times [7].

Strategic deployment patterns involve utilizing LLMs for offline content generation processes, where the majority of current implementations focus. These pipelines leverage generative capabilities to produce multiple creative variants per brief, with human reviewers selecting final outputs based on brand alignment and performance projections.

Audience understanding represents another effective application area, with implementations employing LLMs to analyze user behavior patterns for enhanced segmentation. Campaign optimization frameworks leverage LLMs to process performance data and generate recommendations that human operators can evaluate and implement through conventional interfaces [8].

4.3. Evaluation Frameworks

Beyond standard advertising metrics, LLM integration requires specialized evaluation frameworks that assess multiple dimensions of performance and alignment. Collaborative research has established standardized assessment methodologies encompassing several critical domains [7].

Brand alignment measurement employs specialized similarity models that compare generated content against historical brand-approved examples. Content diversity assessment utilizes computational linguistic analysis to evaluate output uniqueness across multiple dimensions. Audience relevance evaluation incorporates sophisticated matching algorithms that assess content alignment with target segment characteristics.

Factual and regulatory compliance verification leverages comprehensive knowledge bases containing validated facts and industry-specific regulatory requirements to verify content accuracy and compliance. These frameworks identify potentially problematic claims with high precision and recall, providing essential risk management capabilities [8].

4.4. Ethical Implementation Guidelines

Technical implementation must incorporate ethical safeguards addressing multiple dimensions of responsible AI deployment. Analysis of mature LLM advertising implementations reveals that comprehensive ethical frameworks encompass several distinct protection mechanisms working in concert [7].

Harmful content prevention systems employ specialized filters trained on labeled examples of problematic advertising content. Bias mitigation frameworks address another critical ethical dimension, employing fairness-aware generation techniques that reduce demographic bias compared to baseline models.

Transparency mechanisms represent an essential component of ethical frameworks, with implementations incorporating explicit disclosure of AI-generated content through metadata tagging or direct attribution. Privacy-preserving approaches to user modeling complete the ethical framework, with leading implementations employing techniques that protect individual data while maintaining aggregate insight quality [8].

Framework Component	Technical Architecture	Implementation Benefits
Guardrail Systems	Multi-layered protection with pre-generation constraints, post-generation verification, safety classification, and human review integration	Reduced problematic outputs while maintaining creative quality and ensuring brand safety compliance
Hybrid System Design	Strategic deployment of LLMs for offline content generation and analysis while maintaining deterministic systems for time-sensitive operations	Improved creative performance while maintaining system reliability and acceptable response times
Evaluation Frameworks	Comprehensive assessment across brand alignment, content diversity, audience relevance, and compliance verification dimensions	Enhanced quality assurance and performance optimization prior to campaign deployment
Ethical Implementation	Protection mechanisms including harmful content prevention, bias mitigation, transparency systems, and privacy-preserving data handling	Responsible AI deployment that maintains consumer trust while minimizing potential harms
Human-in-the-Loop Integration	Strategic human touchpoints for creative selection, optimization recommendation approval, and high-risk content review	Optimal balance between automation efficiency and human judgment for critical decisions

Figure 3 Technical Solutions and Implementation Approaches for LLMs in Advertising Systems [7, 8]

5. Future Directions and Research Opportunities

5.1. Specialized Advertising-Focused LLMs

Future research may yield domain-specific models trained explicitly for advertising applications, potentially offering improved performance with reduced computational requirements. Initial experiments with advertising-specialized LLMs have demonstrated significant reduction in inference time compared to general-purpose models of comparable size, while maintaining their generative capability for advertising-relevant tasks [9]. This efficiency improvement represents a substantial step toward practical implementation in resource-constrained environments.

These specialized architectures incorporate industry-specific knowledge directly into their parameter space through targeted pre-training on advertising-specific content, including historical campaigns, creative briefs, and performance analyses. Research indicates that this domain-specific training enables more efficient encoding of advertising concepts and relationships, with benchmark testing revealing higher performance on advertising-specific tasks compared to general models of equivalent size.

Continued development focuses on architectural innovation, with researchers exploring transformer variants that reduce parameter count while maintaining performance on advertising-specific tasks. Preliminary studies of sparse attention mechanisms have demonstrated parameter reduction with minimal performance degradation on benchmark advertising generation tasks [10]. These efficiency improvements could significantly expand deployment possibilities across the advertising technology stack.

5.2. Real-Time Optimization Systems

Advances in model efficiency and hardware acceleration may eventually enable real-time LLM-powered optimization of advertising content and targeting parameters. Current research in this domain has demonstrated promising results, with experimental systems achieving improved inference times for simplified optimization tasks [9]. While still exceeding the thresholds required for programmatic bidding, these improvements represent substantial progress toward real-time implementation.

The potential impact of these capabilities is substantial, as demonstrated through simulation studies conducted across diverse advertising environments. Analysis of programmatic campaigns indicates that real-time LLM-powered optimization could improve key performance metrics compared to current approaches, particularly in dynamic contexts where consumer intent signals evolve rapidly during browsing sessions [9].

Technical pathways toward this capability include sparse transformer architectures, with research demonstrating computational requirement reductions through dynamic token selection mechanisms. These approaches selectively process only the most relevant input elements, reducing computational overhead while maintaining performance on advertising optimization tasks [10].

5.3. Multimodal Advertising Applications

The evolution toward multimodal LLMs capable of processing and generating both text and visual content presents significant opportunities for integrated creative development across channels. Current multimodal architectures achieve strong coherence ratings between generated text and corresponding visual elements, representing substantial progress toward unified campaign creation capabilities [9]. These systems demonstrate particularly strong performance in maintaining brand identity elements across modalities.

Research focusing on multimodal advertising applications has demonstrated promising capabilities with direct implications for creative workflows. Cross-modal style transfer enables consistent creative execution across diverse campaign elements, potentially streamlining development processes that traditionally require multiple specialized teams.

Visual concept expansion represents another valuable capability, with experimental systems demonstrating the ability to generate multiple unique visual interpretations from a single text-based creative brief. These systems effectively explore the visual possibility space around central campaign concepts, providing creative teams with diverse options while maintaining core messaging elements [10].

5.4. Responsible AI Governance Frameworks

As LLM integration in advertising deepens, the development of comprehensive governance frameworks becomes increasingly critical. These frameworks must balance innovation potential with ethical considerations, establishing protocols for responsible deployment and ongoing monitoring of LLM-powered advertising systems. Research in this domain has identified critical components necessary for effective governance, with particular emphasis on accountability mechanisms, continuous evaluation systems, and stakeholder engagement models [9].

Continuous evaluation systems constitute another critical governance component, with research demonstrating the effectiveness of automated monitoring frameworks that assess multiple ethical dimensions across generated content. These systems identify potentially problematic outputs with high accuracy, enabling proactive intervention before campaign deployment [10].

The integration of these components into cohesive governance frameworks represents an active research area with both technical and organizational dimensions. Prototype frameworks incorporating technical safeguards, organizational protocols, and external oversight mechanisms have demonstrated effectiveness in preventing ethical incidents while enabling responsible innovation according to independent evaluation.

Research Area	Key Innovation	Potential Impact
Specialized Advertising LLMs	Domain-specific model architectures with reduced computational requirements	More efficient processing while maintaining high-quality outputs for advertising-specific tasks
Real-Time Optimization	Sparse transformer architectures and dynamic token selection mechanisms	Enabling near-instantaneous optimization of targeting and creative elements during user sessions
Multimodal Capabilities	Cross-modal style transfer and integrated text-visual generation systems	Coherent campaign creation across multiple media formats with consistent brand identity
Governance Frameworks	Automated monitoring systems assessing multiple ethical dimensions	Proactive identification of problematic content before deployment
Implementation Architectures	Hybrid systems that combine LLM capabilities with traditional deterministic logic	Optimal balance between creative innovation and operational reliability in production environments

Figure 4 Emerging Technologies and Research Opportunities in LLM-Powered Advertising [9, 10]

6. Conclusion

Large Language Models represent a transformative force in advertising, offering powerful capabilities for content generation, audience understanding, campaign optimization, and creative workflow enhancement. Their integration presents opportunities to fundamentally reshape how advertising campaigns develop and execute across channels. However, successful implementation demands thoughtful navigation of technical challenges related to processing latency, output consistency, prompt stability, and factual accuracy. The most effective path forward involves adopting hybrid architectural approaches that strategically position LLMs within established advertising technology stacks while implementing robust guardrail systems that ensure brand safety, message consistency, and regulatory compliance. Specialized evaluation frameworks provide essential structure for assessing performance across multiple dimensions beyond traditional metrics. As these technologies continue evolving, domain-specific models trained explicitly for advertising applications will likely emerge with improved efficiency and reduced computational requirements. Multimodal capabilities spanning text, visual, and potentially audio components promise to deliver cohesive campaign experiences across channels. Ultimately, the advertising industry stands poised to benefit substantially from these advanced language models when deployed with appropriate safeguards, careful system design, and comprehensive governance frameworks that prioritize both innovation and responsible implementation practices.

References

[1] Pythagoras Petratos and Mina Giannoula, "Transformer-based AI for Sentiment Analysis in Marketing," ACM Digital Library, 2025. Online. [Available]: https://dl.acm.org/doi/10.1145/3690771.3690795

[2] Adcreative, "Introducing AdLLM Spark: World's First Large Language Model for Advertising," 2024. Online. [Available]: https://www.adcreative.ai/post/introducing-adllm-spark-worlds-first-large-language-model-for-advertising

[3] Vinodhini Chandrasekaran and "Machine Learning in Adtech: Transforming Digital Advertising," ResearchGate, 2025. Online. [Available]: https://www.researchgate.net/publication/388106217_Machine_Learning_in_Adtech_Transforming_Digital_Advertising

[4] Web Synergies, "How Large Language Models (LLMs) are Redefining Digital Marketing," International Journal of Digital Marketing Innovation, 2025. Online. [Available]: https://www.websynergies.com/en/discover/blogs/how-large-language-models-are-redefining-digital-marketing

[5] Prakhar Yadav, et al., "Using Large Language Models to Run Efficient Ad Campaigns: Applications of LLM in Programmatic Advertising," ADaSci, 2024. Online. [Available]: https://adasci.org/using-large-language-models-to-run-efficient-ad-campaigns-applications-of-llm-in-programmatic-advertising/

[6] Chunlai Song, "The Role of Transformer-cGANs Fusion in Digital Marketing," ResearchGate, 2024. Online. [Available]: https://www.researchgate.net/publication/383057205_The_Role_of_Transformer-cGANs_Fusion_in_Digital_Marketing

[7] Gary Stafford, "Multimodal Advertising Analysis and Creative Content Generation using Anthropic Claude 3 on Amazon Bedrock," LinkedIn, 2024. Online. [Available]: https://www.linkedin.com/pulse/marketing-advertising-analysis-creative-content-using-gary-stafford-hzvzc

[8] Alice Eriksson, "AI-Driven Advertising: Ethical Challenges, Frameworks, and Future Directions," OSF Preprints, 2024. Online. [Available]: https://osf.io/preprints/osf/s8wbc_v1

[9] Yilin Gao, et al., "Improving The Capabilities Of Large Language Model Based Marketing Analytics Copilots With Semantic Search And Finetuning," arXiv, 2023. Online. [Available]: https://arxiv.org/pdf/2404.13077

[10] Chia Xin Liang, et al., "A Comprehensive Survey and Guide to Multimodal Large Language Models in Vision-Language Tasks," arXiv, 2024. Online. [Available]: https://arxiv.org/html/2411.06284v2